

HapticHead

Augmenting Reality via Tactile Cues

Von der Fakultät für Elektrotechnik und Informatik
der Gottfried Wilhelm Leibniz Universität Hannover

zur Erlangung des akademischen Grades

Doktor rerum naturalium
(abgekürzt: Dr. rer. nat.)

genehmigte Dissertation

von Herrn

M. Sc.

Oliver Beren Kaul

geboren am 8. Oktober 1987
in Gehrden, Deutschland

2021

Oliver Beren Kaul

HapticHead - Augmenting Reality via Tactile Cues

1. Referent:

Prof. Dr. Michael Rohs

Gottfried Wilhelm Leibniz Universität Hannover

2. Referent:

Prof. Dr. Albrecht Schmidt

Ludwig-Maximilians-Universität München

Tag der Promotion: 28. Mai 2021

Abstract

Information overload is increasingly becoming a challenge in today's world. Humans have only a limited amount of attention to allocate between sensory channels and tend to miss or misjudge critical sensory information when multiple activities are going on at the same time. For example, people may miss the sound of an approaching car when walking across the street while looking at their smartphones. Some sensory channels may also be impaired due to congenital or acquired conditions. Among sensory channels, touch is often experienced as obtrusive, especially when it occurs unexpectedly. Since tactile actuators can simulate touch, targeted tactile stimuli can provide users of virtual reality and augmented reality environments with important information for navigation, guidance, alerts, and notifications.

In this dissertation, a tactile user interface around the head is presented to relieve or replace a potentially impaired visual channel, called *HapticHead*. It is a high-resolution, omnidirectional, vibrotactile display that presents general, 3D directional, and distance information through dynamic tactile patterns. The head is well suited for tactile feedback because it is sensitive to mechanical stimuli and provides a large spherical surface area that enables the display of precise 3D information and allows the user to intuitively rotate the head in the direction of a stimulus based on natural mapping.

Basic research on tactile perception on the head and studies on various use cases of head-based tactile feedback are presented in this thesis. Several investigations and user studies have been conducted on (a) the funneling illusion and localization accuracy of tactile stimuli around the head, (b) the ability of people to discriminate between different tactile patterns on the head, (c) approaches to designing tactile patterns for complex arrays of actuators, (d) increasing the immersion and presence level of virtual reality applications, and (e) assisting people with visual impairments in guidance and micro-navigation.

In summary, tactile feedback around the head was found to be highly valuable as an additional information channel in various application scenarios. Most notable is the navigation of visually impaired individuals through a micro-navigation obstacle course, which is an order of magnitude more accurate than the previous state-of-the-art, which used a tactile belt as a feedback modality. The *HapticHead* tactile user interface's ability to safely navigate people with visual impairments around obstacles and on stairs with a mean deviation from the optimal path of less than 6 cm may ultimately improve the quality of life for many people with visual impairments.

Keywords: Tactile Feedback, Guidance, Accessibility

Zusammenfassung

Die Informationsüberlastung wird in der heutigen Welt zunehmend zu einer Herausforderung. Der Mensch hat nur eine begrenzte Menge an Aufmerksamkeit, die er zwischen den Sinneskanälen aufteilen kann, und neigt dazu, kritische Sinnesinformationen zu verpassen oder falsch einzuschätzen, wenn mehrere Aktivitäten gleichzeitig ablaufen. Zum Beispiel können Menschen das Geräusch eines herannahenden Autos überhören, wenn sie über die Straße gehen und dabei auf ihr Smartphone schauen. Einige Sinneskanäle können auch aufgrund von angeborenen oder erworbenen Erkrankungen beeinträchtigt sein. Unter den Sinneskanälen wird Berührung oft als aufdringlich empfunden, besonders wenn sie unerwartet auftritt. Da taktile Aktoren Berührungen simulieren können, können gezielte taktile Reize den Benutzern von Virtual- und Augmented Reality Anwendungen wichtige Informationen für die Navigation, Führung, Warnungen und Benachrichtigungen liefern.

In dieser Dissertation wird eine taktile Benutzeroberfläche um den Kopf herum präsentiert, um einen möglicherweise beeinträchtigten visuellen Kanal zu entlasten oder zu ersetzen, genannt *HapticHead*. Es handelt sich um ein hochauflösendes, omnidirektionales, vibrotaktil Display, das allgemeine, 3D-Richtungs- und Entfernungsinformationen durch dynamische taktile Muster darstellt. Der Kopf eignet sich gut für taktiles Feedback, da er empfindlich auf mechanische Reize reagiert und eine große sphärische Oberfläche bietet, die die Darstellung präziser 3D-Informationen ermöglicht und es dem Benutzer erlaubt, den Kopf aufgrund der natürlichen Zuordnung intuitiv in die Richtung eines Reizes zu drehen.

Grundlagenforschung zur taktilen Wahrnehmung am Kopf und Studien zu verschiedenen Anwendungsfällen von kopfbasiertem taktilem Feedback werden in dieser Arbeit vorgestellt. Mehrere Untersuchungen und Nutzerstudien wurden durchgeführt zu (a) der Funneling Illusion und der Lokalisierungsgenauigkeit von taktilen Reizen am Kopf, (b) der Fähigkeit von Menschen, zwischen verschiedenen taktilen Mustern am Kopf zu unterscheiden, (c) Ansätzen zur Gestaltung taktiler Muster für komplexe Arrays von Aktoren, (d) der Erhöhung des Immersions- und Präsenzgrades von Virtual-Reality-Anwendungen und (e) der Unterstützung von Menschen mit Sehbehinderungen bei der Führung und Mikronavigation.

Zusammenfassend wurde festgestellt, dass taktiles Feedback um den Kopf herum als zusätzlicher Informationskanal in verschiedenen Anwendungsszenarien sehr wertvoll ist. Am interessantesten ist die Navigation von sehbehinderten Personen durch einen Mikronavigations-Hindernisparcours, welche um eine Größenordnung präziser ist als der bisherige Stand der Technik, der einen taktilen Gürtel als Feedback-Modalität verwendete. Die Fähigkeit der taktilen Benutzerschnittstelle *HapticHead*, Menschen mit Sehbehinderungen mit einer mittleren Abweichung vom optimalen Pfad von weniger als 6 cm sicher um Hindernisse und auf Treppen zu navigieren, kann letztendlich die Lebensqualität vieler Menschen mit Sehbehinderungen verbessern.

Schlagerworte: Taktiles Feedback, Orientierungshilfe, Barrierefreiheit

Acknowledgement

First of all, I would like to thank my Ph.D. supervisor Prof. Dr. Michael Rohs, for allowing me to work on exciting research in the Human-Computer Interaction Group. I highly appreciated the freedom in choosing my research topic and all the fun little side projects I was able to work on during my time as a Ph.D. student. The opportunity to visit so many conferences for research and networking was a fantastic experience that would not have been possible without Michael's support.

Moreover, I appreciate the willingness of Dr. Albrecht Schmidt to be my co-examiner. I fondly remember the winter schools co-organized by Albrecht, which were charming opportunities for networking and discussing new ideas on the latest Human-Computer Interaction research.

I am grateful for all the fantastic students I enjoyed working with while advising their Bachelor's and Master's theses. In particular, I would like to thank Marc Mogalle, Kevin Meier, Andreas Domin, Leonard Hansing, Guido Gardlo, and Jonas Bock for their substantial contributions to my research published in this dissertation. Benjamin Simon and Kamillo Ferry further helped me conducting experiments as research assistants.

I further want to extend my thanks to Professor Hans Gellersen. He provided excellent feedback as an associate editor for the ACM Transactions on Computer-Human Interaction journal, which led to the gratifying publication of our article on micro-navigating people with visual impairments. I am also thankful for all the study participants who took part in the experiments presented in this thesis, particularly for the participants with visual impairments who traveled to our lab independently and gave us helpful feedback on our system.

I am thankful for the fruitful discussions and collaborations with my colleagues at the Human-Computer Interaction Group. I would particularly like to thank Max Pfeiffer for his support and advice in getting me started in research and Tim Duent for his kind personality and his help in designing a much cleaner version of the HapticHead motor driver board. I would like to thank my friends Tamara Bandikova, Abby Wanyu Liu, Alexander Eiselmayer, Marie Laflör, and Jan Balata for being generally awesome people, reviewing parts of this thesis and providing me with helpful suggestions.

My heartfelt thanks to my wonderful girlfriend Tam for her support, positive attitude, and delicious Vietnamese meals. I am excited for the next chapter! Finally, I would like to thank my parents for making my education possible and providing me with continuous support, love, and encouragement.

Contents

1	Introduction and Motivation	1
1.1	Approach and Goals	3
1.2	Contributions	6
2	Foundations of Human Perception and Tactile Feedback	9
2.1	Human Nervous System	9
2.1.1	Nerve Cells	10
2.1.2	Firing of Neurons	11
2.1.3	Perception of External Stimuli	12
2.2	Structure of Human Skin	13
2.3	Tactile Perception	14
2.3.1	Early Research on Tactile Perception	16
2.3.2	Merkel Disks - $A\beta$ SA1 Low-Threshold Mechanoreceptors	16
2.3.3	Ruffini Endings - $A\beta$ SA2 Low-Threshold Mechanoreceptors	17
2.3.4	Meissner Corpuscles - $A\beta$ RA1 Low-Threshold Mechanoreceptors	17
2.3.5	Pacinian Corpuscles - $A\beta$ RA2 Low-Threshold Mechanoreceptors	17
2.3.6	Specialized Afferents in Hairy Skin	18
2.3.7	High-Threshold Mechanoreceptors	19
2.3.8	Tactile Perception around the Head	20
2.4	Existing Approaches to Generate Tactile Stimuli	20
2.4.1	Thermal Stimulation	21
2.4.2	Inadequate Stimulation of Receptors through Electrical Stimuli	21
2.4.3	Adequate Mechanical Stimulation and Actuator Types	22
2.5	Vibrotactile Patterns: Terminology and Characteristics	25
2.6	Displays for Presenting Vibrotactile Stimuli	26
2.6.1	Indicating Direction via Vibrotactile Displays	27
2.6.2	Tactile Displays on the Head	29
2.7	Tactile Illusions	32
2.7.1	The Funneling Illusion	32
2.7.2	The Cutaneous Rabbit Illusion	33
3	Tactile Localization Accuracy and the Funneling Illusion around the Head	35
3.1	Introduction	36
3.1.1	Terminology	37

3.2	Related Work	37
3.2.1	Funneling Illusion through Objects and Out-of-Body	39
3.2.2	Localization and Funneling Illusion on the Head	39
3.3	Tactile Perception Prototype	39
3.4	Tactile Perception Experiment	40
3.4.1	Tactile Perception Experiment – Design	41
3.4.2	Tactile Perception Experiment – Procedure	43
3.5	Tactile Perception Experiment – Results	44
3.5.1	Quantitative Results – Single Actuator Trials	44
3.5.2	Quantitative Results – Multi Actuator Trials	46
3.5.3	Subjective Results	50
3.6	Tactile Perception Experiment – Discussion	51
3.6.1	Single Actuator Localization Performance	51
3.6.2	Heat Map of Localization Accuracies	52
3.6.3	Midline Bias	53
3.6.4	Possible Dominant Hand Effect	53
3.6.5	Occurrence of the Funneling Illusion	54
3.6.6	Occurrence of the Centralizing Bias	54
3.6.7	Head Sensitivities and Subjective Feedback	55
3.7	Limitations	55
3.8	Conclusion and Future Work	56
4	HapticHead – A Tactile User Interface around the Head	57
4.1	Concept	57
4.1.1	Limitations and Anatomical Considerations	58
4.1.2	Tactile Illusion Considerations	60
4.2	Hardware Implementation	61
4.2.1	Initial Main Prototype	61
4.2.2	Second Iteration of the Main Prototype	62
4.2.3	Third Iteration of the Main Prototype	65
4.2.4	Fourth and Final Iteration of the Main Prototype	66
4.3	Other Prototype Versions	66
4.3.1	Powerful, Large Actuator Prototype	66
4.3.2	Low-Frequency Actuator Prototype	67
5	Tactile Patterns around the Head	69
5.1	Introduction and Motivation	69
5.2	Related Work	71
5.3	Distinguishability of Patterns Experiment	72
5.3.1	Vibrotactile Pattern Design	72
5.3.2	Distinguishability of Patterns Experiment – Procedure	74
5.3.3	Distinguishability of Patterns Experiment – Objective Results	75
5.3.4	Distinguishability of Patterns Experiment – Subjective Results	75

5.3.5	Distinguishability of Patterns Experiment – Discussion	76
5.4	Tactile Pattern Notification Experiment	78
5.4.1	Tactile Pattern Notification Experiment – Design & Implementation	79
5.4.2	Tactile Pattern Notification Experiment – Procedure	80
5.4.3	Tactile Pattern Notification Experiment – Objective Results	81
5.4.4	Tactile Pattern Notification Experiment – Subjective Results	82
5.4.5	Tactile Pattern Notification Experiment – Discussion	83
5.5	Limitations	85
5.6	Conclusion and Future Work	85
6	VRTactileDraw – A Tactile Pattern Design Interface for Complex Arrangements of Actuators	87
6.1	Introduction	88
6.2	Background and Related Work	90
6.2.1	Tactile Feedback	90
6.2.2	Tactile Pattern Editors	90
6.3	Iterative Design Process and Implementation	93
6.3.1	Virtual Reality Interface Design	93
6.3.2	Goals and Basic Features	94
6.3.3	Design and Implementation	94
6.4	Preliminary User Study	95
6.4.1	Design and Study Tasks	95
6.4.2	Implementation	96
6.4.3	Procedure	96
6.4.4	Participants	97
6.4.5	Results	97
6.4.6	Discussion and Implemented Improvements	98
6.5	Main User Study	99
6.5.1	Implementation and Procedure	99
6.5.2	Participants	99
6.5.3	Results	99
6.5.4	Discussion	100
6.5.5	Improvements After the Main Study	103
6.6	Limitations	105
6.7	Open Source Release and Extension to Tactile Interfaces	106
6.8	Conclusion and Future Work	107
7	Increasing Presence in Virtual Reality	109
7.1	Introduction	109
7.2	Related Work	110
7.3	Presence Experiment	110
7.3.1	Presence Experiment – Measuring Presence	111
7.3.2	Presence Experiment – Virtual Reality Environments	111

7.3.3	Presence Experiment – Procedure	113
7.3.4	Presence Experiment – Results	113
7.3.5	Presence Experiment – Discussion	115
7.4	Limitations	115
7.5	Conclusion	116
8	Intuitive and Precise 3D Guidance	117
8.1	Introduction	118
8.2	Related Work	119
8.3	Guidance Comparison Experiment vs. Visual and Auditory Stimuli	119
8.3.1	Guidance Comparison Experiment – Results	121
8.3.2	Guidance Comparison Experiment – Subjective Results	126
8.3.3	Guidance Comparison Experiment – Discussion	126
8.4	Refining the Algorithm for Precise Guidance	128
8.5	Guidance Precision Experiment	129
8.5.1	Guidance Precision Experiment – Objective Results	131
8.5.2	Guidance Precision Experiment – Subjective Results	131
8.5.3	Guidance Precision Experiment – Discussion	131
8.6	Real-World 3D Guidance Experiment	132
8.6.1	Real-World 3D Guidance Experiment – Results and Discussion	133
8.7	Limitations	134
8.8	Conclusion	134
9	Assistance of People with Visual Impairments	137
9.1	Introduction	138
9.1.1	Approach	139
9.1.2	Contributions	140
9.2	Related Work	141
9.2.1	Assistive Technologies for VIPs	141
9.2.2	VIP Guidance via Tactile Feedback	142
9.3	Initial Design of Tactile Patterns for Special Navigation Instructions	143
9.3.1	Informal Interview at an Educational Center for VIPs	143
9.3.2	Elicitation Study: Designing the Initial Tactile Patterns	144
9.3.3	Pattern Recognition Performance Study	145
9.4	Implementing the Micro-Navigation System	147
9.5	Obstacle Course Experiment	148
9.5.1	Obstacle Course Experiment – Design and Implementation	148
9.5.2	Obstacle Course Experiment – Procedure	149
9.5.3	Obstacle Course Experiment – Objective Results	151
9.5.4	Obstacle Course Experiment – Subjective Results	152
9.5.5	Obstacle Course Experiment – Discussion	154
9.6	Static Patterns Refinement Study	156
9.6.1	Static Patterns Refinement Study – Implementation	157

9.6.2	Static Patterns Refinement Study – Procedure	158
9.6.3	Static Patterns Refinement Study – Participants	159
9.6.4	Static Patterns Refinement Study – Results	159
9.6.5	Static Patterns Refinement Study – Discussion	160
9.7	Improved System Validation Experiment with VIPs	161
9.7.1	Improved System Validation Experiment – Implementation . . .	162
9.7.2	Improved System Validation Experiment – Procedure	163
9.7.3	Improved System Validation Experiment – Participants	164
9.7.4	Improved System Validation Experiment – Objective Results . .	164
9.7.5	Improved System Validation Experiment – Subjective results . .	166
9.7.6	Improved System Validation Experiment – Discussion	167
9.8	Overall Discussion	170
9.8.1	Limitations	172
9.9	Conclusion and Future Work	173
10	Summary and Conclusion	175
10.1	Limitations	178
11	Future Work	179
11.1	Vision of a Hidden and Precise 2D Tactile Micro-Navigation System . .	180
	Bibliography	183
	Acronyms	205
	List of Figures	207
	List of Tables	213
A	Appendix	215
A.1	Exemplary User Study Questionnaire	215
	Curriculum Vitae	225
	Declaration	229

Introduction and Motivation

Humans mainly perceive their environment via a set of five basic senses: *sight*, *hearing*, *touch*, *taste*, and *smell* [199]. While sight and hearing are omnipresent in today's world of information, the other senses are more subtle but just as important to grasp the whole experience of a specific location or setting. Imagine a woman visiting a Christmas market with her friends. Her sense of *sight* will lead her through the tightly packed street in search of wassail, food stands, and presents while making sure her friends follow along. Her sense of *hearing* will absorb the atmosphere of people chatting in the streets, pitchmen advertising their goods, and her friends happily talking about the approaching holidays. *Touching* and *tasting* a hot cup of wassail and feeling the comforting *touch* of her warm mantle and scarf while *smelling* her wassail, freshly baked bread, candy, and fir needles completes her experience.

Not all humans can perceive the five basic senses in the same way or even at all. People with hearing impairments might either hear a set or all frequencies attenuated or not at all. People with visual impairments experience their surroundings differently as their sense of sight is either impaired or completely absent. While the senses of touch, taste, and smell may also be impaired in one way or another, humans rarely suffer from impairment in multiple senses at the same time. People with impairments generally rely more on their remaining other senses. For example, people with visual impairments rely more on their sense of hearing and touch for identifying their location in the environment. Minor impairments of a basic sense may sometimes be cured by reasonably simple measures, e.g., reading glasses for hyperopia or a hearing aid for mild hearing loss. However, the more severe an impairment becomes, an improvement or a cure of the sense becomes quite expensive (e.g., eye surgery for glaucoma or cochlea implants for certain kinds of complete hearing loss) or is simply not yet possible to achieve.

Since researchers, notably Geldard et al. [54], started seriously experimenting with tactile feedback in 1957, significant progress was made in partially replacing the perception of a severely impaired sight or hearing sense through touch. People who suffer from hearing impairments may perceive certain sounds, phonemics, or words through touch via a set of dedicated tactile actuators [169, 204]. People with visual impairments or complete blindness may get navigation instructions and perceive points of interest around them via spatial tactile feedback [206].

State-of-the-art navigation and guidance systems use various technologies to stimulate the human visual, auditory, or haptic sensory channels. The visual channel is usually

the channel of choice, as it typically has a higher bandwidth than the other channels in terms of bits per second perceived [217]. However, sometimes the visual channel is not the desired primary channel for some kinds of feedback or in special situations such as when driving a car [212]. The visual channel may be severely limited for people with visual impairments, or it might be overtaxed, and essential feedback can be overlooked, or lighting conditions may prevent the user from seeing the feedback at all. Another reason to use the tactile instead of the visual or auditory feedback channels is faster initial reaction times, as shown in several studies such as [188].

The **overarching goal** of this thesis is *to establish the head as a means for tactile communication* by replacing parts of the senses of sight or hearing with the sense of touch, e.g., for spatial awareness, target searching, and micro-navigation, at a higher precision and accuracy compared to prior work thanks to the relatively high localization accuracy of tactile stimuli and natural mapping occurring on the head [159].

This thesis introduces a tactile grid around the head to relieve or replace an impaired visual channel, which I invented and named *HapticHead* (see Figure 1.1). It is a high-resolution, omnidirectional vibrotactile display worn on the head that presents general, 3D directional, and distance information through moving tactile cues and patterns. It consists of a grid of vibrotactile actuators arranged in three concentric ellipses around the head for relatively uniform coverage, optimized for head shape and user comfort. The head is well suited for guidance applications and tactile feedback, as it is sensitive to mechanical stimuli [57, 145] and provides a large spherical surface. This spherical surface allows displaying precise 3D information and allows the user to intuitively turn the head in a stimulus's direction due to natural mapping [159].

There are several application possibilities of the HapticHead concept, some of which are explored in this thesis. Along with all varieties of precise *2D* guidance and navigation use cases, possible applications for *3D* vibrotactile feedback around the head include:

- 3D guidance for visually impaired people,
- increasing the level of presence in VR with fitting tactile feedback,
- indicating plane positions and distances around air traffic controllers,
- displaying enemy positions and distances around fighter jet pilots,
- providing obstacle position warnings around drone pilots,
- signifying fish and obstacle positions around scuba divers,
- helping to find sought constellation or star positions above stargazers,
- indicating approaching traffic towards skiers (collision feed-forward),
- displaying object positions around crane operators, and
- enriching 360° videos with haptic feedback.




Figure 1.1.: HapticHead prototype, side and front view. The blue markers indicate ear openings, and the red markers indicate the locations of 10 actuators that form a “ring” around the face. The prototype contains a total of 24 actuators on three concentric ellipses.

Most of the use cases mentioned above require specially designed tactile patterns (TPs). For designing TPs, prior work only focused on simple arrangements of tactile actuators, which are configured in rings or 2D grids. This thesis presents the design and implementation of a sophisticated TP editor called *VRTactileDraw* for complex arrangements of tactile actuators, enabling even novices the rapid and cost-effective design of TPs for all kinds of tactile interfaces and use cases. *VRTactileDraw* is released alongside this thesis as open-source software¹.

The following subsection will introduce the underlying hypotheses and goals as well as the resulting approach taken in this thesis.

1.1 Approach and Goals

The following underlying **hypotheses** set a direction for the approach taken in this thesis:

H1 Tactile perception on the head is perceived as strongly *present* for the user and can thus hardly be missed as related work already found the head to be susceptible to tactile stimuli [143, 144, 145].

H2 Tactile guidance on the head may provide natural and *intuitive* guidance in three dimensions due to the spherical shape of the head, humans focus on the senses

¹*VRTactileDraw* open-source release: <https://github.com/obkaul/VRTactileDraw>

of sight and hearing, which are both located on the head, and natural mappings occurring from directional tactile patterns applied to the head [159].

H3 Tactile localization performance on the head and especially on the forehead is sufficiently *precise* to enable high-precision 3D guidance and micro-navigation as related work indicated a high localization precision on the forehead [105].

This thesis first examines the fundamentals of human perception and tactile feedback to establish a common ground. We also discuss various related work in tactile perception, tactile feedback, tactile displays, and tactile illusions in chapter 2.

To reach the overarching goal of this thesis, to establish the head as a means for tactile communication, the first subgoal **G1** is to research tactile sensitivities and relevant tactile illusions on the head to gain an understanding of human tactile perception on the head with implications for tactile user interface design. Chapter 3 provides an in-depth investigation into tactile perception, localization accuracies, and the tactile funneling illusion on the head. This chapter results from a CHI'20 publication [103] and clears up several misleading conclusions of related work in terms of localization performance and the tactile funneling illusion on the head.

The second subgoal **G2** of this thesis is to *construct a tactile interface around the head*, based on what has been learned from researching fundamentals and related work on tactile perception. The design approach, its limitations, and the HapticHead tactile user interface evolution through multiple iterations are shown in chapter 4. This chapter partly results from CHI'16 [98], and CHI'17 [99] publications.

Chapter 5 deals with the third subgoal **G3**, which is to *research how many and what kinds of tactile patterns on the head users can distinguish and which kinds of patterns are accepted among the user base*. We can show that the addition of spatial location dramatically increases the total distinguishable number of patterns available to the user compared to sole intensity and rhythm modulation and which kind of patterns are more accepted amongst the experiment participants in terms of recognition accuracy and comfort. This chapter is primarily based on a MUM'20 [101] publication.

With our experience working with TPs, we learned that designing them is a daunting task when defining the intensities over time for many actuators individually (e.g., for HapticHead). The need for an easy-to-use TP editor for tactile user interfaces featuring complex arrangements of many actuators could not be adequately satisfied by related work. Thus, the final subgoal **G4** of this thesis is to *build a pattern design software that can deal with a large number of actuators in complex configurations*, allowing even novices to design their own TPs for complex tactile user interfaces. Chapter 6 is dedicated to fulfilling G4 by describing the process of the design and implementation of our tactile design software VRTactileDraw and its evaluation. We found that even novice users with non-technical backgrounds can work with VRTactileDraw to design

relatively complex TPs for various use cases. This chapter is primarily based on an INTERACT'21 publication [93].

The following three chapters investigate several different use cases of HapticHead, all with the overarching goal to establish the head as a means for tactile communication.

Chapter 7 is dedicated to increasing the feeling of presence in Virtual Reality scenarios by attempting to increase immersion through HapticHead. We show that adding vibrotactile feedback around the head can significantly increase the level of presence users experience in certain VR scenes. This can be applied in immersive games and specific VR training simulations where the level of presence or spatial awareness (including collision prevention) is essential such as in complex maintenance jobs, anxiety therapy, or flight training. This chapter is based on an INTERACT'17 publication [95].

Several 3D guidance use cases of HapticHead are introduced in chapter 8. Our experiments indicate that HapticHead tactile guidance is significantly faster and more precise than spatial audio for finding visible and invisible virtual objects in 3D space around the user. We further successfully navigated blindfolded users to real household items at different heights using HapticHead as a 3D guidance Augmented Reality device. This chapter is based on the same CHI'16 [98], and CHI'17 [99] publications as chapter 4.

With HapticHead established as a highly precise guidance system, another compelling use case is the assistance of people with visual impairments in micro-navigation, one of the main focuses of this thesis and discussed in-depth in chapter 9. In this chapter, we first elicit a set of four fundamental navigation instructions and accompanying tactile patterns. Subsequently, we test these patterns in two micro-navigation experiments with seeing but blindfolded and visually impaired users. We found that our final implemented micro-navigation guidance system performs an order of magnitude more precisely than a state-of-the-art tactile belt from related work. In addition, we can signify important navigation instructions such as “stairs up” or “STOP” to the user through highly intuitive TPs powered by natural mappings. This chapter is mainly based on a TOCHI'21 publication [102].

Finally, chapter 10 provides an overall discussion and conclusion of using tactile feedback around the head, while chapter 11 discusses various possible future work directions.

The author has supervised several Bachelor's and Master's theses. These theses contributed to several concepts, systems, and ideas presented in this dissertation [14, 21, 39, 44, 51, 53, 64, 65, 115, 134, 139].

1.2 Contributions

With the overarching goal of this thesis to establish the head as a means for tactile communication, it presents the following primary contributions to the field of tactile research:

- **An in-depth investigation into tactile localization accuracy and the occurrence of the funneling illusion around the head**

This investigation highlights areas around the head more suited for dense tactile feedback (e.g., the forehead and chin) and tests the occurrence of the funneling illusion, which may potentially be used to provide more comfortable and accurate tactile stimuli to users. The results show strongly varying localization performances, funneling illusion occurrences, and other characteristics of tactile feedback across different head regions. We computed a detailed heat map of vibrotactile localization accuracies on the head. Our results inform the design of future tactile head-mounted displays with several cues that were not available so detailed in prior work. We further clear up several misleading conclusions of related work regarding localization performance and tactile funneling illusion on the head.

- **HapticHead, a sophisticated tactile display around the head**

With influences from related work and our fundamental research on head-based tactile perception, this thesis presents the design, implementation, and evolution of a sophisticated tactile display around the head, which may be employed in many use cases (see above). The HapticHead tactile display consists of 24 vibrotactile actuators spread across the head, neck, and chin. It was optimized through multiple iterations for comfort and guidance performance. The final hardware design includes the 24 actuators mounted on a fiber bathing cap with a chin-strap, a custom circuit board to drive the actuators, and a Raspberry Pi 4, which accepts commands wirelessly. Prior state-of-the-art tactile guidance systems did not allow the same level of precision and accuracy achieved by HapticHead in several guidance and navigation studies.

- **Fundamental research on tactile patterns around the head**

While prior work focused mainly on tactile sensitivities or tactile systems with specific use cases on the head, our investigation fundamentally evaluates recall and distinguishability of 30 tactile patterns around the head and agreement on meaning without a predetermined context. The agreement is low, yet the recognition rate is surprisingly high. We identify which kinds of patterns users recognize well and which ones they prefer. Static patterns with a static stimulus location have a higher recognition rate than static patterns with a dynamic stimulus location, which move across the head as they play. However, participants

prefer the latter type for comfort. We also show that participants can distinguish substantially more around-the-head spatial patterns than smartphone-based patterns. The spatial location has the highest positive impact on accuracy among the examined features, so this parameter allows for a large number of levels.

- **A Virtual Reality tactile pattern design software**

We contribute a Virtual Reality interface called *VRTactileDraw* that enables designers to rapidly prototype complex tactile interfaces. It allows for painting strokes on a modeled body part and translates these strokes into continuous tactile patterns using an interpolation algorithm. The presented VR approach avoids several problems of traditional 2D editors. It realizes spatial 3D input using VR controllers with natural mapping and intuitive spatial movements, which are significant advantages over prior work. To evaluate this approach in detail, we conducted two user studies and iteratively improved the system. The study participants gave predominantly positive feedback on the presented VR interface (SUS score 79.7, AttrakDiff “desirable”). The final system is released alongside this thesis as an open-source Unity project for various tactile hardware.

- **An investigation on improving the feeling of presence in Virtual Reality through head-based tactile feedback**

We use HapticHead with specially designed vibrotactile feedback to increase users’ perceived presence in virtual reality scenes. In a between-groups comparison study, the vibrotactile group scored significantly higher in a standardized presence questionnaire compared to the baseline of no tactile feedback. While related work has achieved an increased level of presence with various approaches, our system is one of the few that can easily be integrated into VR or AR HMDs.

- **An investigation on HapticHead’s 3D guidance performance in Virtual and Augmented Reality**

We conducted three experiments on HapticHead’s 3D guidance performance in Virtual and Augmented Reality, which indicate that HapticHead vibrotactile feedback is both faster (2.6 s vs. 6.9 s) and more precise (96.4 % vs. 54.2 % success rate) than spatial audio for finding visible virtual objects in 3D space around the user. The baseline of visual feedback is – as expected – more precise (99.7 % success rate) and faster (1.3 s) in comparison, but there are many applications in which visual feedback is not desirable or available due to lighting conditions, visual overload, or visual impairments. Mean final precision with HapticHead feedback on invisible targets is 2.3° compared to 0.8° with visual feedback. We further successfully navigated blindfolded users to real household items at different heights using HapticHead as an Augmented Reality device in a lab study. Compared to prior state-of-the-art tactile-only guidance systems,

HapticHead is the first to offer genuine intuitive *3D guidance* thanks to natural mapping, and it also achieves substantially higher accuracy and precision compared to, e.g., tactile belt guidance systems.

- **A tactile system to micro-navigate people with visual impairments**

In advancing the fields of tactile micro-navigation and assistive technologies for people with visual impairments (including fully blind people), we make the following two contributions: First, a set of fundamental navigation instructions and associated *intuitive* tactile patterns on the head powered by natural mappings, optimized in three consecutive studies for micro-navigation. Secondly, we contribute a system using these optimized tactile patterns for essential navigation instructions combined with a continuous tactile guidance stimulus to provide precise micro-navigation instructions. The system supports navigation around obstacles and on stairs. The precision of our system is substantially better (5.7 cm mean deviation from the optimal path) than prior state-of-the-art tactile micro-navigation systems (e.g., 49 cm mean deviation from the optimal path [52]).

Foundations of Human Perception and Tactile Feedback

This chapter provides an overview of the human neuroanatomy and receptors responsible for detecting vibrations, what tactile feedback entails, and by what means it is generated. We focus primarily on aspects related to tactile interfaces around the head.

We start by describing the human nervous systems, individual nerve cells, and how they interact with each other. Subsequently, we introduce different specialized receptors (afferents) to perceive external mechanoreceptive stimuli on the human skin. The human skin structure is essential because the number of afferents present in the skin varies widely depending on body parts and whether it is glabrous or hairy skin. Next, we present ways to generate mechanoreceptive stimuli using various tactile actuators. Then we discuss tactile patterns, their terminology, and thorough background of existing tactile displays in related work with specific regards to the use case of guidance via tactile displays. Finally, research in the area of tactile illusions is presented to provide a seamless transition to the following chapter on tactile illusions and localization accuracy of tactile stimuli on the head.

The content of this chapter is partly based on the books “*Atlas of Functional Neuroanatomy*” (2015) by Walter Hendelman [68], “*Stevens’ Handbook of Experimental Psychology*” (2002) by Pashler and Yantis [155], “*Funktionelle Neuroanatomie : Lehrbuch und Atlas*” (1994) by Zilles and Rehkämper [230], “*Neuro- und Sinnesphysiologie*” (2006) by Schmidt and Schaible [182], “*Scholarpedia of Touch*” (2016) by Prescott et al. [165], and further a 2018 doctoral thesis by Andreas Tarnowsky [205]. In the following sections, only the fundamental’s basic aspects regarding this thesis are described. For further information, please refer to the sources mentioned above.

2.1 Human Nervous System

The nervous system of humans and other vertebrates can be divided topographically into two areas. While the central nervous system (CNS) refers to the unit consisting of the brain and the spinal cord, the peripheral nervous system (PNS) comprises the body’s remaining nerve conduits. Two kinds of connections exist between the PNS and the brain: direct connections of specific body areas through 12 pairs of cranial

nerves to the brain and indirect connections of the other body areas to the spinal cord through 31 pairs of spinal nerves. Through these connections, all areas of the body may be reached with information flowing in both directions.

2.1.1 Nerve Cells

Depending on the direction of the information flow, a distinction is made between *afferent* and *efferent* nerves. Motor body functions, for example, are initiated by the CNS and transmitted to the corresponding target system via the efferents. Sensory information, such as ground surface properties perceived by the feet, is transmitted to the CNS by afferents.




Figure 2.1.: Fundamental structure of human nerve cells. Note that the structure changes based on the specialization of the cell. Some specialized cell elements may be altered or not present compared to this general structure (e.g., not all axons spawn axon collaterals). Figure source: own work, with background from [68, 205, 230]

Nerve cells (*neurons*) are the foundation of any nervous system. They collect, process, and forward external stimuli, depending on their specialization. Figure 2.1 shows the general structure shared by most human neurons with slight discrepancies depending on the neuron's particular function.

The cell core (*nucleus*) is enclosed by the cell body (*soma*). Several afferents usually spawn from the soma as branches (*dendrites*). These dendrites collect external stimuli and may be connected to other neurons via *synapses*. The soma further spawns an axon hillock, which allows the neuron to communicate with other neurons through electrical signals over potentially large distances via an *axon*.

An axon is essentially a string of *Schwann cells* surrounded by *Myelin sheaths* which are responsible for electrical insulation between axons. *Ranvier nodes* connect Schwann cells and may serve as a connection point for *axon collaterals* which directly branch off axons in some cases (see [230, p. 44]). Axons vary in length between a few micrometers and several meters in larger animals.

Eventually, the axon leads into an *axon terminal* which is a hub of connectors to either the dendrites of other neurons or directly to the soma of other neurons. The connection points are called *synapses* and transmit information either via chemical *neurotransmitters* or *electrical stimuli*. Details of these processes may be found in [68, 230].

2.1.2 Firing of Neurons

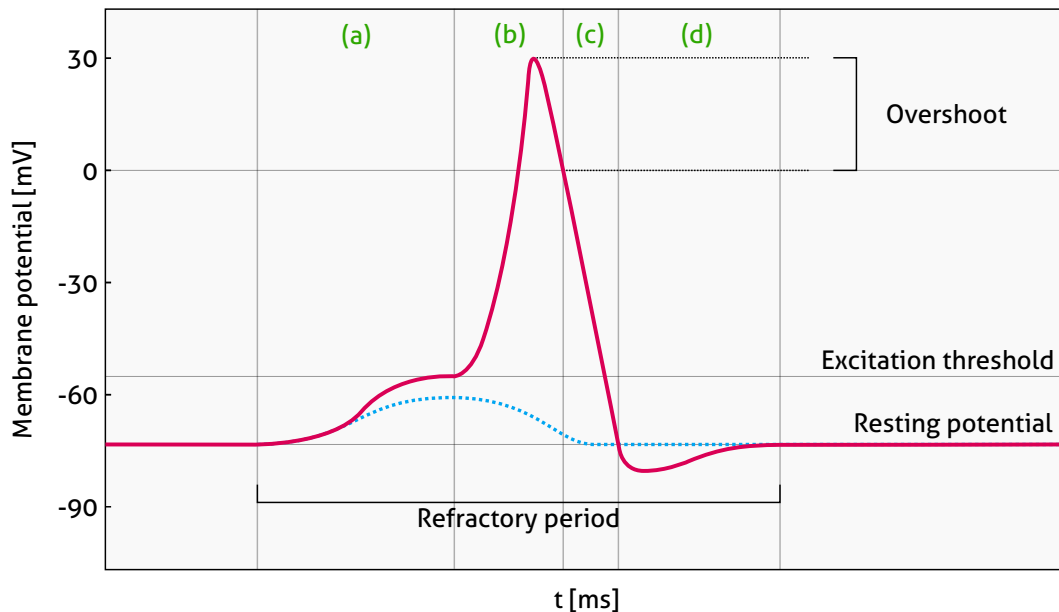


Figure 2.2.: Schematic curve of the action potential of a neuron in case of a successful firing (red curve) and a curve where the excitation threshold was not reached (blue curve). The successful firing can be split into four phases: (a) initiation phase, (b) depolarization, (c) repolarization, and (d) hyperpolarization. Figure adapted from [205, Fig. 2.3]

Neurons actively generate electric signals based on one or multiple external stimuli collected by their dendrites or through direct synaptic connections to other neurons. These stimuli may either increase or decrease the neuron membrane potentials. The firing of an abstract neuron is shown in Figure 2.2. This process can be divided into the following phases:

- (a) **Initiation phase:** The membrane potential of a neuron continuously changes based on the sum of positive or negative input stimuli from its dendrites and potential direct synaptic connections to other neurons. In case the membrane potential reaches a certain cell-specific *excitation threshold*, it enters the next depolarization phase. Else, the membrane potential returns to its *resting potential* without entering the next phase.

- (b) Depolarization phase: The neuron quickly builds up an *action potential*, sending an electrical signal down its axon. This process is called the firing of a neuron. Neurons always fire at their full strength; there is no such thing as a neuron's partial firing. Depending on the neuron type, the potential may even become positive for a short while (*overshoot*).
- (c) Repolarization phase: After a period of time that depends on the specific cell type, its membrane potential returns to its *resting potential*.
- (d) Hyperpolarization phase: After repolarization, the membrane potential may drop lower than the resting potential of a cell. This phase may last between a few milliseconds to several seconds, depending on cell type [182, p. 28].

After the membrane potential returns to its resting potential in the hyperpolarization phase, the cell is ready to fire again. The entire duration of firing until returning to its resting potential is called *refractory period* which is highly cell type-dependent and limits the maximum firing rate of the neuron. For example, a human heart muscle is actuated only a few times per minute depending on the firing of the *sinoatrial node*. However, the hearing nerve (*Nervus acusticus*) fires at around 1000 Hz [205, p. 9].

2.1.3 Perception of External Stimuli

The perception of external stimuli requires specialized nerve cells which are collectively called *receptors*. In literature, four different kinds of physicochemical receptor classes are mentioned [182, p. 182]:

- mechanical deformation: haptic/tactile perception and hearing,
- chemical stimuli: smell, taste, endogenous substances,
- temperature: warming and cooling,
- light: visual perception.

Most receptors are highly specialized and react most strongly to *adequate stimuli* within their scope and less or not at all to stimuli outside their scope. Some receptors may perceive *inadequate stimuli* from a different stimuli class outside their scope. For example, the mechanoreceptors of the skin, which usually sense vibration or surface properties, can be stimulated using *electrical* instead of mechanical stimuli.

The skin is the largest sensory organ of humans by surface area. Thirteen different afferents contribute to its sensitivity and these afferents can be categorized into five groups [155, p. 537]:

- two *pain* afferent types,
- one *itch* afferent type,
- two *thermoreceptive* afferent types,

- four *proprioceptive* afferent types (sense muscle length, muscle force, and joint angle),
- four *mechanoreceptive* afferent types (sense skin deformation, mechanical pressure, vibration).

The following section will describe the general structure of human skin as with special regards to the mechanoreceptive afferent types.

2.2 Structure of Human Skin

The general structure of human skin is shown in Figure 2.3. The outmost layer is known as *stratum corneum* (sometimes referred to as part of the adjoining *epidermis*). It acts as a natural barrier against environmental influences and is resistant to most bacteria, viruses, and even some chemical substances. The stratum corneum consists of several layers of dead skin cells continually regrowing from the epidermis below. It heavily varies in the number of layers and thickness depending on the body part, mechanical stress, and various other individual factors [72, 142, 222]. For example, the stratum corneum on skin areas around the head consists of roughly 7 - 12 layers with a few μm thickness. In contrast, the palms and heels consist of about 50 to 86 layers with a thickness of approximately 0.3 - 1 mm [222]. Individual mechanical stress significantly changes the stratum corneum over time. People who often walk barefoot, for example, grow a very thick stratum corneum under their feet.

Below the stratum corneum lies the *epidermis*, which contains live cells responsible for constantly growing more stratum corneum layers from below. This layer also contains dendrites from various afferent cells from below the epidermis (see Fig. 2.3). The epidermis measures roughly 0.1 mm in thickness on areas around the head [223].

The epidermis is connected to the underlying *dermis* layer through the *papillary layer*, a wave and cone structure that mechanically binds epidermis and dermis together and improves the sustenance of the epidermis due to the larger surface area compared to a flat connection between these layers [205]. The dermis itself consists primarily of fascia and collagen fibers, which causes a high degree of elasticity and tearing strength. Depending on body area, the dermis is between 0.5 and 2 mm thick (roughly 0.6 to 1.2 mm on areas around the head) [223].

Below the dermis resides the final skin layer called the *hypodermis* or *subcutis*. It consists of loose fascia chambers filled with fat cells and varies widely in thickness, depending on the body area and underlying tissue. On the head, the hypodermis is strongly connected to the periosteum of the head in most areas. This almost immovable composition is called the *epicranium* or scalp.

The following section will introduce details about tactile perception and the mechanoreceptive afferent types present in human skin.

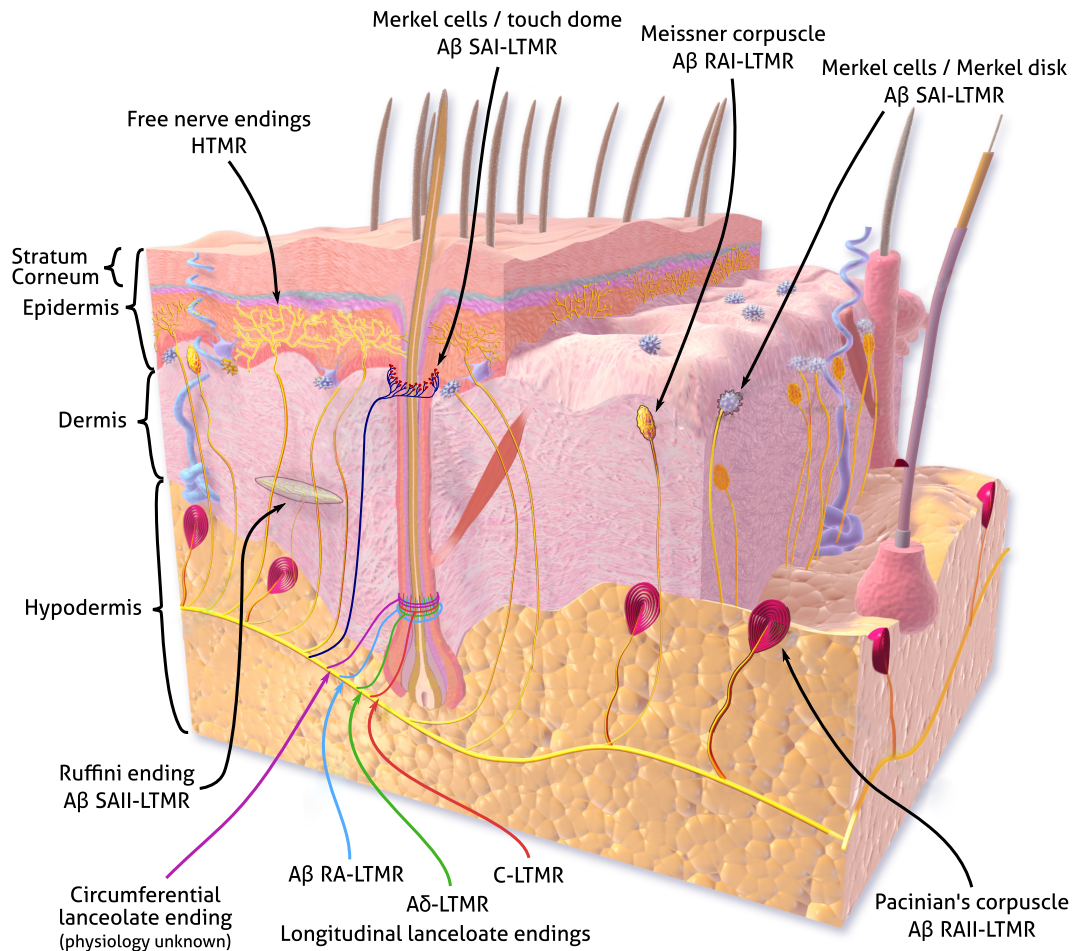


Figure 2.3.: Mechanoreceptive afferents (tactile receptors) in the human skin. Adapted from [201]; enriched by additional information on hair afferents by [1]. Note that there are substantial differences between glabrous and hairy skin. This image shows afferents present in both glabrous and hairy skin. Not all of the shown hair afferents are present on all types of hair (see text).

2.3 Tactile Perception

Developing a tactile display requires knowledge about tactile perception and the afferents responsible for sensing tactile stimuli regarding afferent properties (e.g., peak sensitivity) and the application area on the human body. Tactile perception depends primarily on the mechanoreceptive afferent types, and they are the main focus of this section.

Afferent type	A β SA1-LTMR	A β SA2-LTMR	A β RA1-LTMR	A β RA2-LTMR
Receptor	Merkel	Ruffini	Meissner	Pacinian
First discovered (year)	1875	1894	1853	1741
Ending type	A β (40-65 m/s)	A β (35-70 m/s)	A β (35-70 m/s)	A β (35-70 m/s)
Location	Tip of epidermal sweat ridges	Dermis and deeper tissues	Dermis	Dermal papillae (close to surface)
Sensory function	Form and texture perception	Tangential force, hand shape, motion direction	Motion detection, grip control	Tool use, perception of distant events through vibrations
Effective Stimulus	Edges, points, corners, curvature	Skin stretch	Skin motion	Vibration
Response to sustained indentation	Sustained with slow adaption	Sustained with slow adaption	None	None
Innervation density (finger pad)	100/cm ²	10/cm ²	150/cm ²	20/cm ²
Frequency range	0 - 100 Hz	0 - ? Hz	1 - 300 Hz	5 - 1000 Hz
Peak sensitivity	5 Hz	0.5 Hz	50 Hz	200 Hz
Spatial acuity	0.5 mm	> 7 mm	3 mm	> 10 mm

Table 2.1.: Low-Threshold Mechanoreceptors (LTMRs) in glabrous skin and their properties. Excerpt from [155, p. 541].

The perception of harmless and harmful touch sensations relies on specialized mechano-sensitive sensory neurons that fall into two general categories: low-threshold mechanoreceptors (LTMRs), which respond to harmless mechanical stimuli, and high-threshold mechanoreceptors (HTMRs), which respond to harmful mechanical stimuli [1].

The human skin structure varies between bald/hairy and thick/thin skin and features different receptor densities, depending on the area. The highest density of mechanoreceptors can be found on the fingertips. Table 2.1 shows properties of the four mechanoreceptive afferents present in glabrous skin. For hairy skin, a partly different set of afferents is responsible for sensing touch. Their properties are shown in Table 2.2. Figure 2.3 shows the mechanoreceptive afferents roughly in their respective depths in the human skin.

Not all of the afferents shown in Figure 2.3 are present in all skin types. For example, glabrous skin does not contain any hair afferents, and hairy skin does not contain Meissner corpuscles (A β RA1-LTMR). Furthermore, not all hair types contain all the hair afferents shown in Figure 2.3 (see subsection 2.3.6).

Before going into the details of the skins' mechanoreceptors, the following subsection is dedicated to David Katz's findings, one of the first researchers in tactile perception.

2.3.1 Early Research on Tactile Perception

A pioneer of tactile research is the experimental psychologist David Katz who started working on tactile perception in 1914 after finishing his habilitation thesis on color perception. While most of the mechanoreceptors in glabrous skin were discovered and medically described before David Katz was born in 1884, he was the first researcher who conducted experiments on tactile perception [205].

Several findings of Katz's initial book on tactile perception "Der Aufbau der Tastwelt" (1925) [90] (see [116] for an English review) are still relevant today. For example, Katz estimated the human tactile perception to be most sensitive for stimuli in the range of 50 to 500 Hz, which is close to current study results [90, p. 201]. He further described specific properties of slowly adapting mechanoreceptors as "fatigue effect" and first noted "masking effects" of mechanoreceptors [90, p. 205].

Katz also found the "duplex theory of tactile perception", which states that the skins' impressions of vibration and pressure are two distinct and separate perception channels [90, p. 187]. This theory has also been repeatedly confirmed in modern tactile research [56, 73].

The following subsection will introduce details about the mechanoreceptive afferent types present in glabrous and hairy human skin.

2.3.2 Merkel Disks - $A\beta$ SA1 Low-Threshold Mechanoreceptors



Figure 2.4.: Detail view of tactile receptors in the skin. From left to right: Merkel disk ($A\beta$ SA1-LTMR), Ruffini corpuscle ($A\beta$ SA2-LTMR), Meissner corpuscle ($A\beta$ RA1-LTMR), Pacinian corpuscle ($A\beta$ RA2-LTMR). Adapted from [201].

Merkel cells were first discovered and investigated by Friedrich Merkel in 1875 [137]. These afferents play a crucial role in perceiving forms and textures [155, p. 543]. As they are in the group of slowly adapting (SA) receptors, Merkel cells can detect both short and long-lasting stimuli with a peak sensitivity around 5 Hz.

In human skin, Merkel cells are usually situated in small, synaptically interconnected groups in the papillary layer (between epidermis and dermis) in glabrous skin and around the hair follicles in hairy skin. These interconnected groups are known as *Merkel-cell neurite complexes* or *Merkel disks* in glabrous skin [155] and as *touch domes*

in hairy skin [29]. Touch domes are solely present around a special type of hair, classified as *guard* hair [1].

At around 100 afferents/cm², the density of Merkel cells is highest on the fingers. Regions around the head feature an afferent density of around 16 (scalp), 18 (forehead), and 33 (neck) afferents/cm² [117].

2.3.3 Ruffini Endings - A β SA2 Low-Threshold Mechanoreceptors

First discovered by Angelo Ruffini [170] in 1894, Ruffini endings are also in the group of slowly adapting receptors. As they are significantly less sensitive and feature a lower spatial acuity than Merkel cells, their role in tactile perception is controversial [155, p. 541]. Based on these properties, it is unlikely that Ruffini endings contribute to detecting local surface details. Since they respond primarily to stretching within the skin, it seems more plausible that they serve to detect movement and register the position of the hand and fingers (proprioception) [155, p. 571].

2.3.4 Meissner Corpuscles - A β RA1 Low-Threshold Mechanoreceptors

Georg Meissner was the first to present a description of Meissner corpuscles, a rapidly adapting receptor type to detect skin motion and low-frequency vibrations in 1853 [135]. As a member of the group of fast adapting receptors, Meissner corpuscles are nearly impervious to static stimuli. They are biologically structured in a way so that frequencies below 2 Hz are mechanically filtered out. Like Merkel cells, Meissner corpuscles can also be found in the papillary layer between the epidermis and dermis.

With a spatial acuity of 3 cm, they are significantly less accurate than Merkel cells but significantly more accurate than Pacinian corpuscles. Meissner corpuscles are considered to play an essential role in the perception of subtle movements on the skin. Since they can detect slippage or small changes in contact force in this way, they are essential for grip control [155, p. 566]. Consequently, Meissner corpuscles are found in glabrous skin, mainly in the fingertips, palms, and soles. Hairy skin does not contain Meissner corpuscles [29].

2.3.5 Pacinian Corpuscles - A β RA2 Low-Threshold Mechanoreceptors

Initially discovered by Abraham Vater in 1741 [214], (Vater-)Pacinian corpuscles were later “re-discovered” and investigated by Filippo Pacini (see [69]). They are also in

the group of rapidly adapting receptors. Their biological structure features a unique onion-like structure around the nerve, which acts as a mechanical high-pass filter (see Figure 2.4). This high-pass filter fully cancels out any pressure stimuli below a frequency of 5 Hz, strongly attenuates stimuli below 150 Hz, and allows the receptor to showcase a relatively high peak sensitivity at around 200 Hz (see [155, p. 569]).

Recent studies show that Pacinian corpuscles are the only mechanoreceptors capable of detecting and encoding complex vibrotactile stimuli, i.e., the temporal structure of high-frequency signals [155, p. 570].

With a relatively low innervation density of 20 afferents/cm² even on the finger pads, PC receptors feature a significantly lower spatial acuity than the other receptor types. Pacinian corpuscles rarely appear in hairy skin, including large areas of the head [1, 29, 211].

2.3.6 Specialized Afferents in Hairy Skin

Humans and other mammals rely heavily on the perception of a variety of touch sensations through hairy skin, including social interactions and the presence of foreign objects or organisms on our skin. However, far fewer tactile perception studies have been carried out on hairy skin compared to glabrous skin. The perception of tactile stimuli in hairy skin differs significantly in physiology from that of glabrous skin.

Several LTMR sub-types innervate hairy skin and are physically and functionally associated with hair follicles. These LTMR sub-types fall into A β -, A δ -, and C-type classes depending on conduction speeds (see Table 2.2) [1]. In a review by Abraira and Ginty [1] from 2013, three different types of hair are described: guard, awl/auchene, and zigzag hair. Not all hair types contain all the hair afferents shown in Figure 2.3. An exception is the circumferential lanceolate ending afferent, which is present on all types of hair, but its physiological properties remained unknown at the time of Abraira and Ginty's review in 2013.

Just like on glabrous skin, A β -LTMRs are divided into the two groups of slowly (SA) and rapidly adapting (RA) receptors, based on their firing adaption rates [1]. A β SA1-LTMRs are associated with Merkel cells which are organized as touch domes around the hair shafts between the epidermis and dermis (see Figure 2.3). These A β SA1-LTMRs are only present on specialized **guard hairs**, which are usually thicker and longer than the other hair types. Guard hairs are also surrounded by A β RA-LTMR longitudinal lanceolate endings (LLEs) with similar characteristics as Pacinian corpuscles, generating a response to mechanical stimuli at up to 400 Hz in cats and rabbits [23].

Just like guard hairs, **Awl/auchene hairs** are surrounded by A β RA-LTMR LLEs and further by A δ -LTMR LLEs, and C-LTMR LLEs [1]. A δ -LTMRs were initially described as D-Hair units, reflecting their specific response to small sinus movements and down

Afferent type	A β SA1-LTMR	A β RA-LTMR	A δ -LTMR	C-LTMR
Ending type	A β (40-65 m/s)	A β (26-91 m/s)	A δ (5-30 m/s)	C (0.2-2 m/s)
Associated fiber conduction speed	Merkel cell (touch dome)	LLE	LLE	LLE
Location	Surrounding guard hair follicles, between epidermis and dermis	Surrounding hair follicles in the dermis	Around hair follicles in the dermis	Surrounding hair follicles in the dermis
Hair types	Guard hair	Guard hair, Awl/auchene hair	Awl/auchene hair, Zigzag hair	Awl/auchene hair, Zigzag hair
Optimal stimulus	Indentation	Hair follicle deflection (fast)	Hair follicle deflection (slow)	Hair follicle deflection (slow)

Table 2.2.: Low-Threshold Mechanoreceptors (LTMRs) longitudinal lanceolate endings (LLEs) surrounding hair follicles in hairy skin. Table generated from data available in [1, 155].

hairs in cats and rabbits. A δ -LTMR-like responses can be found in humans, but they do not always correlate to hair follicle movements, and it remains unclear whether and how A δ -LTMRs influence touch perception in humans. C-LTMRs are suggested to be associated with tickling sensations, and C fibers are three to four times more numerous than A fibers. C-LTMRs are especially sensitive to skin indentation and are maximally activated by slowly moving stimuli (e.g., caressing). In humans, C-LTMRs are present in hairy skin and are speculated to play a role in mediating “emotional touch”.

Finally, **Zigzag hairs** are surrounded by the previously described A δ -LTMR LLEs, and C-LTMR LLEs [1].

A recent review of innervation densities across the whole body reports an innervation density of 17 units/cm² for the neck and scalp, with 45 % of these units being RA receptors. In comparison, the forehead (including eyes and nose) has an innervation density of 48 units/cm² with 35 % being RA receptors [29]. With the higher density of RA receptors on the forehead than the scalp and neck, a higher spatial acuity of vibration localization can be expected.

2.3.7 High-Threshold Mechanoreceptors

High-threshold mechanoreceptors (HTMRs) respond to harmful mechanical stimuli (e.g., tearing of the skin). These receptors are present in the lower layers of the epidermis, closer to the surface than all LTMRs (see Figure 2.3). HTMRs are usually not relevant in tactile feedback applications, as they do not react to non-harmful stimuli on the skin.

2.3.8 Tactile Perception around the Head

In a series of experiments, Myles et al. [143, 144, 145, 146] investigated the vibrotactile sensitivity of different head regions and hair densities and used a headband with seven C2 Tactors to provide vibrotactile stimuli with varying characteristics in intensity and frequency to soldiers. They found that soldiers preferred a tactile to a visual or auditory display for directional cueing and that the forehead, frontal, parietal, and temple regions were most sensitive to tactile stimuli. Dobrzynski et al. [43] investigated information transfer capabilities of a ring of 12 vibrotactile actuators around the head concerning the maximum number of active actuators and the maximum comfortable vibration intensity. They strongly suggest avoiding the use of multiple simultaneous actuators to show different directions. The more actuators were active at a time in their experiment, the lower the correct recognition rates for the number of active actuators became. They recorded a 94 % recognition rate for a single actuator, 40 % for two actuators, and only 5 % for five simultaneously active actuators.

While the perception threshold of tactile stimuli was thus already researched in-depth by Myles et al., the localization accuracy of tactile stimuli with amplitudes stronger than the perception threshold on the respective head regions was still mostly unclear before research presented in this thesis was started, even though Dobrzynski et al. conducted some research in this area. However, their conclusions were vague [43]. In the next chapter (chapter 3), this thesis presents research on the localization performance of vibrotactile stimuli on the head. We were able to show that the forehead achieves the highest precision in the localization of vibrotactile stimuli (0.7 cm mean absolute deviation), followed by the frontal top of the head (0.9 cm deviation), the chin (1.2 cm deviation), and the bottom back of the head (1.2 cm deviation). The rear top (1.4 cm deviation) and sides of the head (1.5 cm deviation) relatively scored the worst in localization performance.

Together with the work mentioned above by Myles et al. [143, 144, 145, 146], the forehead presents itself as a great candidate for tactile patterns in high-precision guidance applications. The following section will introduce approaches to generate adequate (e.g., skin deformation) and inadequate (e.g., electrical) tactile stimuli.

2.4 Existing Approaches to Generate Tactile Stimuli

This section introduces various tactile systems and displays to generate a variety of tactile stimuli. First, we look at displays stimulating the *thermoreceptive afferents* through Peltier elements. Secondly, *Electrotactile displays* are presented, which may inadequately stimulate a skin area and underlying nerves through two or more electrodes. A defined voltage curve is applied to the skin area to be stimulated, whereby the receptors' nerve endings are induced to build up an action potential. Finally, typical

mechanical tactile actuators are presented, which indirectly but adequately stimulate the mechanoreceptors by local deformation of the skin.

The head contains no Meissner corpuscles on the hairy parts, and using an electrotactile display on the head is highly questionable in regards to safety. For the system developed in this thesis, we are looking to generate mechanical tactile stimuli at a frequency of around 200 Hz because Pacinian-type receptors feature a peak sensitivity around 200 Hz. The exact peak sensitivity frequency of afferents surrounding hair follicles mentioned in the previous section is unknown at the point of writing this thesis [1].

2.4.1 Thermal Stimulation

The human skin can perceive coldness and warmth through thermal afferents [182]. Researchers use temperature perception to develop thermal-tactile displays which feature Peltier elements [77, 224]. Peltier elements are thermoelectric converters that allow transporting thermal energy using an electric current to achieve either a heating or a cooling effect. An early thermal-tactile display featured a Peltier element with direct skin contact on the fingertip to present a predetermined temperature curve [78]. It is further possible to indirectly heat or cool the skin through Peltier elements attached to metal or other mechanical actuators consisting of metal. Yang et al. combined mechanical and thermal tactile feedback in a 5x6 pin-array piezoelectric bimorph display with attached Peltier elements [77]. Their users were able to sense virtual materials through their thermal and surface properties.

Typical thermal-tactile displays achieve cooling and warming rates of 6 and 12° C per second [77], which is significantly too slow for applications involving real-time alarms or guidance. Further disadvantages include the large power consumption and size/weight of thermal-tactile feedback systems (e.g., Yang et al. had to use a jacket filled with water for their cooling effect [77]).

2.4.2 Inadequate Stimulation of Receptors through Electrical Stimuli

Apart from adequate mechanical stimuli, inadequate electrical stimuli can also be used on the skin's mechanoreceptors using two or more electrodes. Typical electro-tactile displays consist of a grid of electrodes and a current-limiting circuit that outputs up to 320 V of alternating current to overcome the skin's natural resistance in order to apply a current of around 0.5 - 2 mA [84]. The skins' resistance varies widely across the body and is usually more extensive, the thicker the stratum corneum is and can also be influenced by hair density. It may effectively be decreased by using electrode gel or wet electrodes.

Hiroyuki Kajimoto is an advocate for grid-based electro-tactile displays. He presented works on electro-tactile displays for enhancing the touchscreen experience with electro-tactile feedback using a transparent electro-tactile display [85] and as a vision substitution system using a 512-electrode electro-tactile display on the back of a smartphone [86]. He even applied the same 512-electrode display on the forehead for vision substitution [83].

While the latter system with a large number of electrodes on the forehead first seems comparable to the mechano-tactile around-the-head display presented in this work, a few key differences prevent the use of similar electro-tactile feedback for the Haptic-Head concept. First, Kajimoto applied his display solely on glabrous skin. Its efficiency and connectivity on hairy skin are unknown and would likely require significantly different electrodes to conduct electricity through hairy skin reliably. Furthermore, the HapticHead concept relies on stimulating various sites all *around* the head, not just on the forehead. Thus, the electrical circuit has to prevent current flowing through brain areas by design, which is difficult to achieve. Kajimoto et al. did not discuss potential safety issues of applying 300 V stimuli on the forehead in their paper [83].

2.4.3 Adequate Mechanical Stimulation and Actuator Types

Adequate stimulation of mechanoreceptors through mechanical stimuli can be achieved either through quasi-static stimuli using solenoids or servo motors (e.g., braille displays for visually impaired individuals) or through vibrotactile stimulation using various actuator types. The literature distinguishes between three general types and five subvariants of vibrotactile actuators [28] (exemplary physical appearance shown in Figure 2.5):




Figure 2.5.: Overview of different actuator types. 1 - solenoid, 2a - generic voice coil actuator, 2b - bone conduction speaker, 3a linear resonant actuator, 3b - Lofelt actuator [125], 4 - 8 and 12 mm eccentric rotating mass actuators, 5 - two different sizes of piezoelectric actuators.

- **Linear Electromagnetic Actuators – Solenoids:**
 - Advantages: Easy to operate electrically, can produce strong forces at low frequencies.
 - Disadvantages: Solenoids typically produce loud click noises when switching their state; they may heat up during use.

- **Linear Electromagnetic Actuators** – *Generic Voice Coils*:
 - Advantages: Amplitude and frequency can be adjusted separately, quick reactions to changes in the input signal
 - Disadvantages: Usually feature peak amplitudes greater than 500 Hz (outside of peak human mechanoreceptor’s sensitivity), difficult to operate electrically as they need a clean input signal (e.g., sinus wave at a certain frequency).
- **Linear Electromagnetic Actuators** – *Vibrotactile Voice Coils*:
 - Advantages: Amplitude and frequency can be adjusted separately, quick reactions to changes in the input signal.
 - Disadvantages: Difficult to operate electrically as they need a clean input signal (e.g., sinus wave at a certain frequency).
- **Rotary Electromagnetic Actuators** – *Eccentric Rotating Mass (ERM) Actuators*:
 - Advantages: Cheap, easy to operate electrically, can produce stronger amplitudes in relation to volume and weight compared to the other actuator types.
 - Disadvantages: Amplitude and frequency are linked and cannot be changed separately, relatively slow to react to changes in voltage.
- **Nonelectromagnetic Actuators** – *Piezoelectric Actuators*:
 - Advantages: Amplitude and frequency can be adjusted separately, quick reactions to changes in the input signal.
 - Disadvantages: Require high operating voltages (~100 V), require a large number of layers stacked on top of each other for produce large enough displacements to be picked up by human mechanoreceptors.

Linear Electromagnetic Actuators

The electromagnet is the most commonly used physical phenomenon to generate vibrotactile stimuli: An electrically conductive wire (e.g., copper) is covered with an electrically insulating material (e.g., lacquer) and wound into a continuous coil. If a constant electric current is allowed to flow through the wound wire, a uniform magnetic field is created that is strongest inside the coil. A piece of ferromagnetic material (e.g., steel) placed near the energized coil is physically pulled in the direction of the electromagnetic field. The force disappears once the current is turned off. A permanent magnet (which has its own magnetic field) is either attracted or repelled by the energized coil, depending on its physical orientation and the direction in which the current flows through the wire. Applying an oscillating electric current to the coil creates a magnetic field that changes over time, providing a simple method for

creating tactile vibrations. This physical principle is also used in audio speakers to create broad-frequency air pressure fluctuations that humans perceive as music or sounds. By modulating the electric current's amplitude and frequency, both amplitude and frequency can be controlled individually, generating highly customizable tactile effects [28].

Subcategories of Linear Electromagnetic Actuators include solenoids (a movable piece of ferromagnetic material inside a coil), voice coils (a movable permanent magnet inside a coil), and specialized vibrotactile voice coils, which are like voice coils but the amplitude to frequency curve is tuned to provide the most substantial amplitudes at frequencies which human mechanoreceptors can best perceive (e.g., Linear Resonant Actuators (LRAs) or the Lofelt actuator [125]).

Rotary Electromagnetic Actuators

Rotating direct current (DC) motors are designed to rotate continuously when a constant voltage or current is applied; therefore, their internal structure is more complicated than that of a solenoid or voice coil. An off-center mass is often attached to the output shaft so that its rotation exerts large radial forces on the motor body. Thus, these actuators are also called Eccentric Rotating Mass (ERM) actuators. This type of actuator can produce oscillatory sensations when driven with a steady voltage, which significantly decreases electrical control complexity compared to other actuator types. However, this design essentially couples both the frequency and amplitude of the resulting vibration to the motor's rotational velocity (in cycles per second or hertz), which is a significant disadvantage compared to voice coils as it disallows producing vibrations with any combination of frequency and amplitude. Besides, internal static friction typically prevents such motors from rotating when the applied voltage is below a certain threshold. They also take a short time to start rotating, causing a delay in starting the vibrotactile cue (typically around 75 ms for smaller actuators and up to 250 ms for larger actuators) [28].

Nonelectromagnetic Actuators

Another approach to generating vibrotactile sensations uses the piezoelectric effect, in which certain solid materials change shape when subjected to an electrical voltage. Vibrotactile display applications typically use multilayer ceramic piezoelectric actuators shaped like a disk or bar as a single layer only changes thickness in the order of a few μm . Such actuators respond very quickly to applied inputs. They can output arbitrary waveforms, but they also typically require inputs on the order of 100 V, making system integration challenging and potentially unsafe for use on humans [28].

Some other technologies, such as shape memory alloy (SMA), may also be used for vibrotactile rendering. SMA actuators are metals that remember their initial shapes and alter their mechanical characteristics according to temperature changes. The effect relies on reversible phase variations in the solid metal alloy. Although SMA actuators may be small and typically have a favorable power-to-weight ratio, they are also known to have high energy consumption, slow reaction times to changes in voltage, and unsafe temperatures for use on the human body while changing shape [28].

2.5 Vibrotactile Patterns: Terminology and Characteristics

(Vibro-)tactile Patterns (TPs) are essentially a set of carefully designed commands for one or multiple actuators over time. Predefined or dynamic tactile patterns can be played back on (vibro-) tactile displays. A simple example of a tactile pattern is a tactile phone notification pattern that turns an actuator on and off rapidly.

The actuator type used for the primary hardware prototypes of the research presented in this thesis, Eccentric rotating mass (ERM) actuators (e.g., [164]), can essentially only be turned on ($\sim 0.5 \text{ V} < \text{supply voltage} < 3.6 \text{ V}$) and off ($\text{supply voltage} < \sim 0.5 \text{ V}$). The input voltage can also be supplied by modulating a 5 V input signal with pulse-width modulation (PWM), which is the usual approach when working with ERMs. Specific actuator characteristics such as spin-on-time and full-stop-time determine what kind of stimuli can be created and how they feel.

When talking about modulating the input signal to create a vibrotactile pattern, we use the following terms:

- **intensity** – modulating the input signal at a frequency above 1 kHz. This is usually the PWM frequency and leads to decreased perceived intensity and amplitude when decreasing the PWM signal’s on-off-ratio. The PWM frequency should ideally be higher than 20 KHz to prevent perceivable auditory noise.
- **roughness** – modulating the input signal between 12 and 50 Hz. This leads to the stimulus feeling “rough” and uneven, especially when using square waveforms.
- **rhythm** – modulating the input signal at a frequency of less than 12 Hz. This can create special rhythms such as 500 ms on, 500 ms off, repeatedly for a defined instruction or meaning.

It is possible to chain multiple actuators together and thus modify the **location** of the stimulus. This chaining can result in a *static pattern with a static stimulus location* – e.g., three motors vibrate together on the left side of the head. It may also result in a

static pattern with a dynamic stimulus location where multiple actuators work together to create a moving tactile stimulus – e.g., through smoothly interpolating between four actuators from the backside of the head over the top to the forehead using an algorithm such as Tactile Brush [80]. For simplicity, we will refer to both of these kinds of TPs as *static patterns* as they do not change based on the environment. Modifying the location of a tactile stimulus on the head has a significantly higher positive impact on recognition performance than modifying the rhythm or intensity/roughness of standard ERM actuators (see chapter 5).

Finally, it is also possible to create dynamic stimulus location patterns that change and react to, e.g., a target in 3D space around the user, indicating the user's target location. An example of this is the target acquisition task in chapter 8. We will refer to these patterns as *dynamic* or *continuous patterns*, as they change based on the environment.

The following section introduces various existing tactile displays for various use cases from prior work with particular regard to navigation and guidance use cases and tactile displays on the head.

2.6 Displays for Presenting Vibrotactile Stimuli

Apart from the regular tactile feedback known by the general population (e.g., tactile smartphone notifications and game controller vibrations), researchers invented a large variety of tactile displays for various use cases, including situational awareness, navigation and guidance, vision substitution, obstacle avoidance, notification, target acquisition, and others.

Early work on tactile displays appeared in 1957 by Geldard et al. [54] who ran experiments on vibrotactile intensity and temporal discrimination and was neatly summarized alongside newer work and general guidelines by Jones and Sarter [82]. Their research review in the area concludes that different levels of vibrotactile intensity and frequency are hard to distinguish and even interfere with each other. At the same time, stimulus location and duration are easier to identify and can thus achieve a higher bandwidth of communicated information.

A large variety of tactile feedback systems appeared after the initial steps in this domain, summarized in the subsections below:

- vibrotactile belts: [30, 48, 49, 52, 71, 81, 151, 158, 206],
- vibrotactile systems on the arm, wrist, or hand: [2, 18, 20, 62, 122, 152, 160],
- vibrotactile systems on the head and neck: [16, 26, 34, 43, 99, 104, 105, 129, 132, 150, 178, 221],
- vibrotactile systems on the back: [9, 79, 80, 183, 202],

- vibrotactile systems on the foot: [133, 181],
- full body suits: [17, 38, 47, 114, 124].

Most of these systems have in common that they use simple ERM actuators in simple configurations (e.g., a ring or grid). The advantages of LRAs over ERMs in terms of being able to control the frequency and amplitude separately and the generally lower response times are usually not deemed important enough by researchers developing systems for body parts other than the feet, palm, or fingertips. A possible reason for this is likely that mechanoreceptors' density on the feet, fingertips, and palm is significantly higher than on any other body part. Thus, highly specialized tactile feedback with varying frequencies and intensities is best presented at body locations with high mechanoreceptor densities for acceptable recognition performance.

On skin areas with lower mechanoreceptor densities, it makes more sense not to modulate the frequency and intensity in too many increments but instead increase the number of actuators and spread them over a larger skin area. This approach allows using spatial location instead of frequency or intensity as a modulator of tactile patterns. Spatial location was proven to have a higher positive impact on TP recognition accuracy than intensity, and frequency (see [82] and chapter 5). Thus, in body areas with low mechanoreceptor densities, ERM actuators' advantages over vibrotactile voice coils generally outweigh the negatives. Thus, the system developed in this thesis also uses standard ERM actuators for the main prototypes used in the user studies. Nevertheless, one of the prototypes presented in this thesis uses specialized vibrotactile voice coil actuators [125] with the goal of lowering audible noise through using lower frequencies at higher amplitudes. However, this prototype was not used in a user study due to the actuators' clunkiness, power drain, and actuator noise issues resulting from unclean sinus input signals (see subsection 4.3.2).

2.6.1 Indicating Direction via Vibrotactile Displays

Several recent works have focused on indicating direction via a variety of vibrotactile outputs. The most promising systems use vibrotactile feedback on the feet, wrist, waist, neck, or head, and usually focus on macro-navigation instructions (e.g. "turn left at the next intersection"), while we focus on precise guidance and micro-navigation in multiple chapters of this work. Thus, the related works presented below are fundamental to this thesis. Systems explicitly operating on the head are presented in the following subsection.

Foot *Shoe me the way* is a tactile foot-based macro-navigation system for eyes-free urban navigation by Schirmer et al. [181]. They developed simple vibrotactile patterns

for the navigation instructions “left”, “right”, and “behind”. Their experiment showed the feasibility of their system and that users enjoyed the way of navigating.

In a comparison experiment, Meier et al. [133] compared three different setups of tactile systems on the foot to a tactile wristband and a tactile belt for macro-navigation in a city. They found that the tactile feedback on the foot provided the most favorable recognition rates of tactile patterns while walking. However, they also conclude that a pedestrian navigation system based on their tactile feedback for guidance might not be sufficient in an urban context with complex geographical situations and should thus be combined with other means of guidance. In contrast, the micro-navigation system for visually impaired people we present in chapter 9 can micro-navigate users through complex geographical situations on its own.

Hand A hand-based tactile system presented by Lehtinen et al. [122] allows guiding users in simple 2D display search tasks (e.g., to find a single letter or word in a large text). They report that their participants in a user study found targets significantly faster with tactile feedback than visual feedback alone. Bial et al. [18] proposed using tactile actuators mounted inside a glove or directly on the fingertips for macro-navigation while driving a car or motorcycle. Recently, Guenther et al. investigated 3D guidance via a vibrotactile glove with a total of 10 spatially positioned actuators [62]. They found that using a higher number of actuators reduced the target acquisition time in their experiment.

Wrist Paneels et al. [152] investigated tactile patterns on a bracelet for indicating directions in macro-navigation. They found that static patterns with a static stimulus location are not well recognized due to the actuators being too close and being recognized as one impulse instead of multiple impulses (funneling illusion). In contrast, static patterns with a dynamic stimulus location are recognized with higher accuracy. We use mostly static patterns with a dynamic stimulus location or fully dynamic patterns in this thesis. Paneels et al. also conducted a vibrotactile pattern discrimination test for their wristband and found detection accuracies for four directions (and three other meanings) after training to be 66 % (and 73 % in a second iteration after further training).

Waist *ActiveBelt* is a vibrotactile belt for directional macro-navigation [206]. This belt was proposed for various use cases such as macro-navigation in a city or notifications of valuables left behind. In a study, participants had to discriminate between the 8 actuators on the belt. Five of the six study participants answered that they could easily discriminate between the actuators. They also report that participants often failed to recognize vibration with a pulse length of less than 500 ms when walking. One of the patterns we use in chapter 9 has a pulse length of only 75 ms (“ATTENTION”)

and is still recognized with near-perfect accuracy, which indicates that vibrations on the head are perceived much more strongly and are more “present” to the user than vibrations on the waist.

Van Erp [49] suggested using spatial vibrotactile cues for navigation directions and tested spatial accuracy with a vibrotactile display mounted on different locations around the torso. He found that the spatial accuracy of vibrotactile stimuli is best in the front-sagittal region with a standard deviation of 4-8° while it is much worse in the other regions with standard deviations around 10-18°. Heuten et al. [71] presented a 6-actuator vibrotactile belt with smooth in-between actuator interpolation to indicate high-resolution walking directions. They found a total average deviation to the indicated angle of 15°, which is comparable to [49]. Using GPS, they performed a continuous macro-navigation task and found their users to deviate 6.6 m from the optimal path on average. The large deviations can be attributed to the inaccuracies caused by GPS and to the relatively simple navigation algorithm. Cosgun et al. [30] used a vibrotactile belt with eight actuators to present simple direction and rotation patterns. They reached similar detection accuracy for their static direction patterns with static stimulus locations as [49]. Their static rotational patterns with dynamic stimulus locations have somewhat higher accuracies. Ouyang et al. [151] did a follow-up study on [49] and [71] with a 12-actuator tactile belt, utilizing the tactile funneling illusion. They reported a detection rate of 91 % for a resolution of 7.5°, which is better than both predecessors. The 1D detection accuracies in [49, 71, 151] can be compared to the results of the invisible target finding task in chapter 8 (mean 2D sphere deviation to the target of 2.3°, SD=1.8°).

Neck A vibrotactile collar around the neck for macro-navigation was presented by Schaak et al. [178]. Their experiment showed that the concept worked well for simple turn instructions (e.g., right, front-left). Matsuda et al. [132] developed a vibrotactile collar around the t-shirt seam and showed its ability for 3D spherical guidance by reproducing a study from our work (section 8.3). While the guidance performance in [178] was significantly weaker than the guidance performance of HapticHead (section 8.3), a vibrotactile collar is a less complex system and can more easily be hidden under a shirt or jacket for social acceptability.

2.6.2 Tactile Displays on the Head

This thesis presents an exciting tactile display around the head that offers various new use cases and improves the performance of other tactile displays in several existing use cases. Consequently, this subsection presents related work on similar, head-based tactile interfaces for a variety of purposes. Table 2.3 shows an overview of the systems discussed in this subsection.

Overview of head-based tactile systems

Authors & system	System Type	Use Case	Advantages	Limitations compared to HapticHead
Nukarinen et al. [150]	tactile glasses with 2 actuators	simple macro-navigation	hidden from the public	unsuitable for micro-navigation and precise guidance
Wolf and Kuber [221]	cap with 2 C2 tactors (left and right)	fundamental research on tactile perception		unsuitable for micro-navigation and precise guidance
Oliveira et al. [34]	7 actuators integrated in VR HMDs	guidance to objects in VR	can be fully integrated into HMDs	limited guidance area, slower and less precise for VR guidance
Marquardt et al. [129]	5 actuators integrated in AR HMDs combined with audio feedback	guidance to objects in AR	can be fully integrated into HMDs	limited guidance area, slower
Cassinelli et al. [26] <i>Haptic Radar</i>	multiple infrared sensors and tactile actuators in a ring configuration	spatial awareness of approaching objects		unsuitable for 3D guidance, did not report 2D navigation performance
Berning et al. [16] <i>Proximity-Hat</i>	multiple infrared sensors and pressure actuators in a ring configuration	spatial awareness of approaching objects		unsuitable for 3D guidance, did not report 2D navigation performance
Dobrzynski et al. [43]	ring of 12 vibrotactile actuators around the head	fundamental research on tactile perception	can be hidden under a beanie	unsuitable for 3D guidance, did not report 2D navigation performance
Diener et al. [41] <i>Vibration-Cap</i>	19 tactile actuators spread across the head	fundamental research on tactile perception	fully integrated into a beanie (no chin strap)	unsuitable for <i>intuitive</i> 3D guidance (no feedback below head), did not report 2D navigation performance
Kerdegari et al. [104]	12 ultrasound sensors and 7 vibrotactile actuators on the forehead	wall-following in low-vision scenarios for firefighters	independent of an external tracking system	unsuitable for <i>intuitive</i> 3D guidance, 2D micro-navigation precision substantially lower, paper is unclear about the optimal path

Table 2.3.: Overview of head-based tactile systems, their use cases, and their limitations compared to the HapticHead concept presented in this thesis.

In an experiment comparing simple visual to tactile cues, Nukarinen et al. [150] presented a set of tactile glasses, which can indicate simple left or right tactile navigation commands for macro-navigation. Their results showed that the tactile cues achieved significantly faster reaction times compared to visual text cues. Tactile cueing was also evaluated as less frustrating than visual cueing and can decrease cognitive demand in navigation.

Wolf and Kuber [221] developed a head-mounted display consisting of two C2 tactors on the left and right of the forehead to support situational awareness. They conducted two studies where participants had to interpret different tactile parameters (rhythm, amplitude, frequency, and location) and found that participants had higher error rates

and cognitive workload for sine than for square waves, which is expected as square waves feel more abrupt and are thus easier to identify and otherwise confirmed the conclusions of Jones and Sarter [82] (location and duration/rhythm being easier to identify than frequency or intensity) for forehead-based tactile feedback.

For spatial awareness in 3D VR space, de Jesus Oliveira et al. [34] presented a 7-actuator tactile display to be integrated into VR HMDs. They indicated direction through the location of the vibrotactile stimulus and elevation through its intensity and recorded mean localization errors of 7.7° for their best-performing study condition. A very similar 5-actuator audio-tactile spatial awareness and view management display inside an HMD was presented by Marquardt et al. [129] in 2019, this time for Augmented Reality use cases in combination with a Microsoft HoloLens and auditory feedback elements for increased guidance performance. They report an accuracy of 2° on longitude, 3.6° on latitude, and 7 cm in-depth for their audio-tactile feedback.

Haptic Radar [26] is a ring around the head, consisting of multiple infrared sensors and vibrotactile actuators to give users a “spider-sense” of approaching objects. A similar concept for the same use case is *Proximity Hat* [16], which uses pressure instead of vibrotactile actuators and stimulates other receptors (Merkel disks and touch domes).

Dobrzynski et al. [43] presented a ring of 12 vibrotactile actuators around the head to potentially navigate people with visual impairments and conducted fundamental research on the maximum number of active actuators and the maximum comfortable vibration intensity on the head (see subsection 2.3.8).

VibrationCap by Diener et al. [41] is a concept similar to the previously published HapticHead (see chapter 4) but miniaturized into a beanie and without the chin strap. They evaluated tactile sensitivities of stimuli on the head, confirming the conclusions of Myles et al. [143, 144, 145, 146] and further roughly evaluated tactile localization precision, which was later examined in more detail together with the funneling illusion on the head by us in chapter 3.

A tactile helmet developed by Kerdegari et al. [104] uses twelve ultrasound sensors and seven vibrotactile actuators on the forehead for micro-navigation in terms of following a wall in a low-vision scenario involving firefighters. Their experiment shows a slightly lower route deviation for the vibrotactile modality compared to auditory feedback, highlighting the advantages of using vibrotactile over auditory feedback for micro-navigation. However, this scenario is not directly comparable to our micro-navigation experiments (chapter 9) as participants in [104] were following a wall while potentially touching it, which may increase precision. The paper neither reports whether the wall was touched nor the distance of the optimal path from the wall. The starting positions suggest that the optimal path was very close to the wall (see Fig. 5a in [104]).

The final section of this chapter is dedicated to tactile illusions as many of the tactile displays mentioned above can, sometimes unexpectedly, produce tactile illusions, which may impact their use case.

2.7 Tactile Illusions

Tactile illusions occur when the perception of an object's property by the sense of touch does not appear to match the physical stimulus. They can occur in numerous circumstances and provide insight into the mechanisms underlying haptic sensations. Many illusions can be used or avoided to create efficient haptic display systems or study the human nervous system [165, p. 327].

There are over twenty different tactile illusions [67], but most of these result from the combined input of multiple senses [67, 165]. This thesis is explicitly interested in the two tactile illusions occurring from tactile actuators and tactile perception alone, causing mislocalization of the applied stimuli. Mislocalizing tactile stimuli may cause issues in specific use cases, e.g., guidance stimuli may be mislocalized by users. However, a tactile display may intentionally exploit this mislocalization to provide even more accurate and comfortable guidance.

The following chapter 3 will present our foundational work on the funneling illusion, specifically on the head. HapticHead benefits from the funneling illusion on the head both in terms of user comfort and potentially through increasing guidance precision, even though the latter was not researched in detail as it would require a different guidance algorithm for comparison which does not rely on interpolation between actuators and provides a similar guidance precision.

2.7.1 The Funneling Illusion




Figure 2.6.: The tactile funneling illusion. In case two simultaneous vibrotactile stimuli are presented on the skin, the user may feel only a single stimulus somewhere between the stimuli locations if the two sources are close enough together.

One of the oldest known tactile mislocalization illusions is the funneling illusion, first described by von Békésy (see [67]). It requires two or more simultaneous vibration signals at different skin locations, which are not too far from each other, depending on the body part (e.g., around to 5 cm on the forehead, see chapter 3). If the stimuli are not too far from each other, the user will only feel a single stimulus somewhere between the actuators (see Figure 2.6). The amplitude can also vary between the locations, leading to a perceived stimulus closer to the higher-amplitude actuator(s).

Consequently, the funneling illusion can be exploited twofold. First, to let the user only feel a single stimulus between multiple actuators, potentially increasing guidance precision for tactile displays (see chapter 3). Second, to increase comfort for users as they will only feel a single, continuous stimulus between actuators instead of multiple stimuli with abrupt switches between actuators. This increase in comfort through the funneling illusion even occurs if a tactile pattern purposefully creates a moving sensation between different actuators by interpolating their intensities over time.

Tactile Brush by Israr et al. [80] is an interpolation concept for multiple tactile actuators arranged in a grid in order to purposefully generate a moving tactile funneling illusion, which simulates the feeling of a continuous motion with a single localization point even though multiple actuators are active at a time. This thesis's continuous guidance algorithm is related to the original Tactile Brush algorithm as summarized in section 8.4.

2.7.2 The Cutaneous Rabbit Illusion

Another well-known mislocalization illusion is called the “Cutaneous Rabbit Illusion”, first described by Geldard and Sherrick (see [67]). This illusion requires actuators that can fire extremely short, well-timed impulses (e.g., LRAs) at discrete locations on the skin with, e.g., 10 cm separation on the forearm. Their firing must be spread over time at intervals of only a few milliseconds, typically one burst of five impulses at one location and then another five impulses at the following location and so-on [67]. The user may then perceive a progression of pulses on the path from one location to the other. Thus, this illusion can also be viewed as a kinetic variant of the funneling illusion [67].

Since the cutaneous rabbit illusion requires actuators capable of delivering extremely short yet intense impulses in a few milliseconds, the actuators used in the prototypes of the research presented in this thesis are incapable of producing this illusion. Even the voice coil actuator type used in the low-frequency actuator prototype needs at least 15 ms to finish one cycle at its resonance frequency of 65 Hz (see subsection 4.3.2).

Tactile Localization Accuracy and the Funneling Illusion around the Head

The first subgoal **G1** of this thesis on the way to establish the head as a means for tactile communication is to research tactile sensitivities and relevant tactile illusions on the head to gain an understanding of human tactile perception on the head with implications for tactile user interface design. This chapter provides an in-depth investigation into tactile perception, localization accuracies, and the tactile funneling illusion on the head. It also clears up several misleading conclusions of related work regarding localization performance and the tactile funneling illusion on the head.

As explained in the previous chapter, the vibrotactile funneling illusion is the sensation of a single (non-existing) stimulus somewhere in-between the actual stimulus locations. Its occurrence depends upon body location, the distance between the actuators, signal synchronization, and intensity. Related work has shown that the funneling illusion may occur on the forehead. We were able to reproduce these findings and explored five other regions to get a more complete picture of the funneling illusion's occurrence on the head. Our study results (24 participants) show that the actuator distance, for which the funneling illusion occurs, strongly depends upon the head region. Moreover, we evaluated the centralizing bias (smaller perceived than actual actuator distances) for different head regions, which also showed widely varying characteristics. We computed a detailed heat map of vibrotactile localization accuracies on the head. The results inform the design of future tactile head-mounted displays that aim to support the funneling illusion.

This chapter is based on the CHI' 2020 paper "Design and Evaluation of On-the-Head Spatial Tactile Patterns" [103], written in collaboration with Michael Rohs. The experiment in this chapter used some influences from a study in the Bachelor thesis of Kerem Can Demir [39] but was eventually re-designed by me and conducted with a newly constructed prototype compared to the Bachelor thesis. Benjamin Simon and Kamillo Ferry helped to conduct the experiment in this chapter as assistant researchers.



Figure 3.1.: Measuring the FI and centralizing bias on the head. A study participant showing two perceived actuator locations on the forehead.

3.1 Introduction

Tactile feedback on the head has been explored in detail (e.g., [16, 34, 35, 43, 99, 104]). In particular, Kerdegari et al. [105] investigated a tactile sensation known as the funneling illusion (FI) or phantom sensation on the forehead (see subsection 2.7.1). This phenomenon emerges when multiple vibrotactile actuators are within close proximity of each other on the human skin. Depending on the intensities chosen, the user may only feel a single stimulation point in-between the actuators with a tendency towards the higher-intensity actuator(s) [13, 105]. Kerdegari et al. found that the FI appears when the distance between actuators is less than 5 cm on the forehead and that there is a centralizing bias, where users systematically underestimate the distance between two actuators, even when the FI does not occur [105].

This chapter aims to extend this investigation to five other regions all around the head to get a more complete picture of the conditions for the occurrence of the FI on the head. We aim to answer the following research questions:

- Can the results reported in Kerdegari et al. [105] be reproduced and validated?
- At which actuator distances does the FI occur for different regions on the head?
- What characteristics does the centralizing bias show for distances between 2.5 and 15.0 cm at these regions?
- How well are users able to localize single actuators all around the head?

The results of this chapter can be used in a variety of existing (e.g., [26, 34, 43, 57, 99, 104, 106, 143, 150]) and future works in Virtual and Augmented Reality as they inform the design of any tactile display on the head in terms of required actuator density depending on the task and head region. For example, the work of Dobrzynski et al. [43] who presented a 12-actuator vibrotactile headband for localization, could increase or decrease their actuator density depending on the head region in such a way that the FI is felt by more than half of the users for any location on the headband and further use the resulting FI to make users localize positions in-between the actuators. Another example would be the guidance algorithm presented in section 8.4 which could be enhanced by incorporating knowledge about the localization precision on the different

head regions. Instead of actuating three actuators, it could be modified to stimulate one to four actuators depending on the head region's localization performance. This modification could potentially increase guidance performance by reducing stimulation overload on head regions with low localization accuracy.

3.1.1 Terminology

As the terms *phantom sensation* and *funneling illusion* (FI) are often used interchangeably [105, 153], we feel the need to define these terms and motivate why we settled on using the term *funneling illusion*. A *vibrotactile phantom sensation* is created by presenting multiple vibrotactile stimuli at nearby locations on the skin. If the locations are close enough (depending on body site), the user may perceive only a single sensation somewhere in-between the stimulus locations, depending on actuator intensities [4]. While the sensation that a user feels is called *phantom sensation*, the phenomenon is called *funneling illusion*, which is one of the human sensory illusions [27].

We will consistently use the term *funneling illusion* (FI) [27, 105], as *phantom sensation* is occupied in medical research to refer to phantom sensations and phantom pain in amputated limbs, which is an entirely different area. Thus, when we refer to measuring the FI, we asked our participants whether they felt single or multiple stimuli during a trial.

In line with [105], we define the *tactile midline bias* as the phenomenon where humans perceive tactile stimuli as closer to the mid-sagittal body plane than they actually are in some but not all body regions. The *mid-sagittal body plane* splits the mostly symmetric left and right hemispheres of the human body.

In line with [105], we define the *tactile centralizing bias* as the phenomenon in which humans tend to perceive multiple tactile stimuli as closer together than they actually are, even if the FI did not occur.

3.2 Related Work

The general related work on tactile perception around the head presented in subsection 2.3.8 and the funneling illusion in subsection 2.7.1 also applies to this chapter. Specific related work concerning the FI is presented below.

Alles first investigated the tactile FI through direct skin stimulation in 1970 (he referred to them as phantom sensations) to convey non-audiovisual data to users [4]. In his experiment, he used two vibrotactile actuators on the forearm and upper arm to find vibration amplitude profiles, maintaining even perceptual strengths for several intended locations in-between the actuators. He found that log intensity profiles are

superior to linear profiles for encoding locations between actuators, among several other observations about the FI's nature.

Cha et al. [27] evaluated the perception of smooth motion with two vibration actuators stimulating the forearm to create a 1D moving FI. They varied the distance between the two actuators and the FI's movement speed, i.e., how fast the sensation traveled between actuators. Their goal was to find suitable stimulation parameters for a *smoothly* moving FI. In a follow-up study, Barghout et al. [13] investigated the accuracy of perceiving intermediate locations in-between four actuators placed in a line on the forearm using stationary and moving FIs.

Using three voice-coil actuators on the forearm, Raisamo et al. [167] investigated saltation perception, pleasantness, and the effects of temporal variables. They also published a follow-up work on three stimulation methods (linear amplitude modulation between actuators, saltation implemented as three fast pulses per actuator, and a hybrid version) using the identical three Tacton C2 actuators as in the first work on the forearm [166]. They found the modulation method to be significantly less arousing and more pleasant than the saltation and the hybrid version.

T-hive [171, 226] is a hemispherical device with a total of 13 independent vibrotactile actuators. It uses the vibrotactile FI to display directions on the hands of users. Their study found that participants were able to discriminate between nine illusory locations in-between three actuators with an accuracy of 76.8 %. A later system by the same authors [225] has a 3×4 actuator grid, attached to the back of a smartphone, for the same use case of displaying directions. Study results show an accuracy of 81.2 % for discriminating between nine illusory locations generated by three actuators. Yatani et al. [227] attached five vibration actuators to the back of a smartphone and explored the pattern recognition accuracy for five static and six moving FI patterns, which were easy to discriminate for the participants of their study with an accuracy of 85 % for the six non-static patterns.

Israr et al. presented two works on a moving FI to increase movies and video games' enjoyment. The first work [79] investigated the effects of various parameters (including body site, on the forearm and back) on rendering FIs. *Tactile Brush* [80] is a 2D interpolation concept for multiple tactile actuators arranged in a grid to generate a moving FI purposefully. Recently, the Tactile Brush algorithm was improved by J.Park et al. [154] in terms of similarity to the target trajectory and uniformity of the stroke motion and was tested on the palm. G.Park et al. [153] continued further work on stationary FIs in 2D cases, quantifying information transfer capacity and measuring the accuracy of perceived positions.

3.2.1 Funneling Illusion through Objects and Out-of-Body

Apart from creating a FI through direct skin stimulation, there are also several other works on utilizing the FI through rigid objects (e.g., smartphones) [87, 88, 107, 108, 191, 193, 200] or even out-of-the-body sensations (e.g., feeling something between two fingers or two hands when holding a tablet or smartphone without touching the actuators) [15, 109, 110, 111, 120, 121, 138, 160, 228]. However, these investigations are more distant from this work, as the explored methods are less applicable to the head.

3.2.2 Localization and Funneling Illusion on the Head

The investigations by Kerdegari et al. [105] are closely related to this work, as they explore the FI and the centralizing bias on the forehead. They found that while the FI almost always occurs when the two actuators are 2.5 cm apart on the forehead, it only occurs in 20 % of the trials at a distance of 5 cm. They also reported a centralizing bias, meaning that the perceived distance between two simultaneously active actuators was around 2.5 cm less than the actual distance. Besides, they claim a strong midline bias when locating single actuators on the forehead. However, we believe that the study design regarding the claimed midline bias is problematic. The scale extended from 0 to 17 cm (actuators at 0 to 15 cm), and the participants saw their forehead, the scale, and the actuators in a mirror. Hence through visual feedback, the participants could exclude actuator locations beyond positions 0 and 15 cm, respectively.

In the study reported below, we corrected this issue and explored five additional locations on the head, beyond the forehead. Our study had more than twice the number of participants in order to get more reliable results. Moreover, we give an error estimation for each measurement and discuss a potential bias through the participants' dominant hands when pointing to the locations.

3.3 Tactile Perception Prototype

Our hardware prototype is a reconstructed version of Kerdegari et al.'s [105] prototype. We contacted the authors on their prototype's specifics (e.g., foam material) and built a slightly enhanced but otherwise very similar version. Just like [105], our prototype consists of seven vibrotactile coin actuators (A1-A7) mounted 2.5 cm apart (less than 1 mm tolerance) on a scale that is attached to a stretchable Velcro fastener. The actuators used are 10×3.4 mm coin-style actuators (Precision Microdrives 310-117, frequency 250 Hz at 3.3 V, 1.9 g normalized amplitude). Kerdegari et al. used the PM 310-113, which is a slightly weaker, no longer available predecessor of the PM




Figure 3.2.: Our prototype built after the template of Kerdegari et al. [105]. Seven actuators (A1-A7) are placed at scale positions 11 to 26 cm, with 2.5 cm distance between the centers. Tolerances are less than 1 mm.

310-117. We used a similar 10 mm thick \times 17 mm diameter neoprene polymer, which is naturally vibration absorbent, to isolate the actuators from each other.

Different to [105], we stamped round forms from the neoprene polymer instead of square ones. We also did not use an additional polymer layer apart from the neoprene polymer as the other plastic polymer used in [105] acted as a glue (our actuators came with a self-adhesive surface). We also moved our actuators on the scale so that participants could point to locations up to 4 cm outside the actuator area whereas [105] mounted the first actuator on 0 cm and the last on 15 cm on a 0 to 17 cm scale, making it impossible to point at a location outside that area.

Our prototype's actuators are driven by a Raspberry Pi 3 using the pigpio library [5] and a custom-built actuator driver board. The Raspberry Pi was updated at 100 Hz through Wi-Fi using a Unity v5.6.6f scene [208] for the experiment. For further details on the implementation of this driver board, we refer to subsection 4.2.2.

3.4 Tactile Perception Experiment

Based on our prototype, constructed after the example of [105], we designed an experiment to answer the research questions posed in the introduction. In addition, we formulated the following hypotheses for the experiment:

- H1** We expect our results for the forehead to be comparable to Fig. 4 left in [105], in terms of localization performance of single actuators.
- H2** We expect a significantly smaller midline bias for the forehead compared to Fig. 4 right in [105], as we use a scale and study design that does not limit participants to show locations only between 0 to 15 cm.
- H3** We expect a similar occurrence of the FI compared to Fig. 5 left in [105] for the distances 0 to 10 cm.
- H4** We expect a similar occurrence of the centralizing bias compared to Fig. 5 right in [105] for the distances 2.5 to 10 cm.

- H5 We expect the localization accuracy of single actuators for the forehead to be significantly better than for any other head region, as Myles et al. found the forehead to have the smallest vibration perception threshold [143, 144, 145].
- H6 We expect the FI to appear at larger distances when comparing the forehead to any other head region, for the same reasons as given in H5.
- H7 We expect the centralizing bias to be larger, the larger the maximum distance for the FI for a given head region, as the data in [105] suggests a correlation between these two parameters.

We invited 24 participants with technical backgrounds from around our university (22 male, 2 female, mean age 24.5 years, SD = 5.8 y) for the experiment. The experiment took 46 minutes on average per participant, excluding filling out questionnaires and the introduction. One of the participants (P5) was left-handed, and another one (P10) ambidextrous.

3.4.1 Tactile Perception Experiment – Design

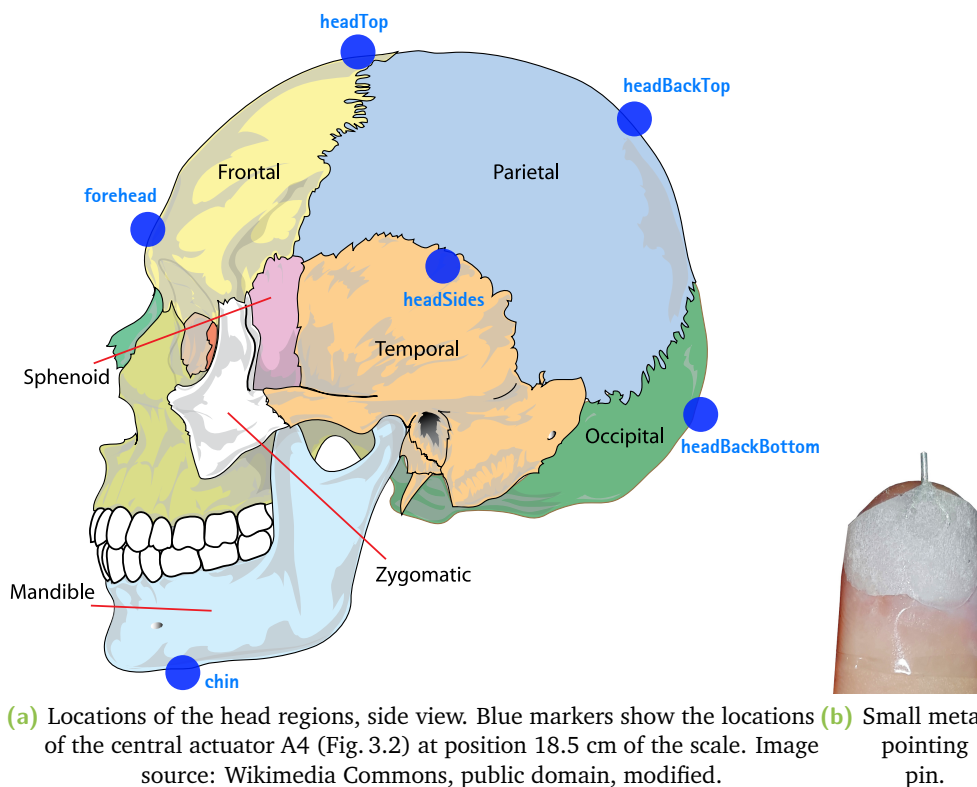


Figure 3.3.: Locations of the selected head regions and exemplary small metal pointing pin.

Since a major goal of this experiment is to map the entire head in terms of the localization precision of single actuators and the occurrence of the FI for different



Figure 3.4.: Our prototype on different head regions for the user study. It was always attached in the same orientations, as shown in this figure for the different head regions so that the scale always went from low to high from the perspective of the experimenter to prevent reading errors.

distances between actuators, we chose a total of six head regions for this evaluation, as shown in Figures 3.3a and 3.4.

Counterbalancing on the six head regions was applied through a balanced Latin square. For every head region, there were a total of 42 trials shuffled randomly to prevent order effects. These 42 trials consisted of single actuations for each of the seven actuators (repeated three times) and all 21 possible combinations between pairs of two actuators (2.5 to 15.0 cm distance between the actuators). The head's side to which the actuators were applied (right or left) was counterbalanced between participants. This was inverted after every six participants as we would otherwise generate order effects due to using a Latin-square-6. We made sure not to introduce confounds between the head side assignment and the head regions. To keep the counterbalancing fully operational, we had to recruit a multiple of 12 participants.

The participants wore Sony WH-1000XM3 headphones playing white noise at 74 dB during the experiment to mask actuator noise. The participant pointed to the single location or the two locations at which they felt a stimulus. Pins were mounted on their index fingers to improve pointing precision (see Fig. 3.3b). The experimenter read the positions on the scale and entered them into the study app (see Fig. 3.5). The reading accuracy is expected to be about ± 0.5 mm. We also argue that this is a more realistic




Figure 3.5.: The experiment UI was used by the experimenter to guide a participant through the study and record data. The two vertical lines can be moved by clicking and dragging them along the scale.

use case as a tactile system should be able to give feedback without the user seeing a visual representation of the system while using it (Kerdegari et al. [105] used a mirror and the participants saw themselves in the mirror pointing at locations). We did not mention or explain the FI effect to the participants, so they were not aware of what we were looking to prove in the experiment.

3.4.2 Tactile Perception Experiment – Procedure

The participants filled out an informed consent form, an optional photographic-release form, and an introductory questionnaire. They were subsequently made familiar with the prototype before putting it on at the first of six possible head locations. We ensured a similarly tight but comfortable fit for all our participants by adjusting the straps ourselves to a pressure of about 4.5 N when pulling the center of the prototype strap 1 cm away from the participant using a BaseTech HS-11 scale. Before starting the trials for a head region, all vibration actuators were tested one after another on proper functioning. Finally, they put on headphones that played white noise.

Each individual trial consisted of the participant signaling readiness through a hand gesture. The experimenter verified that the participant did not touch the scale with his or her fingers and then pressed a start button, upon which a 1 s vibrotactile stimulation played on either one or two actuators at full intensity. The participant then pointed to one or two locations on the scale with their index fingers and the attached pins, depending on the number of perceived stimuli. If desired, a trial could be repeated, which happened only in 0.4 % of the trials. For each trial, the system logged a timestamp, the participant id, the head site, the active actuator(s), the perceived location(s) (read and entered by the experimenter), and the repetition count. Most of the participants performed the trials while having their eyes closed as per our recommendation, so they were less distracted by the environment.

In the end, every participant filled out a final questionnaire with subjective questions on the experience of tactile feedback on the head. The prototype and earpieces of the headphones were disinfected between participants for hygienic reasons.

We implemented this study design and procedure into a Unity [208] application as seen in Fig. 3.5. When designing this experiment’s user interface, we made sure that only the correct procedure through the experiment could be followed and that entering a location could not be forgotten due to being reminded by the test application.

3.5 Tactile Perception Experiment – Results

Analyzing the results, we first take a look at the quantitative results of the single-actuator trials. Second, we focus on the localization performance and occurrence of the FI in the multi-actuator trials. Third, we discuss subjective results and feedback of our participants.

To isolate effects caused by possibly larger errors when pointing to locations on the left side of the head (side of the non-dominant hand) using the right hand (dominant hand) and to ward against a possible influence on the midline bias effect, we reordered all data so that actuator A1 is always on the left of the head and A7 on the right for the evaluations (see Figures 3.2, 3.3a and 3.4 for actuator locations). We chose to flip the data for the left-handed participant P5 and leave it unchanged for our ambidextrous participant P10 since an evaluation of left vs. right-handed individuals makes no sense with only a single left-handed person in the experiment.

In an ideal case, without the influence of the dominant hand and noise, data for symmetric head regions should be symmetric for the sets of opposing actuators (e.g., A2, A6; if participants point X mm to the left for A2, they should also point X mm to the right for A6). Realistically, however, an influence from pointing with the dominant hand is expected, and thus we provide an evaluation on how well single actuators were localized on the left and right sides of the head (see Table 3.2).

For the head side locations, we oriented the data so that A1 is always located towards the front and A7 towards the head’s back. Furthermore, we merged the data of the left and right head sides into a *headSides* location (see Fig. 3.3a, blue labels) for most evaluations.

3.5.1 Quantitative Results – Single Actuator Trials

Table 3.1 shows the mean absolute deviations for each of the head regions. The forehead offers the most precise localization performance while all other head regions were much less precise. A one-way ANOVA with Holm-Bonferroni corrected comparisons showed that the forehead location was significantly more precise than all other head regions ($F_{5,162} = 15.44, p < 0.0001$, confirms hypothesis H5) and all other head region combinations were also significantly different from each other ($p < 0.05$)

Head region	Mean absolute deviation of all actuators [cm]	SD [cm]	Misclassification as multiple actuators [%]
forehead	0.72	0.56	0.20
headTop	0.94	0.81	1.79
chin	1.17	0.89	2.38
headBackBottom	1.21	0.92	5.56
headBackTop	1.40	1.12	6.15
headSides	1.50	1.25	2.78
mean	1.16	0.93	3.14
headSide(left)	1.53	1.24	2.38
headSide(right)	1.47	1.26	3.17

Table 3.1.: Mean absolute localization accuracy and misclassification of single actuators for different head regions (Fig. 3.3a, blue labels).

Head region	Mean deviation A1&A2&A3 (left) [cm]	SD [cm]	Mean deviation A5&A6&A7 (right) [cm]	SD [cm]
forehead	0.14	0.85	0.17	1.00
headTop	-0.02	1.15	0.00	1.09
chin *	1.11	1.23	-0.85	1.15
headBackBottom	-0.15	1.46	0.10	1.64
headBackTop *	-0.81	1.60	-0.04	1.52
mean (symmetric HR)	0.05	1.26	-0.12	1.28
headSides	0.51	2.04	-0.06	1.73
headSide(left)	0.38	2.20	0.38	1.45
headSide(right) *	0.65	1.87	-0.50	1.86

Table 3.2.: Midline bias for different head regions. A positive value for the left actuators (negative for right) shows the average bias towards the midline (red background). A blue background shows a bias away from the midline. Head regions with significant differences between the groups of left and right actuators are marked with a star. For both headSide regions, A1 is on the front and A7 on the back of the head.

in terms of localization performance except for the following combinations: chin–(headTop,headBackTop,headBackBottom), headBackBottom–(headBackTop,headSides), and headBackTop–headSides (Fig. 3.3a shows the locations). Furthermore, Table 3.1 also depicts the percentage of misclassifying a single actuator as multiple stimulation locations. If this happened, we took the midpoint between the two locations for all further evaluations.

Table 3.2 shows the midline bias of different head regions. Multiple one-way ANOVAs comparing the left actuators (A1,A2,A3) vs. the right actuators (A5,A6,A7) for every single head region found significant effects of actuator group on deviation for the chin ($F_{1,70} = 90.3$, $p < 0.0001$), headBackTop ($F_{1,70} = 8.67$, $p < 0.005$), and headSide(right) ($F_{1,34} = 6.64$, $p < 0.05$). The differences for all other head regions were not significant ($p > 0.05$).

Fig. 3.6 shows a heat map of absolute localization accuracies. We merged the data from the symmetric head regions so that, e.g., A1 and A7 show the average of those

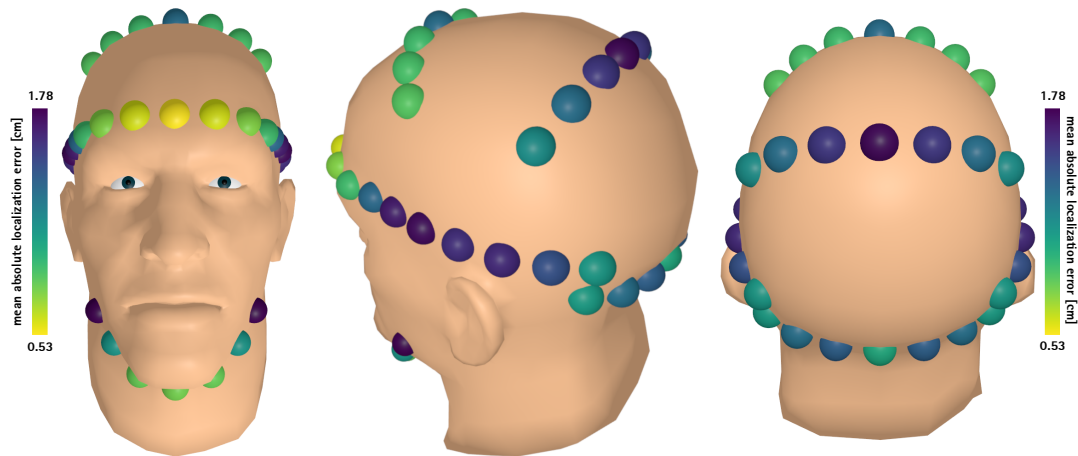


Figure 3.6.: Heat map of mean absolute localization accuracies for all single actuator locations. From left to right: frontal, left, and back view of the model head. The data from symmetric head regions is merged to reduce noise. The color scale is viridis (a perceptually uniform color scale), ranging from 0.53 cm as the minimum error on the forehead to 1.78 cm on the headSides.

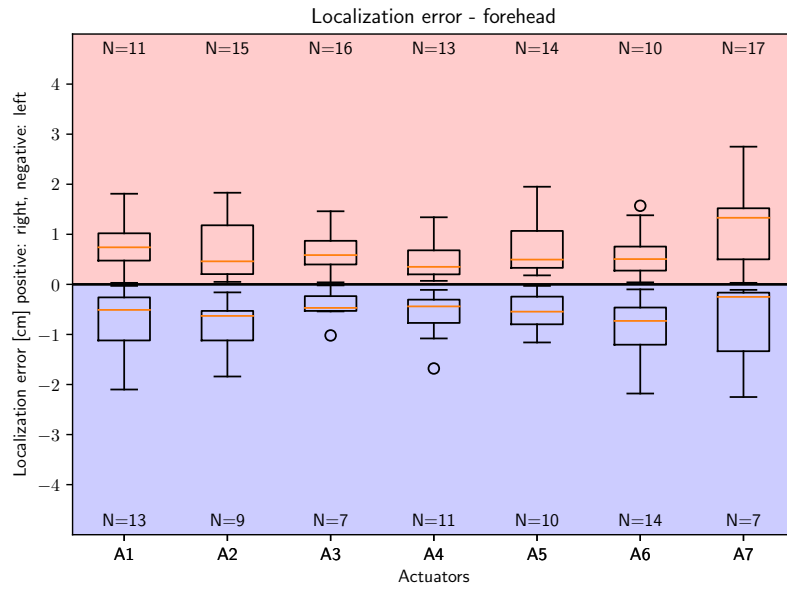
two actuators to reduce noise. For the headSide regions, we merged the data of the individual actuators of both sides.

Fig. 3.7 presents boxplots of the left and right deviations of actuators on the forehead. This Figure is directly comparable to Fig. 4 right in [105], but shows very different data. We suspect that this results from a methodological error in the study design of [105], as discussed below. We decided only to highlight the detailed evaluation of left and right deviations for the forehead and the chin, as these are the two extremes of having no (or minor, slight negative) midline bias (forehead, Fig. 3.7a) and the highest midline bias (chin, Fig. 3.7b).

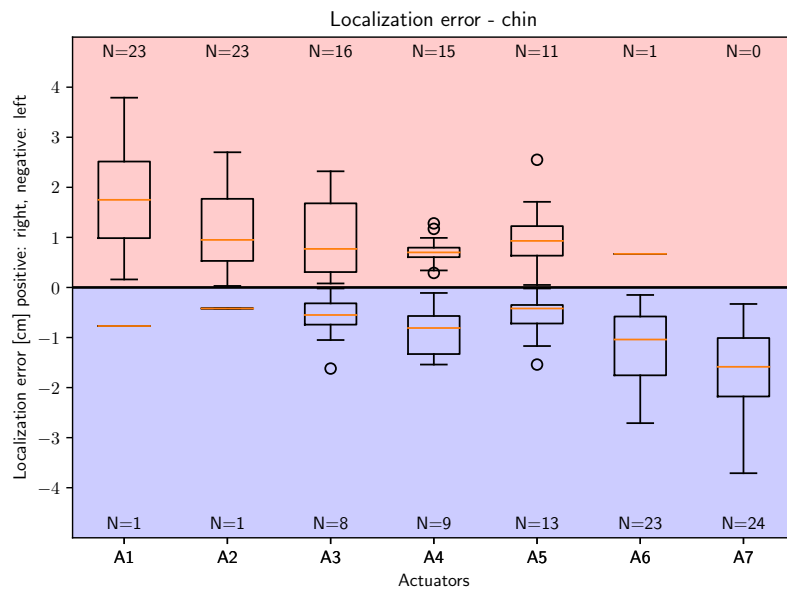
3.5.2 Quantitative Results – Multi Actuator Trials

This section focuses on the localization performance and occurrence of the FI in the multi-actuator trials. Fig. 3.8 shows the FI's occurrence by distance for all head regions. The error bars on these bar charts represent the standard deviation. In general, the different head regions feature very different distances for which the FI still occurs for most users. The threshold at which more than 50 % of users still experience the FI varies between about 5.1 cm for headBackTop to about 7.7 cm for headSide.

The centralizing bias of the different head regions is shown in Fig. 3.9. Error bars represent the standard deviation of all data at specific distances between two locations. Interestingly, there are large biases around 2.5 - 5.0 cm for the head regions chin, headTop and headSide. On the other hand, headBackTop and headBackBottom feature comparatively low biases for most distances.

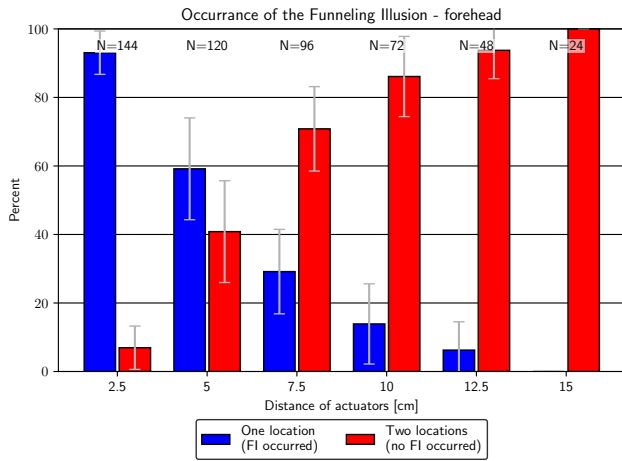


(a) forehead

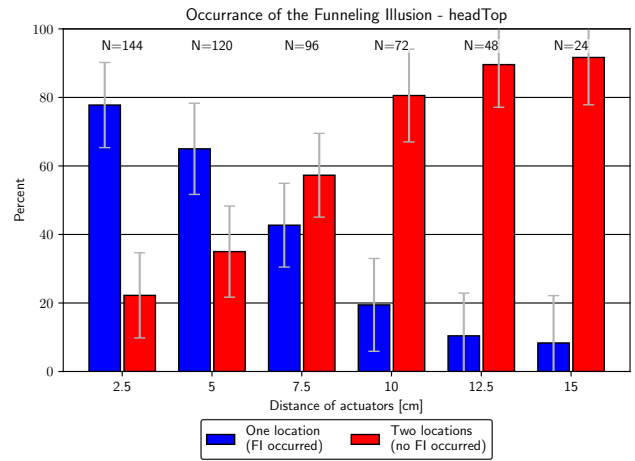


(b) chin

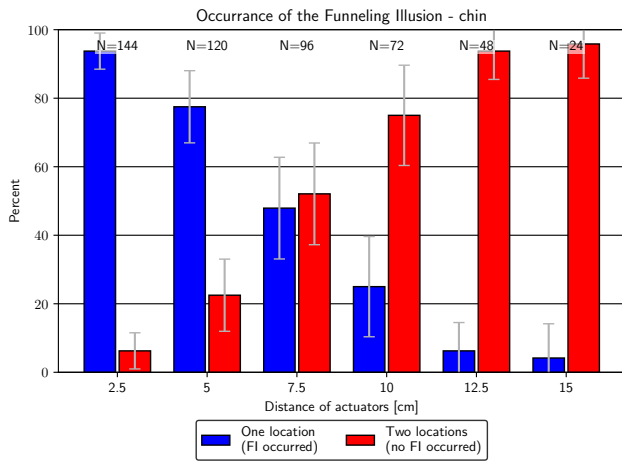
Figure 3.7.: Localization error of each factor on the forehead and chin, comparable to Fig. 4 right in [105]. Deviation towards the left is shown in blue, deviation towards the right in red. If the N for an actuator does not sum up to 24 (number of participants, median of 3 repetitions for each location), this means that the other trials were within ± 0.5 mm of the correct location. A1 is the left-most actuator, and A7 is the right-most actuator.




(a) forehead




(b) headTop




(c) chin



(d) headBackTop

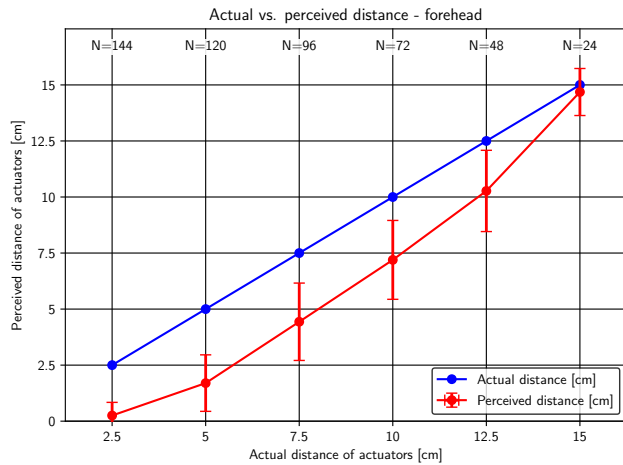


(e) headBackBottom

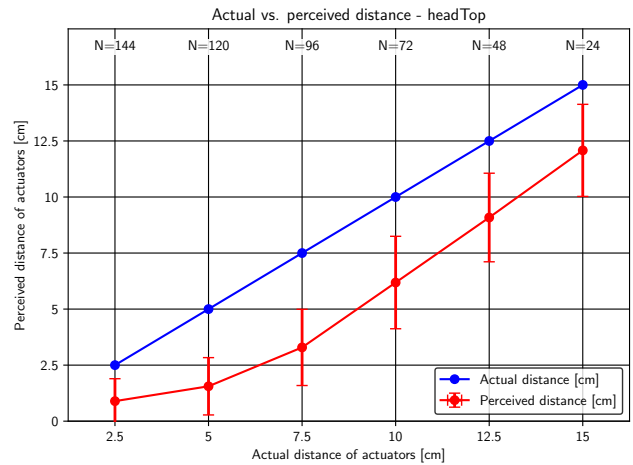


(f) headSides

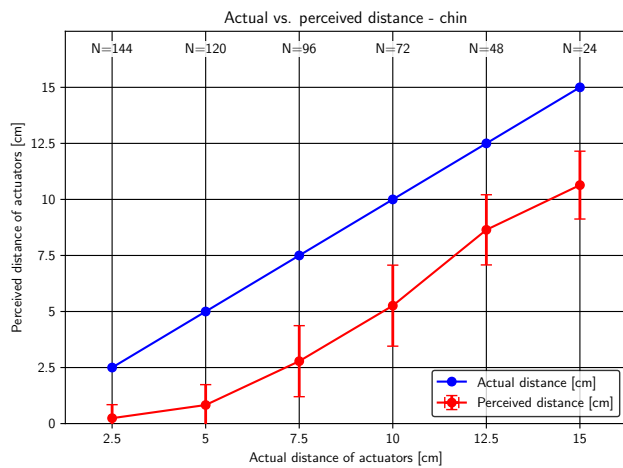
Figure 3.8.: Occurrence frequency of the FI in two-actuator trials for different actuator distances at different head regions. Blue bars show trials in which a participant indicated a single location, red bars show trials in which a participant indicated two locations. Error bars represent the standard deviation between participants.



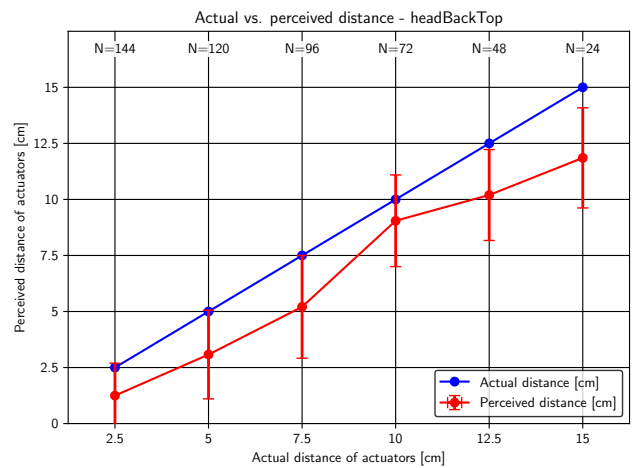
(a) forehead



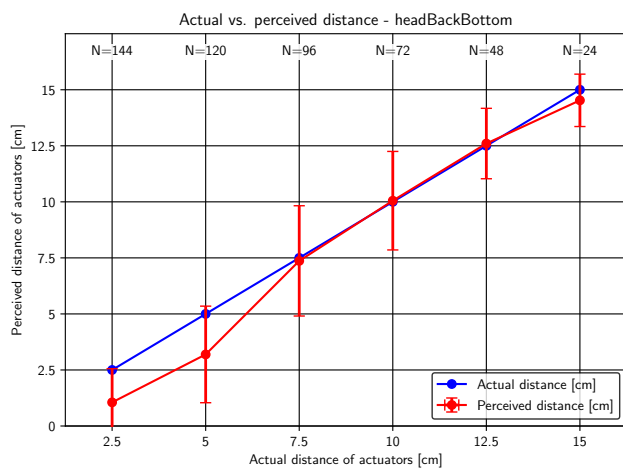
(b) headTop



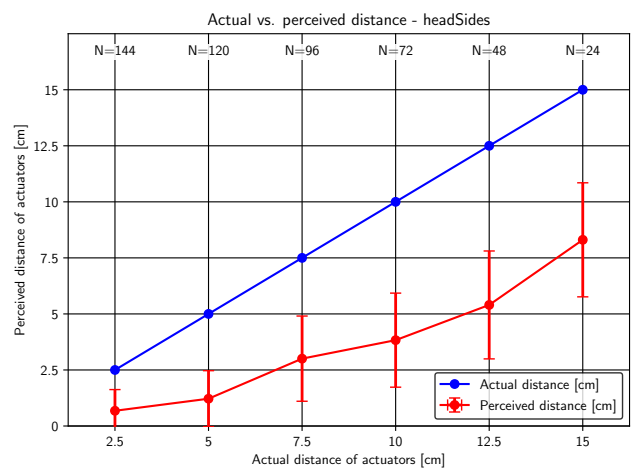
(c) chin



(d) headBackTop



(e) headBackBottom



(f) headSides

Figure 3.9.: Actual and perceived distance between locations with multiple stimulation points. Error bars represent standard deviation between participants. In case a participant indicated only a single location when multiple stimulations were given, a distance of zero is assumed.

3.5.3 Subjective Results

Fig. 3.10 shows the subjectively judged sensitivity. To measure sensitivity, we asked the participants to order the six different head regions by sensitivity (“How sensitive are the different head regions? Please sort them in terms of vibration sensitivity.”). This measurement is influenced by how strongly the actuators (all running at the same intensity) were perceived at different head regions due to more or less hair and nerve density of the skin below. The forehead was judged as most sensitive, followed by the chin. The head sides were judged as least sensitive.




Figure 3.10.: Average subjective head sensitivity rating. Participants were able to sort the head regions by sensitivity from score 6 (highest) to score 1 (lowest).

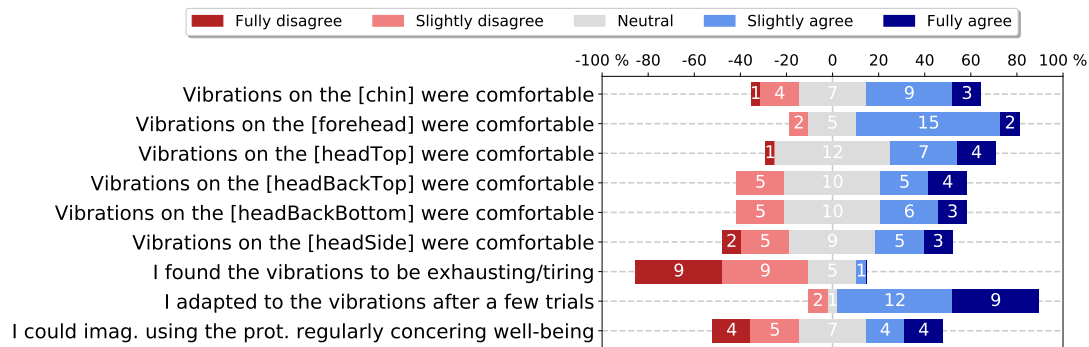


Figure 3.11.: Diverging stacked bar chart of subjective results of our experiment, measured through the final questionnaire.

Fig. 3.11 presents subjective results from the final questionnaire of the experiment. Generally, all head regions were rated as rather comfortable, with the least agreement for headSide. The participants commented that the headSide region was less comfortable because some actuators were close to the ears with some audible noise, despite the white noise being played in the study, and these actuators were generally harder to localize.

In terms of fatigue, the experiment only took 46 minutes on average, and only one out of 24 participants agreed that the vibrations were exhausting, so we do not expect exhaustion or tiredness to influence results. The two participants who indicated that they did not adapt to the vibrations after a few trials commented that the vibration intensity was too high for them, and one found it difficult to relax when concentrating on the localization of the actuators. Participants were split in their opinions on whether they could imagine recurring use of the prototype concerning well-being. The prototype was explicitly designed for the experiment and is rather clunky. For actual use, it would need to be miniaturized and integrated into a VR headset, beanie, or other garments (e.g., [41]). 16 of 24 participants could imagine themselves or other people using a commercial head-based tactile feedback system for various applications. Out of the eight participants who could not imagine themselves or others using a head-based tactile feedback system, six could not think of good use cases, and two thought head-based tactile feedback was uncomfortable in the first place.

We also asked the participants in which scenarios they could imagine using a tactile feedback device on or around the head. They suggested a variety of use cases, such as pedestrian guidance and guidance for the visually impaired, silent notifications, movies and games (VR and AR), head massage and relaxation, firefighting, as well as medical applications.

In other written comments, three participants complained about discomfort when wearing the prototype for a longer time, out of which two specifically mentioned the chin region.

3.6 Tactile Perception Experiment – Discussion

We start the discussion by providing a thorough comparison to [105], first on single actuator localization and midline bias and subsequently on multi-actuator localization and the occurrence of the FI. Furthermore, we compare our work in terms of localization accuracies with [35], and [41].

3.6.1 Single Actuator Localization Performance

When directly comparing the data in Table 3.1 to Kerdegari et al.'s [105] data for the forehead, the midpoint of their range of mean deviations from 0.51 to 0.76 cm per actuator is slightly lower than ours at 0.53 cm to 0.9 cm, mean 0.72 cm. However, since our average still falls within the range, we can accept H1.

With this in mind, we found large differences in localization accuracy for the different head regions (Table 3.1). Localization was most accurate on the forehead, with

the closest contestant (headTop) already much less accurate. Thus we can accept hypothesis H5. Also, the accuracy of headSide was less than half of the accuracy of the forehead, so we recommend avoiding the region close to the ears in tactile system design, not just from a noise perspective but also because localization performance is much worse than for other head regions.

3.6.2 Heat Map of Localization Accuracies

Fig. 3.6 shows a heat map of localization accuracies for all head regions covered in this chapter. Diener et al. [41] created an “accuracy score” heat map for 18 different head regions, excluding the chin. However, they only provide a three-shaded color scale representing a rather coarse accuracy score. Our heat map is more detailed, provides a continuous color scale, and is based on more data: 6 head regions \times 7 actuators \times 24 participants \times 3 repetitions = 3024 trials. [41] is based on 20 participants \times 60 trials = 1200 trials. Furthermore, in [41] participants entered actuator locations in a GUI showing a visual representation of a head from above, and participants could only choose between 19 positions. In our study, participants could point to any location on the scale without the possibility of bias through the GUI design.

When comparing our heat map with that of [41], it is apparent that they are very different. Diener et al. measured the lowest accuracies in the frontal region of the head while our corresponding headTop region performed relatively well in localization accuracy. Furthermore, they measured the best localization accuracy on an area corresponding to headBackBottom, while we measured the best accuracy on the forehead. We attribute these differences to the very different prototype and aforementioned study designs and data input methods. For example, participants in Diener et al. [41] had to choose between two different locations on the forehead or three on the headTop region in a visual interface with a representation of the human head from above; thus, the choice was limited and of varying difficulty. Our study, in contrast, offered a seamless choice on the scale for all head regions and actuator locations and thus posed the same difficulty everywhere.

De Jesus Oliveira et al. [35] studied vibration localization accuracies for four head regions: forehead, frontotemporal (overlaps forehead and headSides in our study), temporal (same as headSides), and occipital (overlaps headBackBottom). Their results mostly agree with what we found. However, we cannot confirm their predictive model of acuity on the head as a function of skin type and distance of stimuli from the head midline. This model does not work for the regions headBackTop and headTop, which were not studied in [35]. These seem to be inverted compared to other head regions in that they show the lowest localization accuracy in the center (see Fig. 3.6).

3.6.3 Midline Bias

A midline bias was noticed by Kerdegari et al. for the forehead (Fig. 4 right in [105]). However, due to a supposed methodological flaw in their study design, the forehead results appear biased and different from ours. In [105], actuator 1 was located at position 0 cm, and the scale ended there, whereas actuator 7 was located at position 15 cm (the scale went on to 17 cm). The participants indicated the perceived stimulus positions themselves in front of a mirror, knowing that the first actuator was at 0 cm and the last actuator at 15 cm. They could not indicate locations less than 0 cm. Therefore, the participants were biased not to indicate locations outside the range 0 to 15 cm, even if they felt a stimulus there. Thus, their conclusion that there is a strong bias towards the forehead midline is invalid, as evident from our corrected study design and Table 3.2 and Fig. 3.7a. These results suggest that we can accept hypothesis H2.

However, we found significant midline biases for the chin and supposedly for headSide(right). For the latter, this bias is most likely caused by pointing difficulties as headSide(left) features almost no midline bias but appears shifted towards the head's back instead (likely also due to pointing issues). The midline bias for the chin is the strongest by a large margin and can be explained by the jaw bones, which transfer part of the vibration intensity from the outside actuators so that it feels like the vibration point is more towards the midline/chin (see also Fig. 3.7b).

headBackTop features a somewhat inexplicable significant bias towards the outside of the scale but only for the left side (non-dominant hand). The median deviation to the outside (0.68 cm) is not much different from the average (0.81 cm, see Table 3.2), so outliers are not the reason for this bias. We see a possible explanation in localized influences from pointing with the non-dominant hand, but since we did not find an overall dominant hand effect (see below), we will leave confirmation of this hypothesis to future work.

3.6.4 Possible Dominant Hand Effect

To find a possible effect of the dominant hand on localization accuracy on the left and right hemispheres of the head, we further analyzed the data shown in Table 3.2. The absolute of the mean of the symmetric head regions between actuator groups shows a possible but minimal effect, as there is a 0.7 mm difference when it should be equal on both sides. Just taking both sides' absolute values and comparing them against each other would erase midline and outside biases. Instead, we inverted the left actuators' data and compared them to the non-inverted data of the right actuator group. A one-way ANOVA shows that there is no statistically significant difference between the actuator groups for inverted vs. non-inverted deviations on all symmetric

head regions ($F_{1,357} = 0.44, p = 0.51 > 0.05$). Thus, we did not find a significant effect of the dominant hand when averaging over symmetric head regions.

3.6.5 Occurrence of the Funneling Illusion

Regarding the occurrence of the FI, Fig. 3.8a is directly comparable to the results of [105] (Fig. 5 left). We measured a much higher occurrence of the FI for the 5 cm distance on the forehead: 59 % vs. 21 %. As this data point is of more importance than the others because it is around the tipping point where participants either feel the FI or not, we have to reject H3. This result might be explained by the low number of 10 participants in [105], with the different study design, or with slight differences in prototype design. For the other head regions, except for headBackTop and headBackBottom, the threshold distance at which a FI still occurs for most participants is always higher than for the forehead, which is expected because of obstruction through hair. Myles et al. [143, 145] found lower absolute tactile detection thresholds for the other head regions. However, headBackTop and headBackBottom actually feature a slightly lesser occurrence of the FI at 5 cm distance than the forehead. Thus, hypothesis H6 has to be rejected. However, due to the small differences between the head regions and due to large variances between users, this result could also be attributed to noise (see error bars at 5 cm distance in Fig. 3.8a vs. Fig. 3.8d and 3.8e).

Using linear interpolation, we estimate the thresholds for the FI to occur for 50 % of the users on the forehead to be around 5.8 cm, for the chin 7.3 cm, for headTop 6.7 cm, for headBackTop 5.1 cm, for headBackBottom 5.2 cm, and headSides around 7.7 cm. These thresholds apply only to prototypes constructed similarly with appropriate vibration insulation between actuators using, e.g., neoprene polymer. Even with a slightly different prototype design, the thresholds will vary, and they will likely decrease the better the insulation between the actuators is. Kerdegari et al. [105] estimated around 3.95 cm for the forehead. However, this result is again most likely influenced by their study design.

3.6.6 Occurrence of the Centralizing Bias

Concerning the centralizing bias, the forehead region shown in Fig. 3.9a is again directly comparable to Fig. 5 right in [105]. We found very similar centralizing biases for distances 2.5 to 10.0 cm; thus, hypothesis H4 can be accepted. For the other head regions, we measured very different centralizing biases (Fig. 3.9). In particular, the headSides regions seem to have the largest centralizing bias, especially at distances 10-15 cm along with the FI occurring for more users even at these distances compared to the other head regions. This can be explained by the closeness of these head regions to an ear.

Interestingly, headTop and chin both feature relatively large centralizing biases especially compared to headBackTop and headBackBottom. This is peculiarly interesting because headTop and chin were both more precise than headBackTop and headBackBottom in single actuator localization (see Table 3.1). We are not entirely sure why this phenomenon occurs and leave this to future research.

To determine whether there is a correlation between the occurrence of the FI and the size of the centralizing bias (H7), we averaged the occurrence frequency of the FI and the size of the centralizing bias over all distances for all head regions. A Pearson correlation test shows that these features are indeed correlated ($r = -0.952, p < 0.005$), so hypothesis H7 can be accepted.

3.6.7 Head Sensitivities and Subjective Feedback

As mentioned in the related work section, Myles et al. [143] found the forehead by far as most sensitive, and the occipital (headBackBottom) and temple (slightly overlaps forehead and headSide) regions to be more sensitive to vibration stimulation, with a lower vibration perception threshold compared to other regions. While it seems obvious to compare our work with [143], they had very different research goals. They studied absolute detection thresholds (ADTs) by head region, while this work uses a vibration intensity well above the ADT to measure other parameters. Still, some of our findings seem to be in line with [143], as our participants also rated the forehead as the most sensitive region and forehead is the most accurate region for localization of single stimuli (see Fig. 3.10 and Table 3.1).

The subjective feedback collected by the final questionnaire and verbal comments suggests that tactile feedback on the head is, on average, well-received with minor differences between the head regions. Certain positions should be avoided, however (ears). Also, a per-user calibration of maximum vibration intensities for the different head regions is desirable to deal with possible discomfort experienced by some participants.

3.7 Limitations

This chapter does not evaluate the effect of different hair densities on accuracy. We had participants with very different hair densities, and we chose not to evaluate a possible effect of different hair densities as related work found no significant effect of hair density on localization performance on the head [41]. Furthermore, while [146] did find a significant effect of hair density on vibration perception threshold, as long as a tactile display operates above that threshold, as our prototype did, this should have no substantial influence on localization performance [41].

3.8 Conclusion and Future Work

In conclusion, this chapter provides several contributions to understand vibrotactile feedback on the head further. These findings can serve as the basis for designing future tactile head-based systems:

- Experiment validation and correction of the problematic midline bias conclusion for the forehead in [105].
- Quantitative results on localization precision and midline bias evaluation of single actuators on the head.
- Quantitative results on maximum distances for the FI at different head locations.
- Characterization of the centralizing bias effect for multiple stimuli and different distances on the head.

When designing head-mounted tactile interfaces at individual or multiple head regions, developers must consider several parameters. There are widely differing localization accuracies for single actuators and maximum distances for the occurrence of the FI on different head regions. If the FI shall be utilized (e.g., for precise guidance), developers must be aware of the widely varying centralizing biases for different head regions, which should be considered in the guidance algorithm.

There are many different systems for use cases in Virtual and Augmented Reality that are enabled with a tactile display on or around the head. These are hinted at in the introduction and related work section. Even if they use just a single actuator, all of these systems can benefit from implementing their tactile feedback, according to this work's findings. For example, hardware prototypes [26, 34, 43, 57, 99, 104, 106, 143, 150] and the guidance algorithm by Israr et al. [80] or our guidance algorithm in section 8.4 can be modified to take into account the varying localization performance, centralizing bias, and FI occurrence for different head regions to implement different actuator densities depending on the task and head region (hardware optimizations) and to provide better guidance performance (algorithm optimizations) potentially.

HapticHead – A Tactile User Interface around the Head

The human head is not only naturally perceived as the “center of attention” for humans but is also close to spherical in shape, which allows placing tactile feedback *all around* the head. The latter has the advantage of providing tactile stimuli, which can *intuitively* be mapped to the point perceived in relation to the user’s current position due to natural mapping [159]. For example, a user feeling a tickle on her chin will intuitively map this tactile signal as “something below [the head/center of attention]”.

The second subgoal **G2** of this thesis on the way to establish the head as a means for tactile communication is to design and construct a tactile interface around the head, based on what was learned from the previous chapters. This chapter will discuss the design approach, its limitations, and the HapticHead tactile user interface evolution through multiple iterations. Finally, this chapter presents two additional HapticHead prototypes featuring different actuator types to deal with some of the limitations encountered along the way.

This chapter partly results from CHI’16 [98], and CHI’17 [99] publications and foundations presented in the two previous chapters.

4.1 Concept

The initial idea for the HapticHead concept originated from reading papers on tactile belts used for various use cases (see section 2.6). Specifically, Tsukada et al.’s *ActiveBelt* consisting of 8 vibration actuators and a GPS unit for navigation [206] gave me the idea of developing a more precise tactile guidance method that would also indicate elevation through natural mapping [159].

While indicating elevation through a tactile belt is possible (e.g., through tactile patterns), it is highly imprecise due to the feedback modality of TPs for elevation on a tactile belt being somewhat unintuitive and imprecise to perceive due to missing natural mapping.

Figure 4.1 shows the first HapticHead concept modeled in Unity [208]. Essentially, it consists of three concentric ellipses angled towards each other, each containing six actuators and another two actuators at the intersections of the three ellipses. This design provides a suitable distribution of actuators around the head. It balances the




Figure 4.1.: First HapticHead concept. Notice that all 20 actuators are located on one of three rings around the head and the actuator above the ear is > 4 cm away from the ear opening.

number of actuators required (20 in total) and the ability to provide tactile feedback on as much surface area of the head as possible without putting actuators on insensible areas of the head. Putting actuators on the face's areas, excluding the forehead, is clearly not an option due to comfort and social acceptability.

4.1.1 Limitations and Anatomical Considerations

The HapticHead concept features some inherent limitations:

- Using HapticHead for any scenario includes buying and wearing additional hardware (the prototype itself and a battery).
- There is audible noise due to vibration occurring close to ear openings at frequencies audible by humans.
- Some users have thick hair, which dampens the effect of vibrotactile stimuli.
- Wearing a tactile interface around the head is socially unacceptable in various social situations, and the prototype is hard to disguise fully.

We were aware of these limitations but still wanted to test what we can achieve with a user interface like this, despite some disadvantages. These disadvantages can also partly be mitigated, as shown in the following subsections.

Audible Noise

Standard tactile actuators similar to those used in smartphones (e.g., [164]) are cheap and readily available. However, they typically run at frequencies of around 50 to 250 Hz, depending on applied voltage [164]. Since humans perceive sounds starting at around 20 Hz [140], putting actuators too close to ear openings is not an option as it may be considered too loud and disturbing by users.

A possible electro-tactile display to deal with potential noise issues was ruled out early in the conceptual phase. An electro-tactile display has a few key advantages over a mechano-tactile display:

- **Noise:** An electro-tactile display produces no audible noise, which is a key advantage especially close to the ears,
- **Efficiency:** Minimal current per electrode pair is needed to generate a significant stimulus, typically less than 2 mA [84]. This high efficiency increases battery life for mobile applications.

However, the disadvantages outweigh the advantages over mechano-tactile displays for applications *around* the head:

- **Safety:** While the cranial base provides high electrical resistance on the head, it is nearly impossible to guarantee no electrical currents pass through brain areas. Even if cathode electrodes are very close to anode electrodes, higher skin resistance may (temporarily) occur at a particular cathode, leading to nearly impossible to predict flows of current to other cathodes, which may even be on the other side of the head. Since an electro-tactile display generates stimuli at around 300 V to breach skin resistance [83], this could potentially hurt or even kill users if the currents flow through brain areas.
- **Calibration:** Electro-tactile stimulation is highly individual in terms of pain threshold, skin type, and the connection of each electrode which in turn is influenced by skin type, stratum corneum thickness, hair, and sweat. The difference between a tickling and a painful sensation is often small, requiring extensive per-user calibration.
- **Complexity:** An electro-tactile display around the head would require different kinds of electrodes for hairy and glabrous areas and a complex electrical circuit that produces high voltages while reducing the chance of currents passing through brain areas by design.
- **Cost:** An electro-tactile display around the head would require medical safety certification for applications on humans, leading to a significantly higher cost over mechano-tactile displays. Without a medical safety certification, an ethics committee would unlikely approve the prototype for use in human user studies due to safety concerns.

Due to the disadvantages mentioned above and potential safety issues, a possible electro-tactile display variant to improve some of the issues encountered with the mechano-tactile display presented in this work (e.g., audible noise) is not further discussed as a possible solution.

Another possible solution to audible noise is to generate mechanical stimuli closer to the lower human hearing frequency threshold. Below 200 Hz, human hearing becomes gradually less sensitive to stimuli; below 20 Hz, most humans only hear sounds if they are played back at very high amplitudes [140]. This option is explored with a specialized low-frequency actuator prototype in subsection 4.3.2.

Thick Hair

The thickness of the user's hair is another possible limitation of the concept, as thick hair may weaken the stimulus received [146]. In our experiments on 3D guidance (see chapter 8), we had two participants with thick hair who indicated that they did not receive sufficiently strong feedback on the top of their heads. These participants needed more time to find the correct targets but had a similar success rate as the others. We attribute this to the frontal vibrotactile actuators on the user's forehead, unimpeded by hair.

However, for specific use cases requiring tactile feedback all around the head, thick hair may exclude some users from experiencing the scenario as intended by the designer. As a possible solution to this, we built a prototype with powerful, large actuators (see subsection 4.3.1). Yet, we did not use this prototype in user studies due to the actuator's clunkiness and substantially increased audible noise.

Social Acceptability

Despite wearing a tactile interface around the head at home for VR or AR use cases seems acceptable, it is highly questionable in use cases involving public situations (e.g., outdoors guidance for visually impaired people). Ideally, HapticHead should be hidden under a beanie or helmet to disguise it and allow users to use it in public situations (see Figure 4.6). However, it is impossible to hide the concept's chin strap, which still renders the system unusable in specific social scenarios, even if the rest of the prototype were entirely hidden under a beanie (e.g., dance balls).

4.1.2 Tactile Illusion Considerations

While our research on the localization performance and funneling illusion on the head was not completed by the time of developing the first HapticHead prototype, we were

already aware of related work by Kerdegari et al. [105] who measured localization performance and the funneling illusion on the forehead. Due to this related work and our observations from using the initial prototype in several studies (see chapter 7 and 8), we doubled the actuator density on the forehead and the chin from the first to the second main HapticHead prototype.

Our work on the funneling illusion later confirmed that with the increased actuator density, a majority of HapticHead's users experience the funneling illusion (at least on the forehead and chin, see chapter 3).

4.2 Hardware Implementation

Throughout this thesis, several iterations of the HapticHead hardware prototype were developed. The first version employed in the first guidance study [98] and its follow-up study [99] had several significant flaws which were improved upon in future versions.

Apart from the main prototype, we also built several prototypes which improved certain aspects of the experience, e.g., stronger vibration or less vibration noise through lower frequency actuators. However, the latter prototypes were not used in studies and remain to be validated in future work.

The following subsections will describe the development of the main prototype in-depth and briefly describe other notable and functioning prototypes.

4.2.1 Initial Main Prototype

Myles and Kalb [144] recommend actuators on the head to operate at frequencies between 32 and 150 Hz because of discomfort above that threshold. For the first version of the prototype, we decided to use actuators operating at 150 Hz at maximum because actuator size increases for equally strong impulses at lower frequencies. The first prototype (Figure 4.2) consists of a bathing cap with 17 vibration motors (Parallax, 12 mm coin type, 3.3 V, 90 mA, 9000 rpm) attached on the inside. The non-stretchable chinstrap hosts another three vibration motors and can be adjusted to different head sizes using a Velcro fastener. The vibration motors are controlled by pulse-width modulation (PWM) signals of four Arduino Nanos on switchboards connected to a stationary PC through USB and updated at 75 Hz.

On the software side, vibration motors are modeled at their corresponding position in a Unity [208] scene (see Figure 4.1). This modeling allows easy spatial activation of selected motors, depending on the task. The user's head is tracked either by the



Figure 4.2.: First HapticHead prototype

internal sensors of an HMD to be used in conjunction with HapticHead or by an external tracking system such as OptiTrack.

We used this first prototype in our first guidance study [98], and its follow-up study in chapter 8 as well as in a study on increasing the feeling of presence in virtual reality (see chapter 7). We found the following deficits, which we aimed to fix in future versions:

- The four Arduinos used to power the tactile actuators were not mobile/wearable due to the switchboards and loose cables.
- The three actuators on the forehead and the chin were spaced too far apart (8 - 10 cm depending on head size) so that the funneling illusion did not occur on the forehead and switches between the actuators could be perceived as abrupt and uncomfortable while decreasing guidance performance at the same time [103, 105].
- The coin vibration actuators we initially used were perceived as too weak on some head regions or because of thick hair [96].
- Attaching actuators on the inside of the bathing cap caused visible marks on the user's foreheads after wearing HapticHead for an extended amount of time.

4.2.2 Second Iteration of the Main Prototype

Based on the experiences from the first experiments with the initial prototype mentioned above, we built a second prototype in order to improve precision for guidance applications and user comfort. As discussed in section 8.3, the first prototype's non-stretchable chin belt seemed to have a substantial negative impact on guidance towards




Figure 4.3.: Revised HapticHead concept. Notice that all 24 actuators are located on one of three rings around the head, and the actuator above the ear is > 4 cm away from the ear opening. The forehead and chin regions contain a higher density of actuators intended to increase precision for guidance applications and enable more users to experience the funneling illusion, increasing comfort due to smooth transitions between actuators.

targets on the diagonals below the user. For the second prototype, we replaced the chin belt with a stretchable one. As shown in Figure 4.3, we increased the number of actuators on the forehead and the chin belt from three to five in order to form a high density "ring" of actuators around the face to potentially increase precision for guidance applications and enable more users to experience the funneling illusion which increases comfort due to smooth transitions between actuators (see chapter 3). We made sure to avoid the ear openings again to minimize noise through bone conduction.

For the second prototype, we placed the actuators on the outside of the bathing cap (see Figure 4.4), this time due to feedback from experiment participants who commented on the vibration motors leaving tiny marks on the forehead (see section 8.3). However, for the motors on the chinstrap, we could not place them on the outside because the vibrotactile impulse would have been attenuated too much and spread over a too large area otherwise.

To make the prototype untethered, we exchanged the four Arduinos with a single Raspberry Pi 2 with a Wi-Fi dongle and a standard 5 V USB battery pack on a custom actuator driver board (see Figure 4.5). The custom actuator driver board contains six four-channel Texas Instruments L293NE motor drivers.

For the vibration actuators, we used a total of 24 Precision Microdrives 310-117 - Pico Vibe [163] (10 mm diameter x 3 mm height, 233 Hz frequency at 3.3 V, low starting voltage of 0.9 V). These actuators were controlled with 500 Hz software PWM by the Raspberry Pi 2, and the Raspberry itself received commands at up to 100 Hz, depending on application.




Figure 4.4.: Second, refined HapticHead prototype, side and front view. Notice actuators located on the outside, the flexible chinstrap, and five instead of three actuators on each, the forehead and chinstrap. Positions of the 10 forehead and chin actuators forming a “ring” around the face are marked in red.

On the software side, a computer or phone take care of playing TPs by sending actuator commands wirelessly to the Raspberry Pi 2 through Bluetooth or Wi-Fi. The Raspberry also accepts basic pattern commands to emulate a rhythm locally (e.g., actuator X, 45 ms on, 60 ms off, repeat) regardless of network latency. Fig. 4.5 depicts the system components and the signal flow.

This prototype was only used in chapter 8.




Figure 4.5.: System overview of the HapticHead tactile around-the-head feedback system.



Figure 4.6.: Third HapticHead prototype, hidden under a beanie for improved aesthetics.

4.2.3 Third Iteration of the Main Prototype

The third iteration of the main prototype is only a slightly altered second version where we exchanged the Precision Microdrives 310-117 10 mm coin actuators with the slightly larger 12 mm version Precision Microdrives 312-101. This larger actuator has the advantages of slightly lowering the maximum frequency from 233 Hz to 208 Hz (closer to the optimal perception frequency of Pacinian type corpuscles) and slightly lowering the audible noise while increasing the normalized amplitude from 1.9 G to 2.6 G. We also switched from a Raspberry Pi version 2 to version 3 for performance reasons.




Figure 4.7.: HapticHead motor driver board v2. Developed in collaboration with Tim Duent.

Thus, the third prototype consists of a bathing cap with 19 vibration actuators (PM-312-101 [164], 12 mm coin type, 3 V, 75 mA, 12500 rpm, 2.6 g normalized amplitude, 40 ms lag time, and 132 ms rise time) attached on the outside and distributed on the whole surface just like the second prototype. The stretchable chinstrap hosts an additional five vibration actuators on the inside and can be adjusted to different head

sizes using a Velcro fastener. The vibration actuators are now controlled by software PWM signals at a frequency of 20 kHz using the pigpio library [5] on a Raspberry Pi 3 [168], connected to a slightly improved custom actuator driver board developed in collaboration with Tim Duentel (see Figure 4.7).

The prototype may optionally be integrated into a beanie due to the naked prototype's questionable aesthetics (see Figure 4.6). It was used in chapters 5, 6, 8, and 9.

4.2.4 Fourth and Final Iteration of the Main Prototype

The previously used 2.4 GHz Wi-Fi and Bluetooth operating in the same frequency range were sometimes subject to massive distortions for 500 ms every 60 seconds in our lab environment. Those distortions appeared during the first visually impaired navigation study in section 9.5 and the resulting delays were deemed unacceptable for the follow-up study in section 9.7.

In the fourth and final version of the main prototype, we exchanged the Raspberry Pi 3 with a Raspberry Pi 4 to use 5 GHz Wi-Fi which proved a lot more stable to use in our lab. This final prototype version was solely used in the Improved System Validation Experiment in section 9.7.

4.3 Other Prototype Versions

Several other prototype versions were developed in the timeframe of this thesis. Notable prototypes and the reasons why there were not used in user studies are presented below.

4.3.1 Powerful, Large Actuator Prototype

In our initial studies, some participants noted that they could perceive some of the actuators less than others, specifically on the head's back due to thick hair. As a potential solution to this, we developed a prototype with stronger actuators (see Figure 4.8). The PM-307-103 [162] has a typical normalized amplitude of 7G, which is 2.7 times more than the PM-312-101 [164] we used in the fourth and final iteration of the main prototype.




Figure 4.8.: Specialized, discontinued prototype with stronger actuators. Actuator type: PM-307-103 [162]

During our initial tests, we noticed that the higher amplitude made it easier to localize actuators on the head's back. However, we decided to discontinue this prototype and not use it in a user study due to several major issues:

- the prototype was a lot more clunky due to the dimensions of the actuators and their 3D printed casings,
- the actuators were significantly louder,
- it was not possible to drive the actuators at very low amplitudes for head regions that do not require strong intensities (e.g., forehead or chin) because of rather high minimal amplitudes of around 1G at the maximum start voltage of 1.2 V.

4.3.2 Low-Frequency Actuator Prototype

The prototype shown in Figure 4.9 is a specialized low-frequency prototype that uses Lofelt L5 actuators, a type of wide-band tactile voice coil actuator [125]. Jonas Bock designed this prototype and the custom 3D printed casings for the actuators during his Master's thesis. The goal of using specialized low-frequency actuators (resonance frequency 65 Hz) for the HapticHead concept was to decrease audible noise by lowering actuators' frequency closer to the lower threshold of human hearing (20 Hz).

Unfortunately, the final prototype had significant issues with noise generated by several glitches in the Multiwave FPGA software [186]. Voice coil actuators such as the Lofelt actuator [125] generate audible noise if the input voltage does not correspond to a clean sinus signal over time. Due to these issues with audible noise that proved hard to fix and the general clunkiness, size, and power requirement of the prototype (60 W




Figure 4.9.: Specialized, discontinued prototype with low-frequency Lofelt actuators [125]. This prototype was developed during a Master’s thesis by Jonas Bock. Left: HapticHead with 24 Lofelt L5 actuators in custom casings; top-right: custom driver board with 12 dual MAX98306 stereo amplifiers and a custom FPGA from the MultiWave project [186] to generate sinus wave inputs for the driver board; bottom-right: Single Lofelt actuator and custom 3D printed casing designed by Jonas Bock. Image shown with friendly permission by Lofelt GmbH.

with all actuators turned on), this prototype was not used in a user study and remains to be validated in future work.

Tactile Patterns around the Head

The upcoming chapter deals with the third subgoal **G3** of this thesis, on the way to establish the head as a means for tactile communication, which is to research how many and what kinds of tactile patterns on the head users can distinguish and which kinds of patterns are accepted among the user base.

Urgent notifications via tactile patterns present the most basic use case of HapticHead, yet these kinds of notifications are helpful in various use cases. A smartphone vibrating in a pocket can easily be missed while a strong vibrotactile pattern on the head is perceived much more robust and thus appears more present to the user.

We propose around-the-head spatial vibrotactile patterns for representing different kinds of notifications. The patterns are defined in terms of stimulus location, intensity profile, rhythm, and roughness modulation. A first study evaluates recall and distinguishability of 30 patterns, as well as agreement on meaning without a predetermined context: Agreement is low, yet the recognition rate is surprisingly high. We identify which kinds of patterns users recognize well and which ones they prefer. Static stimulus location patterns have a higher recognition rate than dynamic patterns, which move across the head as they play. Participants preferred dynamic patterns for comfort. A second study shows that participants can distinguish substantially more around-the-head spatial patterns than smartphone-based patterns. The spatial location has the highest positive impact on accuracy among the examined features, so this parameter allows for a large number of levels.

This chapter's content is primarily based on the MUM 2020 publication "Design and Evaluation of On-the-Head Spatial Tactile Patterns", written in collaboration with Michael Rohs [101]. Parts of this chapter are based on a Bachelor thesis by Marc Mogalle [139].

We used the third main prototype (see subsection 4.2.3) for this chapter's experiments.

5.1 Introduction and Motivation

This chapter introduces a notification scenario and (Vibro-)tactile patterns (TPs) presented at different spatial locations around the head using HapticHead. One of the benefits of tactile notifications (TNs) around the head is that they are much more

“present” and cannot be missed as easily as a smartphone vibrating in a pocket or even a vibrating smartwatch. Thus, they also allow use cases in which receiving an imminent tactile warning is essential, such as when stopping at a red light or taking a turn on a narrow sidewalk in a navigation scenario. Furthermore, TNs around the head feature a larger design space than simple tactile feedback using just one actuator at a specific position, particularly regarding spatial dynamics. The design space of around-the-head TPs includes:

- multiple spatially distributed actuator *locations* on the surface of the head, which might be chained together over time,
- *intensity* curves of each actuator signal over time,
- *roughness* of each actuator signal over time,
- *rhythm* of each actuator signal over time.

This larger design space, especially in terms of using multiple spatially distributed actuators, allows for more sophisticated TPs compared to other tactile feedback solutions and, thus presumably, makes it easier to distinguish between patterns in a set of patterns for different meanings. The tactile notifications presented in this work include smartphone notifications where users can not only identify the kind of notification and its priority (as in work by Brown et al. [24] who used a single actuator for feedback), but also the source of the notification. The possible user groups of such a tactile around-the-head notification system are plentiful. They include people (especially with visual impairments or on bikes) in navigation scenarios, firefighters in low visibility situations, jet pilots, and even smartphone users who do not want to miss notifications while also being able to distinguish various patterns for different meanings.

This chapter explores various kinds of spatial TPs around the head for general notifications. We ran two exploratory user studies on the following research questions:

- RQ1** What kinds and what number of around-the-head spatial TPs can be easily distinguished?
- RQ2** How do around-the-head spatial tactile notification patterns compare to tactile smartphone notifications (especially regarding the amount of information they can convey)?
- RQ3** Can the lower tactile sensitivity of the head compared to the fingertip [82] in terms of TP recognition accuracy be offset by using spatially distributed rhythms (e.g., rhythms that use more than one actuator over time)?

The main contributions of this chapter include:

- The first investigation into a large variety of different possible TPs emphasizing stimulus location all around the head.
- TP recognition accuracy and user acceptance of 30 different TPs around the head.

- Confirmation that the lower tactile sensitivity of the head compared to the fingertip [82] in terms of TP recognition accuracy can be offset by using spatially distributed rhythms.

For this chapter, we use the vibrotactile pattern terminology as defined in section 2.5.

5.2 Related Work

The fundamental related work on tactile perception around the head (subsection 2.3.8) and tactile displays (section 2.6) also applies to this chapter. Specific related work concerning tactile patterns is presented below.

While some of the previously cited works in the previous chapters mention the opportunity that lies in creating rich TPs through using multiple spatial locations around the head, there is very little research in actually creating patterns in this vast design space and evaluating their detection accuracy. This lack of research might be caused by most prototypes not covering the entire head except for Diener et al. [41]. However, even they did not create TPs but only measured detection accuracy between single actuators [41].

With this chapter, we aim to provide a first investigation into the vast design space of creating TPs all around the head and specifically use the stimulus location of TPs in order to create more intuitive, easier to recognize TPs because several of the previously mentioned related works identified stimulus location as more straightforward to identify than frequency or intensity.

Prior Work on Notifications and Vibrotactile Patterns. Sahami Shirazi et al. [195] conducted a large-scale assessment of mobile notifications and concluded that users rated the importance of messenger, voice, mail, social, and calendar notifications on an Android device much higher than nine other possibilities. Therefore it makes sense to concentrate on these categories of applications when developing (tactile) notifications.

Brown et al. [24] investigate the effectiveness of specialized TPs (Tactons) using an actuator type called Tactor (a voice coil) on the fingertip. They found that their pattern “rhythm” was easily identified with an average success rate of 93 % for three different possibilities while the “roughness” of the pattern (another three possibilities) was less easy to identify with an average success rate of 80 %. Our experiments in this chapter were inspired by this work and introduce another factor besides rhythm and roughness: spatial location around the head. The addition of spatial location supposedly increases the total number of distinguishable patterns as pattern location is another dimension in pattern design that is easily recognizable [82].

We carried out two experiments that look at how many TPs users can distinguish on the head (*Distinguishability of Patterns Experiment*) and which patterns make sense for notifications. The *Tactile Pattern Notification Experiment* features a direct comparison to an existing TP study from related work [24], comparing and highlighting the effect of spatial vibrotactile stimuli versus other influenceable factors in TP design.

5.3 Distinguishability of Patterns Experiment

In the first experiment, we aimed to answer two questions: Is there agreement on an intuitive meaning of specified patterns without a context, and how well can participants recognize and recall different patterns? Furthermore, which kinds of patterns perform best in terms of recognition? To answer these questions, a two-part experiment was designed: In part 1, we asked the participants to rate all of the patterns along different scales and come up with a description for each one. In part 2, we repeated the patterns and asked them to match the repetition to one of their descriptions.

We invited 20 participants (14 male, six female, mean age 28.5 years, SD = 7.9 years), of which four had previously participated in other HapticHead studies. Prior participation in other HapticHead studies should be no confound since the current experiment varies significantly from the previous experiments, which were not in the domain of distinguishing TPs.

5.3.1 Vibrotactile Pattern Design

For this study, a total of 30 patterns were designed in an iterative process (pre-experiment) by two researchers, based on the distribution of the actuators in the HapticHead prototype, geometrical patterns, and (head) areas of interest (see Figures 5.2 and 5.1). We initially designed a total of 46 patterns and then excluded those that felt too similar to other patterns or actuated too many motors at the same time (leading to sensory overload), leaving 30 viable patterns. While some of these patterns use only a static stimulus location, others integrate a dynamic stimulus location as described in section 2.5. The designed patterns do not utilize the entire design space of the HapticHead concept (see section 5.1), but they do provide exemplary patterns within this design space, primarily utilizing the actuator spatial *location* aspect, as this tended to work best for recognition in the pre-experiment and related work [63, 82].

Pattern name	Pattern location and description	static / dynamic stimulus location	total # of actuators	active time each actuator [s]	total time [s]
HeadBottom	Simultaneous vibration of all five actuators on the chin-strap	static	5	1.4	1.4
HeadLeft, HeadRight, HeadFront, HeadTop, HeadBack	Simultaneous vibration of six actuators on the left/right/front/top/back side of the head	static	6	1.4	1.4
BackHeadClockwise, BackHeadCounterClockwise	Vibration is moving in a circular shape on the back and back-top of the head in a clockwise/counter-clockwise motion, starting at the central-back actuator.	dynamic	6	0.35	1.6
BackHeadOutsideIn, BackHeadInsideOut	Utilizing a trapezoidal shape of motors on the back and back-top of the head, the vibration starts at both ends and moves to the center above the ears / reversed	dynamic	8	0.35	0.85
CircleHorizontalLeft, CircleHorizontalRight	Vibration is starting on the forehead and moves over the left/right side in a circle back to the start.	dynamic	10	0.35	2.35
CircleHorizontalSplit	Vibration is starting on the forehead and moves synchronously on both sides to the back of the head.	dynamic	10	0.35	1.6
CircleLeftForwardSlow, CircleLeftBackwardSlow, CircleRightForwardSlow, CircleRightBackwardSlow	Vibration of six actuators on the left/right side of the head in a circle, starting at the left/right-most forehead actuator, moving down/up	dynamic	6	0.35	1.6
CircleSyncForwardSlow, CircleSyncBackwardSlow	Vibration of six actuators in a circular arrangement on the each side of the head, starting at the outmost forehead actuator, moving down/up.	dynamic	12	0.35	1.6
NavLeft, NavRight	Vibration moving from the forehead over the temple to the left/right side of the head	dynamic	3	0.35	0.85
NavUp, NavDown	Vibration moving upwards from the forehead to the top of the head / reversed	dynamic	3	0.35	0.85
SpiralLeftBackward, SpiralRightBackward	Vibration is starting on the top of the head and moving downwards to the left/right side until it reaches a 540 degree (one and a half) turn, ending on the neck. / reversed. Actuator time and intensity depends on the distance to a spiral layed between the actuators. Multiple actuators can be active at the same time.	dynamic	15	0.05 - 0.5	2.75
SpiralLeftForward, SpiralRightForward	actuator time and intensity depends on the distance to a spiral layed between the actuators. Multiple actuators can be active at the same time.	dynamic	15	0.05 - 0.5	2.75
TopHeadClockwise, TopHeadCounterClockwise	Vibration is moving in a circular shape on top and top-front of the head in a clockwise/counter-clockwise motion, starting at the top-front actuator.	dynamic	6	0.35	1.6
TopHeadXSimultaneous	The vibration simulates the shape of an "X" on the head, both lines starting at the corresponding temple. The lines end at the opposing side of the neck for each line.	dynamic	9	0.35	1.35

Figure 5.1.: Descriptions of all 30 patterns. Left/right/reversed refer to different variants. These different variants may have different directions (e.g., left/right) or normal/reversed starting/ending points and play direction (e.g., forward/backward). Note that all static patterns with a dynamic stimulus location have a **100 ms overlap** between active actuators to make the patterns feel smoother. The actuators we use need an active time of greater than 40 ms due to their lag time as they would otherwise not produce any stimulus.




Figure 5.2.: Visualizations of representative static pattern with a dynamic stimulus location. From left to right: BackHeadInsideOut, CircleLeftForwardSlow, NavLeft, SpiralLeftForward, TopHeadXSimultaneous. Starting point(s) are marked in blue, ending point(s) are marked in darker green. The order of actuators is marked in bright green numbers and through the arrows on the paths. Note that there are simultaneous activations for some patterns.

5.3.2 Distinguishability of Patterns Experiment – Procedure

After a short greeting, the participants filled out an introductory questionnaire, a mandatory informed consent form, and optionally a photographic release form. Afterward, they were introduced to the HapticHead prototype, and it was adjusted to fit them properly. The experimenter controlled the Android application. Once the participant was ready, the experimenter started with the **training part** as described below and, after a short pause, continued with the **test part**. A post-questionnaire gathered further subjective data.

The **training part** of this experiment served to familiarize the participants with the different patterns. Patterns were played in random order, and participants were encouraged to write down a brief description (intuitive association or pattern shape) on a sheet of paper. They could ask to repeat a pattern as often as they liked to while finishing their descriptions. We chose this approach to get the participants to think about the patterns actively and create a vocabulary for talking about them later. Furthermore, patterns were rated by four statements on 5-point Likert scales: This pattern was [intuitive, concise, useful, pleasant].

In the **test part** of this experiment, the patterns were again randomly shuffled, and each pattern was presented three times (90 trials per participant) again. The participants were to match repeated patterns to their descriptions from the training round. However, this time no repetitions were allowed, and if the participant did not recognize a pattern, it was marked as “unknown” and skipped. We measured the time from the start of the pattern until the participant answered.

This study design was implemented into an Android application that controls the third main HapticHead prototype via Bluetooth and is installed on a Galaxy S6 smartphone (see subsection 4.2.3).

5.3.3 Distinguishability of Patterns Experiment – Objective Results




Figure 5.3.: Overall confusion matrix of all 30 patterns, N=1800 trials (30 patterns x 3 repetitions x 20 participants).

Fig. 5.3 shows a confusion matrix for all 30 patterns, including a column for when a pattern was not recognized (“unknown”). Analyzing the data reveals that the mean recognition rate of all patterns was at 56 %, with the worst recognized pattern at about 27 % and the best with a recognition rate of 98 %. The statistical data of the 15 best patterns in terms of recognition rate and their corresponding Likert ratings are listed in Table 5.1. Fig. 5.3 shows a confusion matrix including all 30 patterns. The average time to answer in the second part of the experiment was 12.7 s (SD=10.2 s, median=9.2 s, min=2.3 s, max=83.6 s).

5.3.4 Distinguishability of Patterns Experiment – Subjective Results

Fig. 5.4 shows the questions and results of the post-study questionnaire, which indicate that the participants found the patterns suitable to convey information. The participants also rated the vibration feedback as rather pleasant and judged the patterns’ length as adequate. Moreover, 75 % of the participants stated that they would use a version of HapticHead integrated into a cap or beanie in real-life situations.

Pattern recognition & Likert ratings

<i>Pattern</i>	<i>Rec. rate [%]</i>	<i>Intuitive</i>	<i>Concise</i>	<i>Useful</i>	<i>Pleasant</i>	<i>Rec. time [s]</i>
HeadBottom	98	4.00	3.94	2.83	1.28	5.68
HeadRight	89	3.68	3.84	2.84	1.58	8.87
HeadFront	88	3.83	3.56	2.67	1.22	8.87
HeadLeft	86	3.83	3.83	2.83	1.39	11.27
HeadBack	79	3.94	3.89	3.00	1.72	12.96
NavRight	79	3.60	3.55	3.15	3.25	12.67
NavLeft	75	3.50	3.33	2.94	3.44	13.13
HeadTop	70	3.44	3.56	3.06	1.44	9.49
NavUp	68	3.53	3.37	3.11	3.63	13.18
NavDown	67	3.68	3.53	3.32	3.58	12.43
CircleRightBackwardsSlow	61	3.00	2.94	2.78	3.06	11.20
CircleLeftForwardsSlow	61	2.45	2.45	2.10	3.10	13.84
CircleHorizontalSplit	58	3.61	3.44	2.72	3.28	13.11
CircleLeftBackwardsSlow	58	3.11	3.06	2.56	3.33	11.91
CirclesSyncForwardsSlow	54	2.89	2.89	1.79	2.58	14.02

Table 5.1.: Means of recognition rates of the top 15 TPs in the Distinguishability of Patterns Experiment and correlating Likert ratings (0: strongly disagree – 4: strongly agree), sorted by rec. rate.

5.3.5 Distinguishability of Patterns Experiment – Discussion

The top 15 patterns in terms of recognition rate, shown in Table 5.1 except for CircleHorizontalSplit and CirclesSyncForwardsSlow, share some common characteristics. For one, they reside on clearly different parts of the head (in the sense of a clearly defined area such as “top of the head” or “right side”) and have a rather short playtime (≤ 1.6 s). The confusion matrix in Fig. 5.3 indicates that the most detailed patterns (spiral patterns SLB, SLF, SRB, and SRF), which use nearly every actuator and have the longest playtime (2.75 s), are often confused for one another. The better-performing patterns are short and reside on clearly distinct areas of the head.

The influence of the pattern position is also visible in the HeadTop pattern in Fig. 5.3. While it still is a better than average pattern, it was mistaken in 25% of the cases for the HeadFront pattern. We presume that the reason for this is their close proximity and the overlap with other pattern areas. Other patterns nearby are HeadLeft (HL), HeadRight (HR), and HeadBack (HBa). Furthermore, related work found that the top of the head is less sensitive to vibrotactile stimuli and that localization precision is lower on the top of the head [41, 144, 145], which becomes apparent in the performance of the patterns: HeadTop had the lowest recognition rate of all Head* patterns. Similar observations can be made for the NavUp and NavDown patterns (residing on the top of the head), which have lower recognition rates than NavLeft and NavRight (residing on the head’s sides).

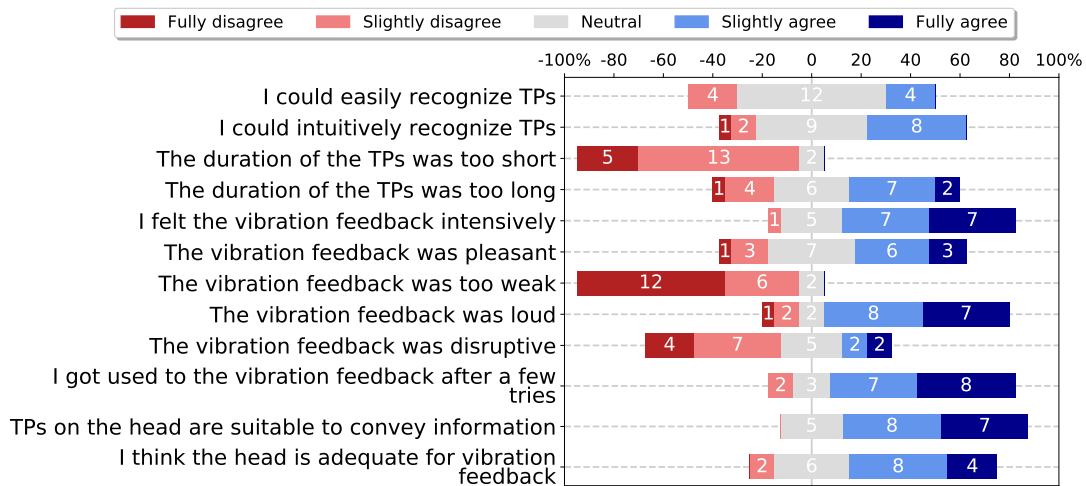


Figure 5.4.: Post-questionnaire results of the Distinguishability of Patterns Experiment about properties of the 30 implemented TPs, $N = 20$.

Table 5.1 indicates that the best patterns in terms of recognition rate were *static* stimulus location patterns (patterns named Head*). However, these patterns were rated much lower in terms of pleasantness than the static pattern with a dynamic stimulus location (all others, see Fig. 5.1). This is likely due to the number of simultaneously active actuators such that the added up vibration amplitude of static patterns is much higher than for dynamic ones, making them less pleasant to feel but easier to identify due to the stronger signal.

For most of the patterns, the agreement on the meaning of a pattern was relatively low. This may be because we did not strongly instruct participants to write down their intuition for each pattern but just asked them to provide a description that captured their personal understanding of the pattern. Some descriptions reflected the participants’ intuitive understanding (e.g., participant P5 on pattern SpiralLeftBackward: “face massage” or P11 on SpiralLeftForward: “spider crawls upwards”). Other descriptions just defined the location (e.g., participant P6 on pattern TopHeadClockwise: “circle on top of the head” or P17 on HeadRight: “right side, very strong”). Most participants created location-based descriptions, so we are not able to conclude common intuitions for most patterns. Thus, we cannot answer the research question RQ1 on whether there is agreement on specified patterns’ intuitive meaning. However, we can answer how well participants could remember and recognize the 30 patterns (see Fig. 5.3).

These results lead us to the second experiment, in which we take some of the patterns that performed well and use them for actual smartphone notifications to measure recall and distinguishability of these patterns in an actual use case and compare our system to related work.

5.4 Tactile Pattern Notification Experiment

The second experiment aims to characterize the performance of vibrotactile notifications via HapticHead regarding distinguishability and ability to deliver information. As a baseline and as a guide to structure the experiment, we used the work by Brown et al. [24] who used a single Tacton C2 actuator mounted on the fingertip. They created three *rhythms* and three *roughness* values to convey the *type* of notification (a call, text or multimedia message) along with the notification *priority* (low, medium, high).

It is debatable whether comparing our interface to Brown et al. is a *fair comparison* as we compare an interface with 24 spatially distributed ERM actuators around the head versus an interface with only one C2 Tacton actuator on the fingertip. However, with this experiment, we essentially want to show that the HapticHead concept can transmit a much higher amount of information using the much higher count of vibrotactile actuators (this is expected due to related work [63, 82]). Even more, the rhythm and roughness conditions may feature only slightly less accuracy on the HapticHead, compared to the much more tactile-sensitive fingertip (due to higher nerve-densities) [82]. To achieve only slightly less accuracy in the rhythm condition, we use *spatially distributed rhythms* (see section 2.5), so that rhythms are not only dependent on modulating the input signal of a single actuator but several actuators one after another (which may still all reside on one identifiable head region).



Figure 5.5.: Internal and front view of the modified Samsung Galaxy S6 dummy.

We recreated Brown et al.'s experiment [24] with a smartphone dummy and mapped Brown et al.'s patterns on a vibration actuator of the same type as used for HapticHead (Fig. 5.5). We chose a smartphone because this represents a widespread use case and allows us to compare a standard ERM actuator which we are also using on the HapticHead prototype, to the Voice Coil Actuator (Tacton C2) that [24] is using. This setup was compared to spatial patterns created by our HapticHead prototype. In the HapticHead condition, we added another dimension to the messages via three different *spatial locations* (left, right, and top of the head) to convey the notification *source* (the person who sent the notification).

We invited 20 participants (16 male, four female, mean age 26.8 years, SD 7.9 years), of whom nine had previously participated in the Distinguishability of Patterns Experiment. This means that 45 % of the participants were familiar with HapticHead and had

previously felt some of the used patterns before, but since there were two months between the experiments and the patterns have different meanings, the confound should be minor.

5.4.1 Tactile Pattern Notification Experiment – Design & Implementation

Brown et al. [24] conveyed notification priority with complex waveforms: They implemented amplitude modulation with a C2 Tactor (sine waves, 250 Hz base frequency, and 0, 30, or 50 Hz modulation frequency) to encode “roughness.” For performance characteristics of the C2 Tactor, we refer to Azadi et al. [7] who compared the C2 Tactor to a common Linear Resonant Actuator (LRA). The LRA in their comparison features similar dimensions and performance characteristics to the ERM vibration actuator that we use.

Our much smaller and less expensive ERM coin vibration actuators are not able to correctly playback the high-frequency complex waveforms used by [24] because of the latency induced by motor spin-up and stop-time. Thus we resorted to a square wave modeled through pulse width modulation as explained in section 2.5. The square wave conveyed “roughness”, which was used to encode notification priority in [24]. This lower-fidelity mapping supposedly decreases accuracy, so our experiment is not a perfect one-to-one comparison because of hardware differences. However, in case the smartphone condition yields similar results as found in [24], then this means that the effect of complex waveforms is not essential and can be substituted by square waveforms on standard vibration actuators. It would also mean that using similar settings for roughness in the HapticHead condition would make the results comparable to those reported in [24]. We included the smartphone condition here to enable this potential comparability.

In the experiment, we used the following square wave settings for modeling notification priorities as roughness. The settings were determined in a pre-experiment.

- High: fully on
- Medium: on/off ratio of 23 ms to 35 ms (17.2 Hz)
- Low: on/off ratio of 25 ms to 50 ms (13.3 Hz)

The off-time generally needs to be higher than the on-time to get the maximum possible roughness feeling because the stop-time of standard vibrotactile coin actuators, such as the ones we used, is generally around twice as high as the spin-up time. Note that modeling roughness in this way also varies the overall perceived intensity of the actuators, as they are not fully on all the time.

The actual experiment was split into two parts: The first part uses the smartphone dummy to test the recognition of message types (rhythm) and priorities (roughness) as in [24], while the second part uses the HapticHead prototype and our patterns to encode message type (spatially distributed rhythm), priority (roughness), and source (location).

The patterns used in this experiment’s HapticHead condition are six of the top ten patterns of the Distinguishability of Patterns Experiment and three new patterns (BackLeft, BackRight, and BackUp). The new patterns are designed, taking into account the lessons learned in the Distinguishability of Patterns Experiment. These new patterns are mirrored versions of NavLeft, NavRight, and NavUp (see Fig. 5.1) and start on the back rather than the front of the head but are otherwise the same. Each side of the head is mapped to a specific sender, with “Emma” on the left, “Ben” using the middle part of the head, and “Alex” on the right side of the head (see Table 5.2).

Assignment of patterns

<i>Type</i>	<i>Source “Emma”</i>	<i>Source “Alex”</i>	<i>Source “Ben”</i>
Call	HeadLeft	HeadRight	HeadTop
Text	NavLeft	NavRight	NavUp
Multimedia	BackLeft	BackRight	BackUp

Table 5.2.: Tactile Pattern Notification Experiment: Assignment of patterns to message type (spatially distributed rhythm) and source (location) for the HapticHead condition.

To keep our study comparable to [24] we also used a training period for each participant of up to 20 minutes duration. The average duration across all participants was 17 minutes. The training period allowed participants to familiarize themselves with both types of output devices (smartphone and HapticHead) and conduct training on all available tactile notifications in a training view within the study app where participants could play back all possible pattern combinations intentionally.

This design was integrated into an Android application for the Tactile Pattern Notification Experiment that controls the third main HapticHead prototype via Bluetooth, installed on a Galaxy S6 smartphone (see subsection 4.2.3).

5.4.2 Tactile Pattern Notification Experiment – Procedure

After a short greeting, the participants filled out an introductory questionnaire, a mandatory informed consent form, and optionally a photographic release form. They were then introduced to the smartphone dummy and the HapticHead prototype. After that, the HapticHead was adjusted to fit. The experimenter explained the Android application to the participant and showed how to use it. Following the introduction, the participants started their 20 minutes training exercise.

Following the training phase, the actual experiment started. In the first part of the experiment, all nine combinations of message type and priority defined in [24] were played three times in random order (seeded by user id). Each of those 27 combinations was played on the smartphone dummy followed by the participant inputting message type and priority through distinct buttons in the experiment app. The participants usually held the smartphone dummy in their non-dominant hand while experimenting and selected items on the screen with their dominant hand.

In the second part of the experiment, all combinations of message type, priority, and source were added three times to a list and randomly shuffled (seeded by user id). Those 81 combinations were played on the HapticHead prototype and were followed by evaluating message type, priority, and source done by the participant in distinct menus.

In both parts of the experiment, the participants could repeat the patterns by pressing a replay button. A post-questionnaire gathered further subjective data on the perceived performance and acceptability of the prototype.

Counterbalancing between the two parts of the experiment was unnecessary as the 20 min training phase can preclude a potential bias through learning before the study, where participants practiced both interfaces, including visual feedback on correctness. After the training phase, participants did not get feedback on correctness anymore. Another potential bias would be forgetting patterns during the experiment, but these would be easy to identify outliers. We did not find evidence of participants forgetting patterns in the data.

5.4.3 Tactile Pattern Notification Experiment – Objective Results




Figure 5.6.: Comparison of the mean success rates (message type and priority separately) of Brown et al. [24] and our experiment. Significant differences between our conditions are marked (Wilcoxon signed-rank test, sign. level $p=0.05$).

As shown in Fig. 5.6, participants correctly identified the call type (encoded by rhythm) in 94.0 % (SD=16.6 %) of the cases in the HapticHead condition and in 85.7 % (SD=17.8 %) of the cases in the smartphone condition. The results reported in [24] are comparable.

The recognition rate of the priority (encoded by roughness) was significantly lower in the HapticHead condition (mean 64.8 %, SD=9.4 %) compared to the smartphone condition (mean 79.2 %, SD=7.6 %).

The source of a message (encoded by location) was recognized with a mean accuracy of 99.3 % (SD=0.9 %) in the HapticHead condition. The patterns that were located on the top of the head (source Ben) scored a bit lower (98.5 %) than the other two sources (mean 99.7 %).

We identified two outlier participants who had a much lower success rate than the others to identify message type/spatially distributed rhythm in the HapticHead condition at 54 % (P8) and 36 % (P16), respectively. Without them, the other participants reached a mean success rate of 99.4 % (SD=1.0 %) for message type, which is on par with the source's success rate. The two outlier participants reached similar success rates in terms of message priority and source, so we did not exclude them from the evaluation.

A t-test did not reveal significant differences in terms of recognition rate for message type/rhythm ($t = 0.66, p > 0.05$) and priority/roughness ($t = -2.1, p > 0.05$) between the participants who had previously taken part in the Distinguishability of Patterns Experiment and those who had not. However, for message source/location there is a significant difference between these groups ($t = -2.74, p = 0.02$). For message source/location, only 12 out of a total 1620 trials failed. The group who previously participated in the Distinguishability of Patterns Experiment had a success rate of 99.7 % (SD=0.5 %), while the group who did not had a success rate of 98.9 % (SD=1.0 %), so the difference is rather small.

Participants repeated the patterns 1.64 times (SD=0.81 times) in the smartphone condition and 1.79 times (SD=0.86 times) in the HapticHead condition.

5.4.4 Tactile Pattern Notification Experiment – Subjective Results

In the post-study questionnaire (Fig. 5.7), participants stated that it was easy and intuitive to recognize TPs on the head and found them easy to remember. The smartphone condition scored lower in all of these statements. A Wilcoxon signed-rank test showed that those differences were significant. The participants had mixed opinions on whether the vibration feedback was too weak in the smartphone condition, even though our smartphone dummy used quite a large vibration actuator compared

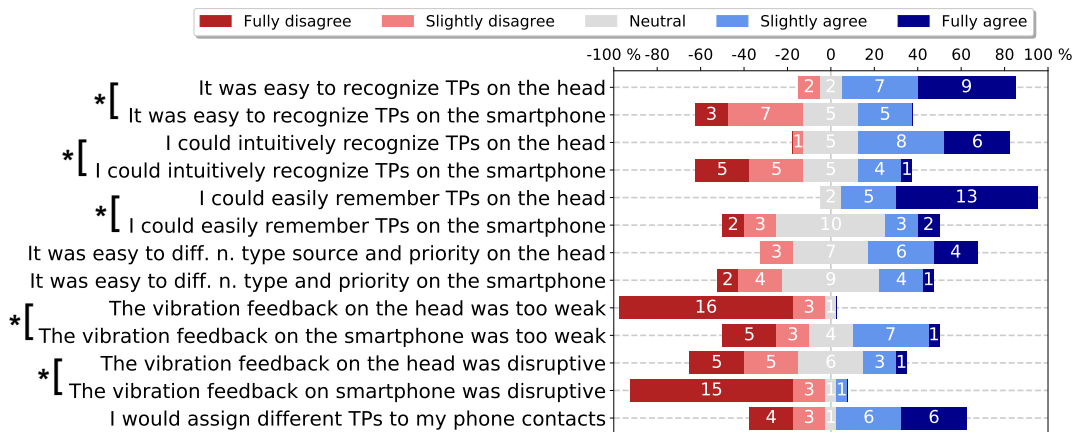


Figure 5.7.: Results of the post-questionnaire of the Tactile Pattern Notification Experiment about recognizing TPs on the head vs. a smartphone (N = 20). Significant differences between statements found using Wilcoxon signed-rank test. Significance level $p=0.005$.

to regular smartphones. Opinions also differed widely on whether vibration feedback on the head was regarded as being disturbing.

5.4.5 Tactile Pattern Notification Experiment – Discussion

When directly comparing our smartphone condition to [24] it is apparent that the participants performed similarly in terms of success rate (see Fig. 5.6). A comparison of significant differences between the works is not possible, as Brown et al. did not publish raw data. However, even though the participants’ performance in [24] is a little higher overall than in our smartphone condition, the differences are marginal. Thus, a comparison of the pattern recognition accuracy between [24] and our HapticHead concept is justifiable, even though the two different works used very different actuator types (ERM actuators vs. Tacton C2).

Message type (call, text, multimedia), encoded by “rhythm” was recognized very well in the HapticHead condition and on par with the performance reported in [24]. This means that the higher vibrotactile sensitivity due to the higher vibrotactile sensitivity of the fingertip [82], and thus expectably higher recognition rate on rhythm, can be offset by introducing spatially distributed rhythms which may still reside on one clearly identifiable area of the head (e.g., left side).

However, priority encoded by “roughness” seems complicated to grasp through the current HapticHead prototype and generally performed worse than in the smartphone condition or in [24]. One of the reasons apart from the lower vibrotactile sensitivity of the head compared to the fingertip [82] may be the difference in the number of actuators active in parallel in each of the patterns, as this affects the perceived

intensity. While the call patterns, which were the *static* stimulus location patterns (HeadLeft, HeadRight, HeadTop from the Distinguishability of Patterns Experiment), used a relatively large area with many actuators activated simultaneously, the message and multimedia patterns only use three actuators activated sequentially (static patterns with a *dynamic* stimulus location). Some participants voiced their concerns that the lowest priority call pattern still felt stronger than the highest priority of a message or multimedia pattern. This may have led to a false assessment of the priorities, especially after the opposing group's pattern (static vs. dynamic stimulus location patterns) was played.

Source (Emma, Ben, Alex), encoded as *location* on the head, was recognized exceptionally well with a mean success rate of 99.3 % across all notification types and priorities. A success rate this high opens up a whole new set of use cases in situations in which users depend on recognizing a pattern correctly, such as when having to stop for a red light at an intersection. A slightly higher success rate for spatial pattern location than for other influenceable factors (intensity, rhythm, roughness) in TP design was expected and confirms previous research by Hameed et al. [63] and other works summarized by Jones and Sarter [82] who found that stimulus location and duration are in general easier to identify than intensity or frequency. Our experiment contributes this finding for the head as a specific body part and finds a higher success rate on the head than related work on other body parts [82], even though this might not be directly comparable due to different experimental setups.

The subjective feedback (shown in Fig. 5.7) suggests that participants rated HapticHead substantially higher than the smartphone in terms of ease, intuitiveness, and recall. This is likely a result of the added location (source) dimension, which was extremely easy to recognize and translated into the overall confidence of identifying patterns correctly. The vibration intensity was also rated significantly stronger for HapticHead than the smartphone condition even though the smartphone used a stronger than usual 12 mm coin vibration actuator rated at 2.6 g normalized amplitude, which is higher than the vibration actuators used in most smartphones currently on the market. Thus, HapticHead also presents itself as a solution to missing important notifications due to not noticing smartphones vibrating in pockets or bags. The subjective results also provide a strong indication that the overlaying hypothesis H1 of this thesis on the perception of tactile patterns feeling strongly *present* to the user (see section 1.1) is accurate as most study participants agreed that they could easily recognize TPs on the head, intuitively recognized them, and strongly disagreed that the vibration feedback on the head was too weak.

RQ2 (see section 5.1) can be answered by evaluating the results of this experiment. Around-the-head spatial tactile notifications can indeed convey a lot more information than tactile smartphone notifications due to the added dimension of multiple spatial actuator locations. While the spatially distributed rhythm was generally easier to identify compared to the smartphone condition (especially when disregarding the

two outlier participants in this experiment), the roughness was more challenging to identify even though the latter might be caused by using very different actuator counts for different notification types.

5.5 Limitations

We did not cover the entirety of possible patterns and accompanying instructions hinted at in the introduction as the design space is too large to cover every possibility. There are most likely even better possible patterns to convey particular instructions (e.g., higher recognition rate, shorter playtime). This is an optimization challenge for future work, which should also investigate letting users design their own patterns with the help of suitable pattern design software (e.g., VRTactileDraw presented in chapter 6). An open question is then whether users can remember their patterns better than those designed by others, which could be the case given related work on user-defined gestures [147].

Thick hair may also be an issue when designing TPs, especially on the back of the head. There is related work that investigates the vibrotactile sensitivity concerning hair density [146]. As long as TP intensities are not too low, hair thickness appears to be a non-issue. While we did have several participants with thicker-than-usual hair in our experiments, we did not find conclusive evidence that they performed worse than others.

5.6 Conclusion and Future Work

Looking back at the research questions in the introduction, we were able to show that participants can identify a large number of different TPs around the head for notification. We identified kinds of TPs that score well in terms of recognition rate and user acceptance. The results also show that the use of locations on the head substantially increases the design space and the number of different tactile notifications that can be successfully identified, especially compared to tactile smartphone notifications. This is due to combining stimulus location with pattern rhythm, intensity, and roughness modulation. Participants were able to distinguish locations more accurately than spatially distributed rhythm or roughness. This was partly expected given related work [63, 82] that found that stimulus location and duration are typically easier to identify than intensity or frequency. However, the ease of identifying stimulus location depends on actuator characteristics and the specific body part, so one experiment on a specific body part does not yield reliable recommendations for the entire body. We contribute our analysis specifically for spatially distributed vibrotactile feedback for locations around the head with standard vibration actuators.

We also contribute our findings from the Tactile Pattern Notification Experiment where we were able to show that when designing TPs, the higher vibrotactile sensitivity due to higher nerve density of the fingertip compared to the head [82] can be offset by introducing localized spatially distributed rhythms instead of using just rhythms on single actuators.

Another interesting finding from this work is that while *static* stimulus location patterns were generally easier to recognize, they were rated lower in terms of pleasantness than static patterns with a *dynamic* stimulus location. This trade-off should be considered when designing TPs for real-world applications.

VRTactileDraw – A Tactile Pattern Design Interface for Complex Arrangements of Actuators

With our experience working with tactile patterns, we learned that designing them is a daunting task when defining the intensities over time for many actuators individually (e.g., for HapticHead). The need for an easy-to-use tactile pattern editor for tactile user interfaces featuring complex arrangements of many actuators could not be adequately satisfied by related work. The next subgoal **G4** of this thesis on the way to establish the head as a means for tactile communication is to build a pattern design software that can deal with a large number of actuators in complex configurations, allowing even novices to design their own tactile patterns for complex tactile user interfaces.

We describe the design and implementation of our tactile design software *VRTactileDraw*, a VR interface that enables designers to prototype TPs for complex tactile interfaces rapidly. It allows for painting strokes on a modeled body part and translates these strokes into continuous tactile patterns using an interpolation algorithm. The presented VR approach avoids several problems of traditional 2D editors. It realizes spatial 3D input using VR controllers with natural mapping and intuitive spatial movements. To evaluate this approach in detail, we conducted two user studies and iteratively improved the system. The study participants gave predominantly positive feedback on the presented VR interface (SUS score 79.7, AttrakDiff “desirable”). The final system is released alongside this chapter as an open-source Unity project for various tactile hardware.

This chapter is based on two Bachelor’s theses by Leonard Hansing [64], and Andreas Domin [44], a CHI’ 2019 late-breaking work [94], and the INTERACT’ 2021 paper “VRTactileDraw: A Virtual Reality Tactile Pattern Designer for Complex Spatial Arrangements of Actuators” [93], which is in publication at the time of writing this thesis. It was written in collaboration with Michael Rohs. Benjamin Simon and Kamillo Ferry helped to conduct the experiment in this chapter as assistant researchers. Maximilian Schrapel helped to integrate his MultiWave prototype as an additional example into the final system. This chapter includes an additional study compared to the INTERACT paper, which was scrapped from the paper due to a conference-enforced page limit.

We used the third main prototype (see subsection 4.2.3) for this chapter’s experiments.



Figure 6.1.: The final version of *VRTactileDraw* in action. Users wear an HTC Vive Pro VR headset and a tactile user interface. A user draws symmetric strokes on the 3D model head in VR (left) and then replays the resulting pattern (right). During drawing, the user can feel the resulting tactile actuation.

6.1 Introduction

Beyond haptic renderings, which may be realized using physics simulations (e.g., contact or impact forces), tactile patterns (TPs) can be used for abstract concepts such as eliciting emotions or guidance during navigation. However, TPs' design for such abstract concepts requires manual exploration of the design space and is also interesting for non-technical people (e.g., for personalized touch sensations between remote humans). With the emergence of high-fidelity haptic feedback, the demand for interfaces that can be used to design TPs and effects rose as well. Several works appeared in the recent past [32, 46, 74, 152, 172, 190] but none of these approaches is designed for a high number of actuators that may be spatially oriented in more complex shapes than just a 2D grid.

This work introduces a pattern design interface for tactile feedback systems that feature many actuators in complex spatial arrangements around the body (e.g., [17, 20, 38, 47, 99, 114, 124, 187]). The need for a pattern designer for systems including many arbitrarily distributed actuators on the human body can be further motivated by the wide range of novel use cases the systems mentioned above enable. For example, full-body suits potentially enable the feeling of physical closeness to a remote person by "distantly touching or brushing" a model body in any desired location, effectively creating a real-time TP. Another example would be creating TPs for an immersive movie where viewers wear a tactile system when watching action-packed scenes and feel specifically designed effects on their body. Imagine feeling the shockwave of an explosion or a giant spider crawling up your spine and over your head before finally becoming visible in the movie scene from above.

In our prior work [94], an iterative design process was followed, including several design and implementation phases and two think-aloud studies with feedback from technical and non-technical users. The goal was to develop two variants of an intuitive TP designer (see Fig. 6.2):

- A *curve interface*, which behaves like an audio/video editor, allowing the user to modify the intensity-over-time curve of each actuator.
- A *drawing interface*, which allows the user to directly draw actuation strokes onto a body part, with interpolation between the actuators.




Figure 6.2.: Curve interface (left) and drawing interface (right) from prior work [94].

With the TP designer from our prior work [94], created patterns can be played while drawing. Heat maps provide a live visual representation of vibration intensity. Users who wear the tactile feedback system can simultaneously feel the created pattern. The two variants have different advantages and disadvantages, but moving on, we focused on improving the drawing interface, as users tended to prefer it over the rather complex curve interface. Some of the most severe disadvantages of the drawing mode are related to the 2D user interface. In particular, the following tasks are difficult to perform with the 2D interface:

- Drawing a stroke on a non-flat body part from a 2D camera perspective as this leads to distortions.
- Moving the camera around the modeled body part while drawing a stroke.
- Adjusting the stroke intensity level while drawing a stroke.

These issues with the 2D interface led us to develop a VR user interface for the same purpose. A VR interface can address the above challenges, as it offers increased spatial awareness, ease of moving around a 3D model by simply walking around, and a more direct spatial mapping when drawing. While developing the VR interface, we had the following research questions in mind:

- RQ1 - Usability: What kinds of interactions are suited best for designing tactile patterns in VR, how can the required interface functionality be made as simple and intuitive as possible, and what levels of and usability does the designed VR interface achieve?
- RQ2 - Comfort: How can the VR experience be made comfortable for the user, and what levels of and comfort does the designed VR interface achieve?

The final version of the resulting *VRTactileDraw* system is shown in Figure 6.1.

6.2 Background and Related Work

We first give an overview of tactile feedback systems without going into too much detail, as the specific actuator configurations and application areas are less relevant to this work. Then we discuss several TP editors and conclude with the specific prior work on which this chapter is based.

6.2.1 Tactile Feedback

The fundamental related work on tactile perception around the head (subsection 2.3.8) and tactile displays (section 2.6) also applies to this chapter. Specific related work concerning various complex tactile systems and the design of tactile patterns is presented below.

Except for full-body, head-worn, and vision substitution systems, most of the systems mentioned in section 2.6 feature a relatively low number of actuators in a simple configuration and thus require only moderate work to define meaningful TPs manually. However, there are systems with large numbers of actuators or complex actuator arrangements [17, 20, 38, 47, 99, 114, 124, 187], which pose an obvious challenge to the design of TPs. They currently require a significant amount of manual work by a pattern designer or algorithmic support to generate meaningful TPs due to their complexity. In such situations and without the support of a suitable interface, creating high-quality TPs is a daunting task. These use cases are likely to profit most from the proposed system, as it is expected to drastically reduce the amount of work needed for generating meaningful TPs, it enables fast prototyping, and allows even non-technical users of the system to define their own TPs.

6.2.2 Tactile Pattern Editors

Seifi et al. [190] evaluated three possible interfaces that allow users to define TPs for a single actuator in a Wizard of Oz study. They conclude that users want more control

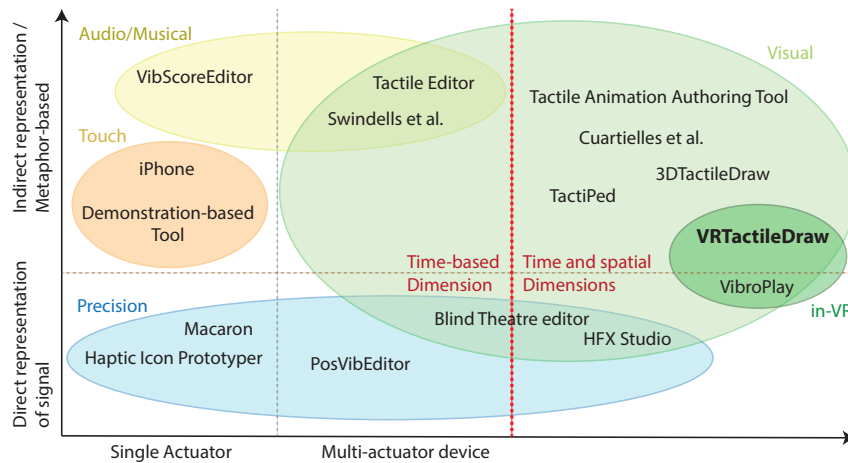


Figure 6.3.: Design space of TP editors after [152], modified to include *VRTactileDraw* and various other recent works.

(intensity, roughness, and rhythm) over TPs than a simple choice among preset patterns. Hong et al. [74] propose the “demonstration-based authoring” method, in which a user taps on the screen and the system generates a corresponding TP, representing the user’s taps. The *Hapticon* Editor [46] can be used to define a special waveform as a TP for a single actuator. *VibScore* [119] is a musical score-like design interface for TPs. The accompanying *VibScoreEditor* can be used to design a TP for a single actuator using musical notes. *Macaron* is a web-based TP editor for single actuator TP design [184]. Swindells et al. published an editor to merge simple haptics with audio and video channels [203].

Cuartielles et al. [32] briefly present four different approaches to TP editors and give recommendations on how a TP editor should allow interaction (touch-based, not overly complex, and allowing immediate replay of the created pattern). The “Blind Theatre Editor” (2009) mentioned in [32] is an approach to generate TPs for up to 64 haptic actuators. However, while it allows defining an intensity curve for each actuator individually on a timeline, like the curve mode in [94], it is also quite complex, shows just a rough indication of spatial actuator position, and can only be used by technicians [32].

The *posVibEditor* [172] aims to support TP design for multiple actuators and arbitrary waveforms but lacks in selecting a particular actuator from a large number of possibilities as the actuators are not visualized on the affected body part. Moreover, each actuator curve has to be defined individually, which leads to a high degree of control on the generated pattern but is, in turn, rather complex. The *bHaptics Designer* [17] by bHaptics Inc. is similar to the *posVibEditor*. It was released alongside a commercial tactile suit and accessorizes for VR gaming and also features grid and timeline editing of TPs using the mouse and keyboard [17].

TactiPEd [152] is a visual interface to easily prototype TPs for multiple actuators. [152] surveys several systems and approaches and discusses their strengths and weaknesses, including some of the ones mentioned above (see Fig. 6.3). The philosophy behind *TactiPEd* is closely related to this work, as the goal is to provide a simple visual interface to define TPs for multiple actuators. While [152] works for actuator arrangements of up to 6 to 8 actuators, such as in a wristband, it is unclear whether it scales to more complex spatial arrangements due to the missing spatial 3D mapping.

The *Tactile Animation Authoring Tool*, presented with an algorithm to interpolate between actuators in grid displays, is similar to the algorithm used in this work. It allows users to draw strokes on a 2D grid representation of actuators, thus enabling rapid prototyping of TPs for 1D or 2D tactile grids [183]. It has not been released publicly nor validated in usability studies with standardized usability measures.

HFX Studio [33] is a haptic editor to design patterns for VR use cases with the human perceptual model in mind. The intention is to design haptic patterns for VR, but the actual pattern editing is performed outside of VR using Unity timelines with a mouse and keyboard. Thus, while the aim is related to the aims of *VRTactileDraw*, the interaction concept is very different. In principle, HFX Studio can design TPs of similar spatial complexity like the ones resulting from *VRTactileDraw*. However, the effort and time investments in designing such complex TPs are considerably higher than with our approach: In Danieau et al.'s [33] study task of drawing a tactile arrow on a model, users took more than three minutes on average, while the same task takes just a few seconds with *VRTactileDraw*. However, HFX studio is one of the most advanced haptic editors, as it aims for high precision and even allows defining effects for other modalities (like thermal and pressure) in addition to vibration.

VibroPlay [76] is the first short concept of an in-VR TP editor that allows direct manipulation of a pattern by touching actuators in a model (only supports binary on or off), which is distinct to our approach. The concept was neither fully documented nor tested in a usability study.

Finally, *3DTactileDraw* [94], mentioned in the introduction, shows a first implementation of two possible user interfaces for designing TPs using strokes on a model for a large number of actuators in arbitrary configurations. The implemented curve interface (see Fig. 6.2, left) is closely related to many of the approaches mentioned above. The drawing interface (see Fig. 6.2, right) pursues a novel approach to define TPs. With the drawing interface, individual actuators no longer have to be defined separately. Instead, the user can draw a stroke on the model of a body part. An interpolation algorithm, first published in [98, 99] and similar to *TactileBrush* [80, 183], takes care of modeling the resulting TP. This drawing interface was preferred by most participants of the user studies that compared the curve interface and the drawing interface [94]. However, due to the nature of a 2D user interface and mouse and keyboard input, this system has some inherent limitations. For this reason, we decided to realize the

TP editor in a VR environment. Unlike most previous related works, the simple and intuitive design of *VRTactileDraw* focuses on novice users.

6.3 Iterative Design Process and Implementation

We started by reviewing TP design interfaces from related work and settled on using the Unity IDE for VR scene modeling with the HTC VIVE Pro, including the wireless add-on, as the basis of our system. The HTC VIVE has left and right-hand controllers, which are used simultaneously. These controllers feature several buttons, including an analog trigger button and a touchpad, which offer various possibilities for designing VR interfaces.

6.3.1 Virtual Reality Interface Design

Apart from appropriate hardware, a suitable user interface design is needed for effective tool selection and to prevent so-called VR sickness (in particular nausea) from occurring [118]. Previous work on VR sketching showed prototypes of, e.g., a color selector menu [37]. We mostly followed the guidelines by Sherman et al. [194] for general guidelines, Alger [3] for recent VR interfaces, and Lin et al. [123] for measures to prevent VR sickness.

Lin et al. [123] recommend using *independent visual backgrounds* to reduce VR sickness. With RQ2 in mind, we implemented this by using a low-poly background from the Unity asset store. We chose a background with few environmental features (e.g., mountains and trees) to prevent distraction and for performance reasons. The VR environment should be rendered at least as fast as the VR headset refresh rate to reduce the likelihood of VR sickness and nausea [194, 209].

Alger [3] defined different zones for placing interfaces in VR around the user in a stationary setup, in which the user cannot freely walk around. Even though we decided to design a free-roaming VR experience, as we need users to move around the model, we still follow his recommendation of placing no static interface elements closer than 0.5 m to the user. Disregarding this recommendation would cause squinting, which in turn leads to lower user comfort (RQ2). Furthermore, we tried placing most of the interface elements in the defined workspace and content zones, which are right in front of the user as outlined by Alger [3].

Alger recommends using radial menus anchored on the controllers or hands themselves combined with a touchpad [3]. We found this to be a helpful recommendation for implementing several different options of the interface for the HTC Vive controllers as

we aimed to design our interface as intuitive and straightforward as possible (RQ1). The final mapping of interface functions to controller buttons is shown in Fig. 6.10.

6.3.2 Goals and Basic Features

We defined the following set of goals for the system, which originate from earlier work [94], VR interface design guidelines mentioned above, our research questions defined in the introduction, and pilot testing:

- G1** Make the VR interface easy and intuitive to use so that non-technical and first-time VR users can still design their imagined TPs for a complex tactile interface on an exploratory basis without having to rely on external documentation or training.
- G2** Implement a state-of-the-art VR interface instead of just adapting an existing 2D interface and putting it into VR.
- G3** All settings within the interface should be within reach and easily usable, comparable to *Tilt Brush* [59] and adhere to the design guidelines by Alger [3].

We identified the following basic features of the VR TP design interface, which originate from pilot testing and adapting our prior work in VR [94]:

- Drawing strokes of varying intensity levels by “laser pointing” at a modeled body part.
- Instantly rendering the stroke that a user is currently drawing as a TP.
- Selecting existing strokes by pointing and clicking or through menu buttons.
- Replaying the entire pattern, which may consist of several strokes that may overlap in time.
- Changing the TP playback speed seamlessly within a certain range (e.g., 0.1-3.0 times the original speed).
- Jumping to a desired position within the TP and looping the TP.
- Changing the start time of a stroke relative to other strokes after drawing it.
- Deleting strokes and patterns.

6.3.3 Design and Implementation

We started designing a VR interface for TPs from scratch based on related work, our goals, and essential functions outlined above. In our first sketching sessions, we came up with the concept of two radial controller menus as suggested by Alger [3], where one “draw” controller in the user’s dominant hand would be responsible for the main functions (e.g., draw, delete stroke/entire pattern, return to the main menu). In contrast, the “select” controller in the user’s non-dominant hand is mapped to

secondary functions, such as stroke selection, pattern playback speed, start/pause, and current position in the pattern. Furthermore, we designed a simple main menu for choosing between the available patterns, specify global settings, and potentially delete patterns.

We initially implemented the same interpolation algorithm as in our prior work [94, 99] (see section 8.4), but finally settled on a different algorithm. The original algorithm was targeted explicitly at guidance, whereas in this work, each actuator's intensity depends on one or more drawn strokes. The original guidance algorithm would sometimes stimulate an actuator farther away to get a person to move their head in that direction. In contrast, the new algorithm always picks the actuators closest to the stroke. The new interpolation algorithm works as follows: For a single position on a stroke, we gather the $N = 3$ closest actuators. These are driven at an intensity proportional to their distance to the stroke position, normalized over all N actuators, and multiplied by the stroke's intensity at the current position (0..1). In case multiple stroke positions of different strokes affect an actuator at the same time, the results are added up per actuator and capped at the maximum intensity. This algorithm is less suitable for tactile interfaces that do not feature a dense layout of tactile actuators. For example, HapticHead has closely spaced actuators to take advantage of the tactile funneling illusion, which may cause users to perceive stimulations as smoother and in-between actuators (see chapter 3). The interpolation algorithm can easily be replaced by a more appropriate algorithm for less dense arrangements in the open-source release of VRTactileDraw.

We used the third main HapticHead prototype (see subsection 4.2.3) and the Unity v5.6.6f IDE for the implementation. It is designed extensible so that it can easily be used with a variety of other tactile interfaces.

6.4 Preliminary User Study

After implementing a first, still rather rough version of the interface, we conducted a preliminary think-aloud user study to fix the most severe issues before conducting the main study.

6.4.1 Design and Study Tasks

We chose the think-aloud method [213] for our user studies to expose usability flaws. In a think-aloud study, the user interface should be self-descriptive so that the user can work on a given task without any advice from the experimenter [213].

To make participants fully explore the possibilities of the interface, we designed a set of 10 tasks. Most of these tasks are open so that the participants may take various approaches and may reach different results.

Specifically, the tasks ask the participants to design a TP which:

1. asks the user to stop,
2. notifies the user of an up-leading staircase,
3. asks the user to turn right,
4. asks the user to crouch,
5. notifies the user of a hazard behind,
6. asks the user to look up,
7. warns the user of a future earthquake,
8. asks the user to run forward,
9. lets the user feel a slow heartbeat, and
10. lets the user feel a simultaneous vibration left and right.

Tasks 1-8 represent general use cases of the tactile interface. Task 9 was chosen so that users experiment with the “Loop” feature and task 10 requires users to experiment with the “set stroke time” feature.

6.4.2 Implementation

We implemented the aforementioned counterbalanced tasks in the interface by embedding the textual instruction statically in the scene’s background so that users are constantly reminded of what they are currently working on.

6.4.3 Procedure

After reading and signing an informed consent form and optionally a photographic release form, we explained the think-aloud study method [213] to the participant. For hygienic reasons, participants wore a balaclava under the HapticHead prototype, which played the TP. On top of that, they wore the HTC Vive Pro, which rendered the VR scenes. The only other instruction to each participant was that they should go back to the main menu after finishing a task, possibly take a break and then start the next task.

Before starting the actual experiment, we tested each of the 24 actuators of the HapticHead individually to make sure there were no defects. The participant then started working on the first task after tapping the appropriate “new pattern” button in the main menu. A balanced Latin Square counterbalanced tasks to distribute order

effects. In case the participant gave us consent, we also recorded the entire session on video for later analysis. All participants gave us consent for video recording, at least for internal uses in this preliminary study. In the end, the participant filled out a final questionnaire containing several Likert scales and comment fields about his or her experiences. The participants received a bar of chocolate as a small sign of gratitude.

6.4.4 Participants

Seven participants with technical backgrounds participated in our preliminary user study (all men, mean age 22 years, SD 2.3 years). Four of them already had some prior experiences with VR headsets, and all of them used to draw less frequently than once a month.

6.4.5 Results

While we received overwhelmingly positive feedback on the software prototype, we focus on the negative aspects and suggestions for improvements as we are looking to enhance the implementation and add missing features.

Some participants asked for a controller mappings figure as they needed a bit of time to realize that the controllers had menus. Others asked for a hint that the left controller can be used to select strokes in addition to the stroke switch option of the right controller's radial menu.

Other participants commented that the extreme values of the "set player time" and "set stroke time" functions were complicated to reach, as the edges of the VIVE controllers' touchpads did not react appropriately to the input.

We noticed that not all participants intuitively found out about the HTC Vive controller trigger button's analog input feature and were thus always drawing strokes at full intensity.

Five out of seven participants missed a better way to know the current playback time or relative position in a pattern. This was only shown as numbers on the select (left-hand) controller in a relatively small font. Three participants suggested a function to delete all strokes at once and to clear a pattern. The initial prototype only allowed for the deletion of strokes one by one and after a confirmation dialog. Two participants requested a function to mirror strokes onto both sides of the head, particularly for task 10 (simultaneous vibration left and right). One participant asked for an option to show actuator positions on the head visually.

6.4.6 Discussion and Implemented Improvements

Our preliminary study uncovered some missing features and design flaws, which we corrected before conducting the main study to improve our system’s usability (RQ1). Since the participants requested a better visualization of the current playback time and relative position in the pattern, we designed a head-up display (HUD) as shown in Fig. 6.4, embedded in the upper center of the user’s field of view. The HUD shows the current and total time in the pattern, an overlay of all strokes with their respective colors, the selected pattern speed, and whether pattern looping is turned on. It also seamlessly shows the current position on the timeline while changing the current time. This makes it much easier to use these functions without looking at the controller. We also decided to show the current stroke intensity percentage as a hint that stroke intensity is proportional to force applied to the trigger button (hidden while no stroke is being drawn).

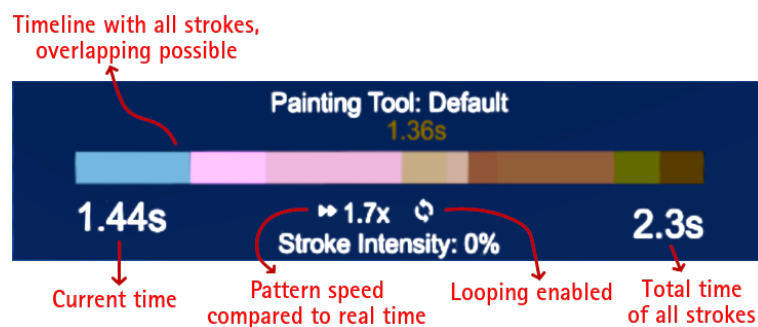


Figure 6.4.: First version of the head-up display with annotated elements.

Instead of providing a function to delete all strokes at once, we implemented a multi-tap for the “delete stroke” button so that users are now able to delete strokes in rapid succession without a confirmation dialog. We decided not to implement an undo function for this feature, as it is easy to draw a new stroke in case of accidental deletion. Deleting all strokes at once is also possible by deleting the entire pattern in the main menu.

Furthermore, we simplified reaching the extreme values of the “set player time” and “set stroke time” functions by limiting the coordinate system of the VIVE controller touchpad to an easily reachable area.

Our initial implementation picked random colors for the strokes. Thus, sometimes strokes could not be distinguished as they were too similar or contrast was too low. We improved color selection by picking colors from a color alphabet designed by Green-Armytage [60]. It ensures a high contrast between the colors and works well on bright backgrounds.

6.5 Main User Study

The main study had 17 participants and was open to participants from non-technical backgrounds. The study design was almost identical to the preliminary study. However, tasks 5 and 8 of the preliminary study had resulted in very similar patterns for most participants. Thus, we exchanged task 5 for a new task, designed to evaluate the actuators' different possible intensities by varying the trigger button pressure. Some of the participants in the preliminary study did not apply this feature. The new task 5 was: "Design a tactile pattern that lets the user feel a growing tension."

Furthermore, since we wanted to measure the final usability (RQ1) and comfort (RQ2) of our system, we added *system usability scale* (SUS) [22] questions to the final questionnaire, and users also filled out an *AttrakDiff* questionnaire [66].

6.5.1 Implementation and Procedure

The study task implementation was the same as in the preliminary study, except for the replaced task 5. Participants also went through the same study procedure as before, except for the additional *AttrakDiff* questionnaire at the end.

6.5.2 Participants

We invited a total of 17 participants (one woman and 16 men, 12 technical and 5 non-technical backgrounds, mean age 24 years, SD 3.4 years) for the main user study. None of these participants had participated in the preliminary study. Eight participants had prior experiences with VR headsets, and 11 frequently use game controllers. Six participants indicated drawing about once a month, five about once a year, and the others never performed any drawing activity.

6.5.3 Results

The final questionnaire was split in SUS questions [22] (Fig. 6.5), questions on intuitiveness and the design of interface elements (Fig. 6.6), and the comfort of using the VR interface, HapticHead, and VR controllers (Fig. 6.7). The system usability scale (SUS) score of our system reached a mean of 79.7 (SD=11.2). The *AttrakDiff* questionnaire [66] was external as we used the official one¹. Results of *AttrakDiff* and a detailed word pair analysis are shown in Fig. 6.8.

¹<http://attrakdiff.de/index-en.html> – accessed September 04, 2020

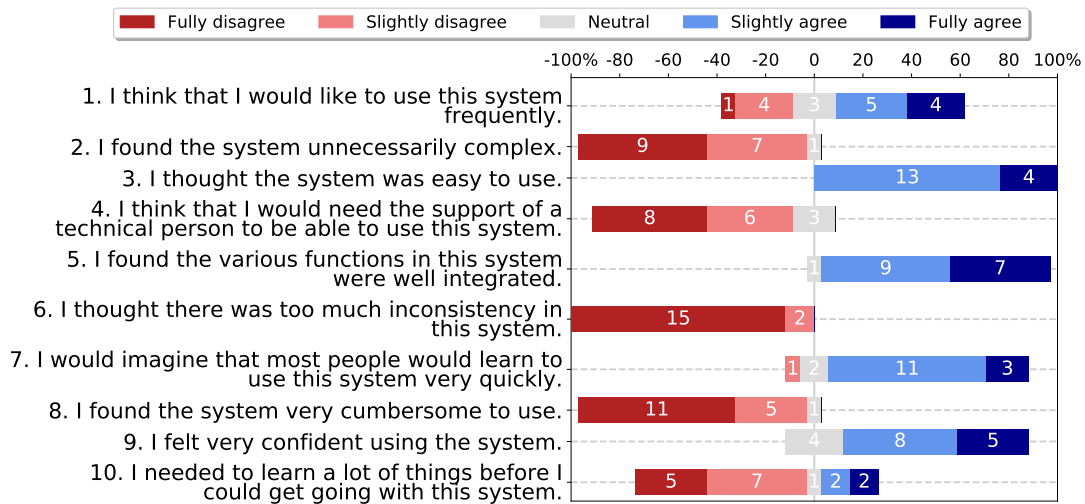


Figure 6.5.: System usability scale (SUS) questions and answers.

Our participants generally designed quite different patterns, but some general similarities could be found among the responses as seen in Fig. 6.9.

6.5.4 Discussion

Generally, the VR TP designer was well received by the study participants. A SUS score of 79.7 is between “good” and “excellent” according to [12]. A meta-analysis of 5000 SUS evaluations showed that a system with a SUS score of 80.3 is better than 90 % of all evaluated systems [22, 177]. The AttrakDiff scores show a similarly positive result. The system is generally rated as desirable (Fig. 6.8), with a little higher pragmatic than hedonic quality (answers RQ1). However, the AttrakDiff result is influenced by two questions that generally penalize VR systems: Since the system is a solo VR experience, it is obviously more isolating than connective (the real world is completely shut out), and it separates the user from the world and other people around, instead of bringing them closer together (see Fig. 6.8, right). Without these two categories, the already good AttrakDiff scores would likely be even higher.

There are possible solutions to make users feel less isolated from the world and other people. For example, it is possible to put the entire experience into an AR context instead of VR, using a device like the Microsoft HoloLens. This would make users feel more connected to the real world and further bring them closer to other people. Multiple users could collaboratively work together on a single pattern: All participating users could feel their collaboratively created pattern. Patterns could be discussed immediately, and new ideas could be brainstormed, explored, and refined together.

Fig. 6.6 shows that the overwhelming majority of the participants liked the concept, implementation, and intuitiveness of the design (answers RQ1). Nonetheless, the

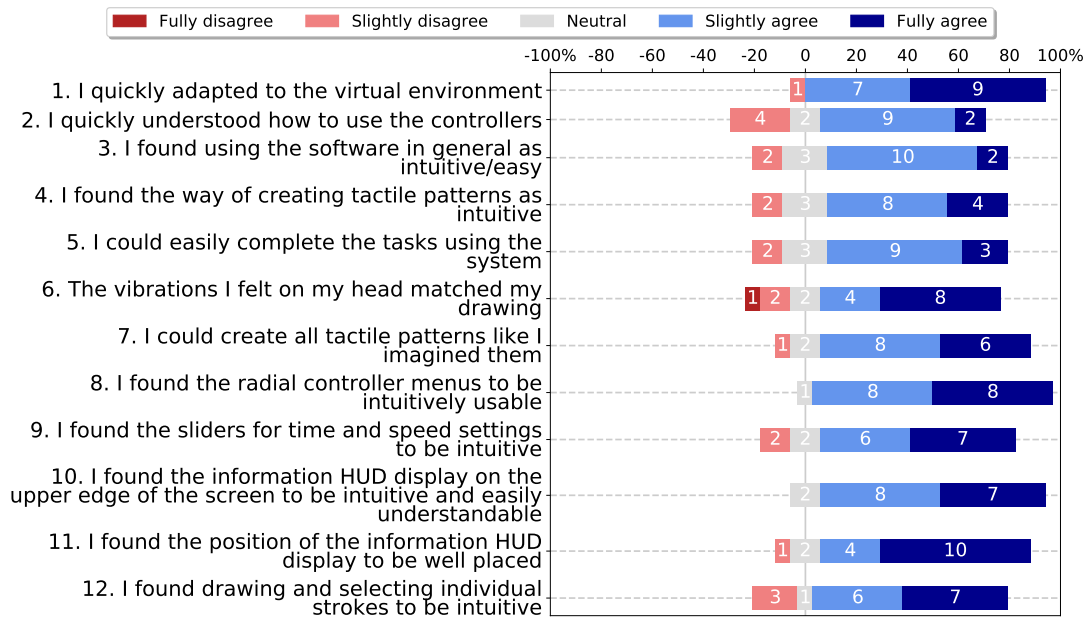


Figure 6.6.: Questionnaire on intuitiveness and interface design.

following analysis concentrates on the few negative responses we got in the questionnaires.

P2 strongly disagreed that the vibrations he felt on his head were matching his drawing. This participant tried to specifically target individual actuators and insert pauses of no actuation, which was difficult with the prototype and chosen interpolation algorithm as it incorporated the $N = 3$ closest actuators instead of just one. While it is possible to add pauses with the “set stroke time” function, this is not as intuitive as some kind of “record mode,” which this participant suggested instead of stopping the time while a stroke is not being drawn. P8 was the other participant who was mostly negative on the questions shown in Fig. 6.6 and answered positively on some of the tiring questions in Fig. 6.7. At 33 years, P8 was the oldest participant and did not like the concept of vibrations on the head.

Regarding comfort, Fig. 6.7 shows that the majority of users were happy with the design and software prototype (answers RQ2). Some participants could not adjust the VIVE Pro headset in such a way that they could sharply see the VR scene and interface. While we made sure that the headset was adjusted correctly at the beginning of the study, it may shift slightly on the head during the study, which leads to the display surface not being in focus. Also, three participants did not like the VR headset with HapticHead below it and described it as disruptive or annoying. However, only one of our 17 participants would have preferred working with a standard 2D UI instead of a VR interface. Five of the participants agreed that working with the controllers was exhausting or tiring. We can relate this to the clunkiness and weight of the VIVE controllers (203 g each). Future systems might offer different controller types, e.g.,

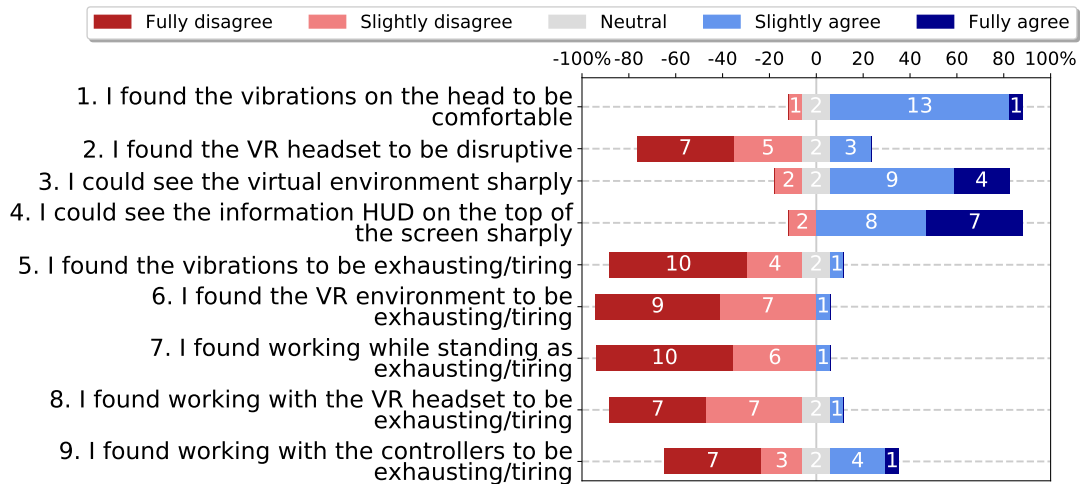


Figure 6.7.: Questionnaire on comfort.

Valve index controllers², which naturally fit around the hands without needing a permanent firm hand grip.

Another interesting finding is the appropriateness, complexity, and similarity of the TPs designed by our participants. While appropriateness is difficult to measure as it is highly individual, pattern complexity varied strongly between participants. Some used only a single short stroke, while others used multiple strokes of varying intensity for the same task. Not all of the designed patterns were similar, but certain tasks had a considerable level of agreement. Fig. 6.9 shows a number of example designs. The task to design a TP for notifying the user of an earthquake evoked mostly high-intensity patterns, which spread over multiple regions of the head, while other tasks, e.g., the task to encourage the user to run forward, produced simpler patterns of lower intensity. An interesting avenue for future research may be a study on how users define and recognize their own TPs and how the TP designs of different users differ in certain relevant aspects (e.g., intensity, complexity, and head region).

In terms of usability comparison vs. related work, we cannot directly compare our system’s usability against other TP editors, as no other work we know of published standardized usability measures such as SUS or AttrakDiff. In terms of performance, our system is clearly faster for rapid prototyping of complex TPs than TP editors based on timelines or intensity curves for individual actuators. For example, HFX Studio [33], as one of the most sophisticated TP editors, conducted a study in which participants had to draw a tactile arrow on the back of a torso model. Their users took more than three minutes on average to complete the task, while the same task using *VRTactileDraw* is as simple as drawing an arrow in any drawing application and only takes a few seconds. The only other work using a similar algorithmic interpolation approach as *VRTactileDraw*, the *Tactile Animation Authoring Tool* [183], is likely to have

²<https://www.valvesoftware.com/en/index/controllers> – accessed August 20, 2020





Figure 6.8.: AttrakDiff: Portfolio of results (left), Diagram of word pairs (right).

a similar TP prototyping performance, but only in the case of a 1D or 2D tactile grid. However, we are not aware of a public release of the tool or any usability evaluation with standardized usability measures.

6.5.5 Improvements After the Main Study

We implemented the following features that the participants suggested in the main study:

- *Set stroke duration* function: Allows changing the total duration of an individual stroke after drawing it with visual feedback (changing bar length) in the HUD. 8 participants suggested this.
- *Symmetric mirror drawing mode*: Strokes on one side of the model are simultaneously drawn on the other side of the model (symmetric relative to the mid-sagittal plane). This was requested by a total of 5 participants and generalizes to other systems beyond HapticHead due to the human body's symmetry.



(a) Task 7, P6 (b) Task 7, P16 (c) Task 8, P7 (d) Task 8, P17

Figure 6.9.: Example TP designs by the participants of our main study. Task 7 was to design a pattern to notify the user of an earthquake, while task 8 encouraged the user to run forward. The heat map in (c) and (d) shows the stroke’s start and its intensity.

- *Show actuator positions* option: This option in the main menu shows the positions of all actuators so that users can consider actuator locations while drawing patterns. 4 participants requested this.
- *Record mode*: While pushing a record menu button on the select (left-hand) controller, time moves on even while not drawing a stroke compared to the regular mode, in which pattern recording time stops when no stroke is being drawn. A red recording icon on the HUD shows whether the record mode is active at a given time. This function was only suggested by a single participant. We still found it to be a worthwhile addition.
- *Overview figure*, shown at startup close to the main menu. It highlights the mapping of several functions to the controller menus and buttons. While this was not explicitly requested, we think that it is necessary to show all possible functions of the interface in a single easy to understand overlay so that there is no confusion about how to set intensity levels while drawing a stroke, how to change drawing modes, or how to return to the main menu.

With these changes, there are two drawing modes: *normal* and *mirror*. Switching between these modes is performed by the VIVE controller grip button. Otherwise, most of the new functions can be reached through the controller menus’ final mapping (see Fig. 6.10).

We also polished the overall look and feel of the HUD by adding a vertical red time indicator, choosing appropriate background transparency, and hiding elements that are not strictly necessary in a given context (see Fig. 6.11).




Figure 6.10.: Mapping of the controller menu in the final prototype. Some functions were not present in the preliminary and main user studies. Left: selection controller, right: drawing controller. Menu items can be selected by moving the thumb in the appropriate direction on the touchpad and then pressing down. Pressing down multiple times without fully lifting the thumb triggers the action repeatedly. Some menu items lead to an adjustable slider, which is confirmed by pressing down (bottom center).




Figure 6.11.: Final version of the HUD. Recording mode is currently active and a stroke is being drawn.

Even though related work [183, 185] found that users might want to manipulate created paths after initial creation, only three participants would have liked an option to manipulate stroke intensity after creation. No participant suggested altering the strokes spatially. Thus, we chose not to include an option to manipulate strokes, as the current implementation allows rapid deletion and re-drawing of a stroke, which is only marginally slower than selecting a different tool and manipulating a stroke.

6.6 Limitations

While our software prototype already received predominantly positive feedback, we did not conduct a follow-up study of the changes made after the main study. Since we implemented features that the participants suggested and that we considered valuable, we estimate the SUS and AttrakDiff scores to improve with these changes.

Nine of our 17 participants in the main study had no prior experience with VR systems. Thus a novelty effect [113] cannot be precluded. Even if we had recruited only participants with prior VR experience, a novelty effect could still occur as the TP editor itself would be novel to them. Future work may implement a similar system and conduct a long-term study, as the novelty effect wears off over time [113].

6.7 Open Source Release and Extension to Tactile Interfaces



Figure 6.12.: *VRTactileDraw* driving a 52-actuator tactile vest mockup (left) and the users' view inside an HTC Vive (right).

The user studies above were conducted with our HapticHead system. However, the prototype can easily be extended to support different body-worn tactile feedback systems, like tactile vests or full-body suits [17, 20, 38, 47, 99, 114, 124, 187]. We provide an open-source release³ of *VRTactileDraw*. The interface is implemented in a modular, extensible way so that it is easy to replace the main components to fit specific requirements. In order to use the *VRTactileDraw* editor with a different tactile output system, a developer has to perform the following steps, further specified in the documentation:

- Use the provided models of the head and human body or replace them with a targeted body part model (e.g., chest).
- Create copies of the *Actuator Prefab* model, assign unique ids, and position the actuators according to the tactile system's specification. This is all done inside the Unity editor.

³*VRTactileDraw* open-source release: <https://github.com/obkaul/VRTactileDraw>



Figure 6.13.: *VRTactileDraw* driving the MultiWave prototype vest [187]: A wearable vibrotactile vest containing 76 actuators (left) and the user’s view inside an HTC Vive (right). The user is replaying a TP and simultaneously feels vibrations as indicated by the heatmap.

- Implement a script to drive the new tactile system, including the system-specific communication protocol. This script has to implement the *CommandSender* interface. The *RaspberryCommandSender* class may serve as an example.

With the open-source release, we provide a documented framework for developers to use on their own tactile prototypes. The original software prototype for HapticHead is included in the framework. Besides, we provide two other examples of two different tactile vests, including the respective 3D model of a human body. Figure 6.12 shows *VRTactileDraw* in action on a virtual tactile vest with 52 actuators around the torso and over the shoulders. The VR software components are fully implemented, but the tactile output systems are currently mockups only.

Figure 6.13 shows *VRTactileDraw* driving a fully implemented tactile vest system [187]. *MultiWave* is an FPGA-based controller for multiple tactile actuators connected to a tactile vest prototype, which is intended for mobile haptic feedback in virtual or augmented reality [187]. It consists of 76 actuators controlled by *MultiWave* via Bluetooth or WiFi. *VRTactileDraw* can be used to generate TPs for *MultiWave* and the tactile vest, e.g., to realize navigation for visually impaired people and to provide orientation feedback and effects in VR environments.

6.8 Conclusion and Future Work

We present the first full-fledged design, implementation, and evaluation of a TP editor where the actual TP editing process is entirely conducted inside a virtual environment.

VR interfaces are well suited for designing TPs for complex spatial arrangements of tactile actuators, as users can freely move around and generally have a better understanding of spatial relationships than in 2D UIs (see [45]). Our system is also more scalable in terms of the number of actuators and spatial actuator configurations than traditional timeline-based TP pattern design UIs. In traditional timeline-based TP editors, a designer has to define an intensity curve for each actuator. Simultaneously, in our approach, an interpolation algorithm calculates intensity curves of all actuators in real-time based on simple strokes drawn by potentially novice users.

VRTactileDraw is generally easy to use and understand, especially when considering the complexity of the created patterns generated out of simple strokes. Thus, we can answer the research questions defined in the introduction as follows: The presented VR interface with its minimalistic environment and self-explanatory interactions is indeed highly suitable to design tactile patterns as it allows rapid prototyping and the SUS and AttrakDiff scores indicate a desirable system. Most of our users were able to design TPs without any prior knowledge freely and rated it as generally intuitive (see Figure 6.6 and 6.8), answers RQ1) and it was highly accepted amongst our study participants in terms of comfort (see Figure 6.7, answers RQ2).

Apart from extending our system to other tactile prototypes, as mentioned in the previous section, another direction for future work was already hinted at in the discussion: Multiple users could collaboratively design TPs, either in the same VR environment and simultaneously experience the created patterns or in an AR variant that shows the natural environment, other users, as well as the designed TPs. This would make users feel less isolated from the world, and their colleagues would facilitate discussion about and refinement of patterns, and would probably lead to a better overall result. This would require that each user is equipped with a tactile feedback device.

These collaborating users would not even need to be at the same physical location but could be represented by avatars and work together at a distance, sharing the created patterns. Furthermore, one user could draw a tactile stroke on a body part, while another user would simultaneously feel the tactile feedback. This would allow for the exploration of novel use cases compared to simple TP design, like feeling the touch of another person, hugs, “cuddles,” and possibly experiencing a faraway person’s emotions.

Future work may also investigate the effects of drawing TPs from a third-person perspective and experiencing the synchronous tactile sensation on the process of designing tactile patterns and the final outcome.

All the use cases mentioned above of extending our concept to multiple users are another considerable advantage of our in-VR (or potentially AR) concept over traditional TP designers. These are not easily extended to multiple collaborating users working in the same environment.

Increasing Presence in Virtual Reality

A high level of presence is an important aspect of immersive virtual reality applications. However, presence is difficult to achieve as it depends on the individual user, immersion capabilities of the system (visual, auditory, and tactile), and the concrete application. We use HapticHead to further increase the level of presence users feel in virtual reality scenes on the way to reach this thesis's overarching goal to establish the head as a means for tactile communication.

In a between-groups comparison study, the vibrotactile group scored significantly higher in a standardized presence questionnaire compared to the baseline of no tactile feedback. This suggests the proposed prototype as an additional tool to increase users' perceived level of presence in virtual reality scenes.

This chapter is based on the INTERACT' 2017 paper "Increasing presence in virtual reality with a vibrotactile grid around the head" [95], written in collaboration with Michael Rohs. This chapter's experiment is based on the Bachelor thesis of Kevin Meier [134].

We used the first main prototype (see subsection 4.2.1) for this chapter's experiments.

7.1 Introduction

Current generation virtual reality (VR) head-mounted displays (HMDs) use high-quality visual displays to stimulate the user's visual sense for a variety of applications. Along with the visuals, typical VR applications also use spatial sound to increase the level of presence in certain situations, such as in horror games. However, the tactile channel is largely neglected due to difficulties in finding a suitable solution acceptable in terms of aesthetics, price, and implementation complexity.

It has already been shown that different kinds of tactile feedback increase the level of presence in a VR scene or game [42, 174, 175]. The purpose of this work is to investigate the effect of a vibrotactile around-the-head HMD on the level of presence that users experience in a VR environment. To do this, we implemented two VR scenarios enhanced by vibrotactile feedback and compared them to no vibrotactile feedback in a between-groups comparison study.

7.2 Related Work

The fundamental related work on tactile perception around the head (subsection 2.3.8) and tactile displays (section 2.6) also applies to this chapter. Specific related work concerning immersion and presence in Virtual Reality is presented below.

In this chapter, we use the terms immersion and presence following the definition by Slater [197]. *Immersion* refers to the objective level of sensor fidelity a VR system provides, while *presence* refers to a user's subjective psychological response to a VR system. Of course, the level of presence is somewhat correlated to the level of immersion a system can provide and depends on the concrete VR scene.

Pausch et al. [156] investigated the performance difference of a system with and without increased visual immersion in a search task scenario and found that users were faster in the immersive condition because they had a better spatial understanding of their virtual surroundings. Dinh et al. [42] evaluated the effect of haptic feedback (wind and heat) on user presence in VR. The effect of a vibrotactile vest and pants on user presence and collision detection in VR was explored by Ryu et al. in [173]. They found that their prototype did enhance the user's sense of presence, especially when combined with spatial sound. Just like the system by Ryu et al., our system aims to increase the level of presence in VR. However, we decided to give vibrotactile feedback around the head as such feedback could easily be integrated into existing VR HMDs instead of requiring additional garments.

7.3 Presence Experiment

We designed an experiment to measure the possible effects of vibrotactile feedback in VR scenes on presence. Our initial hypothesis was that appropriate tactile feedback increases the level of presence in VR scenes. We chose a between-subjects study design where the experiment group experienced the VR scenes with tactile feedback and sound, while the control group experienced the same VR scenes with sound but without tactile feedback.

A total of 20 participants (six female, mean age 24.1, SD 2.5 years) were invited and split up into equally sized experiment and control groups. Eleven participants had previous experience with VR HMDs and applications. All participants filled out a mandatory informed consent form and optionally a photographic release form.

7.3.1 Presence Experiment – Measuring Presence

In a measurement of psycho-physiological parameters for presence (e.g., electromyography or galvanic skin response), Nacke and Lindley showed significant differences for different experiment conditions [148]. However, these significant differences could not be clearly correlated to participants' responses in accompanying questionnaires, which means that the real causes and effects of those differences still need to be researched.

Instead of measuring presence in a quantitative psycho-physiological way, we decided to use a widely used qualitative method presented by Witmer et al. [220]. Their "Immersion Tendency Questionnaire" (ITQ) measures how likely a movie or game immerses a person in general, and the "Presence Questionnaire" (PQ) measures the level of presence a person was experiencing in a previous experimental condition. This method requires a between-subjects study design for comparable results as all other effects need to be canceled out.

Witmer et al. evaluated their original ITQ and PQ questionnaires, identified and removed irrelevant questions which did not correlate with the overall result and the participant's assessment of the level of presence [220]. We used these reduced questionnaires and translated the remaining questions to German (the native language of the participants). We further rephrased all of them as statements to measure agreement with 7-point Likert scales instead of the proposed 7-point semantic differential scales. We used Likert scales because they are easier to answer (always strongly disagree to strongly agree as options) and often less vague depending on the adjectives used in semantic differential scales, which can cause reliability issues [207]. We also added one statement to the PQ, letting participants state how well they could survey their environment due to haptic clues. Witmer et al. did not use any haptic stimulations in their experiments.

7.3.2 Presence Experiment – Virtual Reality Environments

Two VR environments were implemented as exemplary cases. These were augmented by tactile feedback.

Header Simulation

In our soccer header simulation (HS), four catapults and occasionally a plane above the user throw balls in the direction of the user's head. The user's position is right in the center between four goals.

At the beginning of the 20-minute header simulation scene, only a few balls are thrown towards the user one by one. The task of the user is to head the ball into any of the



Figure 7.1.: Soccer Header Simulation in VR.

goals. As time progresses, the frequency of balls being thrown increases, and it may also happen that multiple balls arrive simultaneously. In the last third of the simulation, planes appear occasionally and drop balls on top of the user to increase the challenge. Users can see a goal score indicator for further motivation.

Participants can visually locate the balls and hear spatial sounds from the catapults when a new ball is launched. In the vibrotactile feedback group, participants also experience vibrotactile stimuli interpolated between the three closest motors in the ball's direction. One stimulus is played when the ball is launched (100 ms), and another stimulus is played when the ball impacts the participant's head (200 ms).

Viking Village

In the Viking Village simulation (VV), we took one of Unity's standard environments called "Viking Village" from the Unity asset store and enhanced it with visual, auditory, and vibrotactile rain and snow effects.



Figure 7.2.: Enhanced Viking Village VR scene with different weather effects.

The user experiences a 15-minute camera tour on a wooden wagon, hears effects such as rain (5 minutes), snow (5 minutes), and crackling flames. Participants of the experiment group also feel rain or snow as single droplets hit their heads, causing small

vibrotactile stimuli at the nearest actuator position. A slightly less intense effect for snowflakes (100 ms, 75 % of the maximum intensity) than for rain droplets (200 ms, 100 % of the maximum intensity) is played. It is possible to play “catch the particle” by moving the head towards nearby rain droplets or snowflakes and then feel their impact at the corresponding position. Impacts on facial areas are not correctly modeled as there are no actuators between the forehead and chin.

7.3.3 Presence Experiment – Procedure

After arriving in our lab, participants were introduced to the Oculus Rift CV1. If they were part of the experiment group, they were also introduced to the HapticHead prototype. Besides, they were informed about the experimental procedure and filled out consent forms as mentioned above. Furthermore, they filled out the modified “Immersion Tendency Questionnaire” before starting with the first experimental condition.

Depending on whether the participants were in the experiment or control group, they experienced the soccer header VR scene, as explained above, with or without tactile feedback generated by the first main HapticHead prototype (see subsection 4.2.1). After finishing the first VR scene, there was a pause, and participants were asked to fill out the modified PQ.

In the second part of the experiment, participants experienced the Viking Village VR scene, again with or without tactile feedback, depending on whether they were in the experiment or control group, and filled out another modified PQ. We further asked the participants to estimate how much time had passed after each VR scene because an increased presence level can alter time perception [176]. Participants of the experiment group finally filled out another questionnaire on the tactile feedback’s usefulness and appropriateness.

7.3.4 Presence Experiment – Results

Results from the Likert scales in both the ITQ and PQs were summed up per participant, resulting in the total ITQ and PQ scores per participant as intended by Witmer et al. [220]. Even though Likert scales are usually ordinal scales, we also chose to report the mean and standard deviation due to large variances between participants and for comparability (Witmer et al. also chose to report mean, and standard deviation in their work [220]).

The modified ITQ yielded a median total score of 79.5 (mean 78.6, SD=12.6) for the control group and 91.5 (mean 87.7, SD=12.9) for the experiment group. A Mann-Whitney U test reveals a slight tendency but no statistically significant differences between the groups ($p=0.08 > 0.05$).

After the first VR scenario (soccer header simulation), the modified PQ yielded a median total score of 71.5 (mean 69.7, SD=14.1) for the control group and 91.0 (mean 87.7, SD=14.3) for the experiment group. A Mann-Whitney U test reveals a statistically significant difference between the two groups ($p < 0.01$).

When asked about the total time they thought the 20-minute soccer header simulation took, the control group participants (mean 26.5 minutes) estimated almost the same time as the participants of the experiment group (mean 26.0 minutes).

The modified PQ after the second VR scenario (VV simulation) yielded median total scores of 60.0 (mean 57.2, SD=17.5) for the control group and 77.5 (mean 74.7, SD=22.3) for the experiment group. A Mann-Whitney U test reveals a statistically significant difference between the two groups ($p=0.027 < 0.05$).

When asked about the total time they thought the 15-minute VV simulation took, again the control group participants (mean 18.8 minutes) estimated almost the same time as the participants of the experiment group (mean 19.0 minutes).

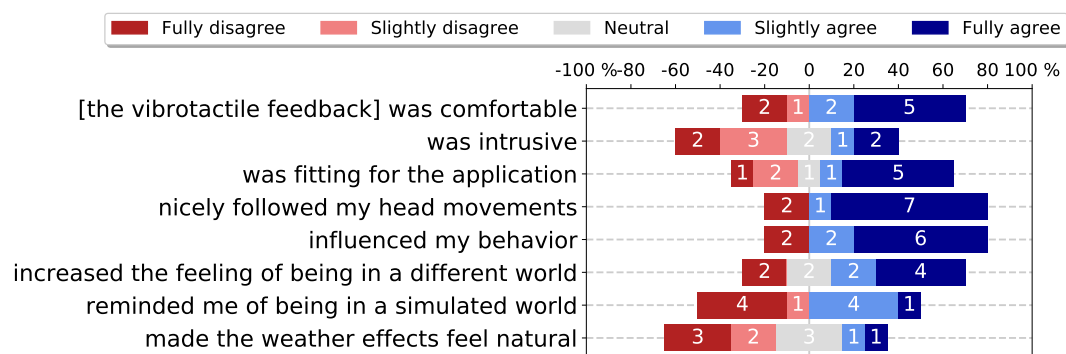


Figure 7.3.: Questionnaire on appropriateness of vibrotactile feedback

The experiment group participants filled out a questionnaire on vibrotactile appropriateness and usefulness at the end of the experiment. Results are shown in Fig. 7.3. Interestingly, there are some statements with very mixed opinions between participants. As a common pattern, some participants rated the vibrotactile feedback negatively in most of the statements, while others rated it positively in most statements which comes down to whether they liked it overall or not. Some participants who disliked the feedback also stated that they disliked the vibration motors' humming sounds. Eight of ten participants of the experiment group agreed that such vibrotactile technology should be incorporated into other applications.

7.3.5 Presence Experiment – Discussion

The difference between the experiment's immersion tendency scores and control groups is unfortunately close to being statistically significant. This indicates that the two groups were not perfectly balanced. Ideally, the ITQ-scores of both groups should be equal, and the p-value should be larger. We suspect that the number of 10 participants per group was too small to cancel out outliers effectively. However, as the difference in ITQ-scores between the groups is not statistically significant, we can directly compare the results of the presence questionnaires.

Both presence questionnaires yielded a statistically significant difference between the two groups. This means that the vibrotactile feedback definitely affected the presence users experienced in the VR scenes. The participants' duration estimates of the scenes were not significantly different between groups, but both longer than the actual time the scenes took. We suppose that this is due to the scenes being rather long, resulting in boredom and exhaustion.

The VV scene's estimated duration values were generally lower in both groups than for the soccer header scene. We suppose that this is caused by the much lower degree of interactivity of the VV scene compared to the soccer header scene. The former was essentially just a camera tour through the scene, and the participant could only move sideways a little and turn the head on the camera wagon. The scene was also quite long at 15 minutes, which led some participants to complain about being bored as to the low degree of interactivity. Furthermore, the vibrotactile augmentation did not make the weather effects feel more natural to most of the experiment group members. This indicates that the simulation itself could be improved.

7.4 Limitations

The relatively low number of 10 participants per group led to a tendency of the groups being almost statistically different in the ITQ results. More balanced groups would have been preferable.

The vibrotactile augmentation of the VV scene's weather effects did not lead to a more natural experience for most participants. Our concrete implementation lacked and could probably be improved using an iterative design approach for the weather effects.

7.5 Conclusion

The presented study shows that adding vibrotactile feedback around the head can significantly increase the level of presence users experience in certain VR scenes. This can be applied in immersive games and specific VR training simulations where the level of presence or spatial awareness (including collision prevention) is essential such as in complex maintenance jobs, anxiety therapy, or flight training.

HapticHead is a relatively complex prototype with questionable aesthetics, but it could be integrated with future VR HMDs to enable 3D tactile feedback in VR environments. This vibrotactile prototype's greater vision is to combine it with other tactile garments and produce a fully tactile and very immersive VR experience. This chapter is also a first step in finding implications for the design of vibrotactile feedback in VR scenes, such as what kind of vibrotactile feedback works best for virtual rain simulation.

Intuitive and Precise 3D Guidance

HapticHead’s initial idea originated from tactile belts used for guidance and navigation. HapticHead can provide 3D guidance and navigation, including depth information intuitively due to natural mapping instead of the 2D guidance that traditional tactile belts provide. This chapter will look at several studies on guidance in 3D, highlighting the outstanding performance a tactile interface around the head can achieve in precision, especially compared to 3D auditory feedback. This chapter presents another step to reach this thesis’s overarching goal to establish the head as a means for tactile communication.

Current generation commercial Virtual and Augmented Reality head-mounted displays usually include no or only a single vibration motor for haptic feedback and do not use it for guidance. We conducted three experiments investigating VR and AR guidance with HapticHead. These experiments indicate that our intuitive tactile guidance is both substantially faster (2.6 s vs. 6.9 s) and more precise (96.4 % vs. 54.2 % success rate) than spatial audio (generic head-related transfer function) for finding visible virtual objects in 3D space around the user. The baseline of visual feedback is – as expected – more precise (99.7 % success rate) and faster (1.3 s) in comparison, but there are many applications in which visual feedback is not desirable or available due to lighting conditions, visual overload, or visual impairments. Mean final precision with HapticHead feedback on invisible targets is 2.3° compared to 0.8° with visual feedback. We also successfully navigated blindfolded users to real household items at different heights using HapticHead vibrotactile feedback in an AR scenario, independently of a head-mounted display.

The content of this chapter is primarily based on the CHI’ 2016 “HapticHead – 3D Guidance and Target Acquisition through a Vibrotactile Grid” [98], and CHI’ 2017 “HapticHead: A Spherical Vibrotactile Grid around the Head for 3D Guidance in Virtual and Augmented Reality” [99] publications which I wrote in collaboration with Michael Rohs who advised me and proofread the papers. Furthermore, Kevin Meier conducted one of the experiments as part of his Bachelor’s thesis [134].

We used the initial prototype (see subsection 4.2.1) for the first experiment and the second prototype (see subsection 4.2.2) for the following two experiments.

8.1 Introduction

Navigation and 3D guidance systems use various technologies to stimulate the visual, auditory, or haptic channels. The visual channel is usually the channel of choice as it typically has a higher bandwidth than the other channels [217]. However, sometimes the visual channel is not the desired primary channel for some kinds of feedback or in special situations such as when driving a car [212]. The visual channel might be overtaxed, and essential feedback can be overlooked, or lighting conditions may prevent the user from seeing the feedback at all. Another reason to use the tactile instead of the visual or auditory feedback channels are faster initial reaction times, as shown in several studies such as [188].

To relieve the visual channel in guidance scenarios, we propose to use HapticHead to present 3D directional and distance information through moving tactile cues and patterns (see subsection 4.2.1). The head is well suited for guidance applications and tactile feedback, as it is sensitive to mechanical stimuli [57, 145] and provides a large spherical surface. These advantages allow displaying precise 3D information and allow the user to turn the head in the direction of a stimulus intuitively due to natural mapping [159].

We conducted three experiments to characterize user performance and refine our 3D guidance concept in virtual and real environments. In all experiments, targets may be located at any position around the head. This includes positions that are not in the visual field initially and positions above, below, and behind the user.

The *Guidance Comparison Experiment* is a follow-up of the experiment in our prior work [98], using the first main HapticHead prototype (see subsection 4.2.1) combined with an Oculus Rift DK2 to find visible virtual targets equally distributed on a sphere around the user. It is designed to compare vibrotactile feedback vs. visual and auditory feedback. In the previous experiment, we found indications that HapticHead feedback might be an exciting alternative to visual and auditory feedback but did not evaluate this in detail [98]. The current experiment includes more participants and additionally records movement trajectories and success rates individually per target. Based on the results, the participants' comments, and our observations, we refined the concept and prototype.

The *Guidance Precision Experiment* evaluates performance differences due to refinements of the prototype (see subsection 4.2.2) and the guidance algorithm. Furthermore, the achievable precision with visual (attention funnels and one-pixel targets) and vibrotactile (invisible targets) feedback is investigated.

The *Real-World 3D Guidance Experiment* is independent of the other experiments and aims to show the usefulness of our concept and second prototype (see subsection 4.2.2) for finding tracked physical objects around blindfolded users, i.e., without visual

feedback. In this experiment, the prototype was modified to provide an additional depth cue to the user via a dynamic vibrotactile pattern for true 3D guidance.

8.2 Related Work

The fundamental related work on tactile perception around the head (subsection 2.3.8), tactile displays (section 2.6), and specifically on indicating direction via tactile displays (subsection 2.6.1) also applies to this chapter.

8.3 Guidance Comparison Experiment vs. Visual and Auditory Stimuli




Figure 8.1.: Unity scene for the guidance to visible targets experiment

As a first step, we evaluate the performance of HapticHead (initial prototype, see subsection 4.2.1) in guiding users who wear an Oculus Rift DK2 towards virtual 3D objects around the head (see Figure 8.1). This allows us to refine the concept and prototype based on the findings.

We built a simple VR environment in Unity 5.3 that spawns 20 small ($r = 1$ m) equidistant spheres on the surface of a larger ($r = 5$ m) invisible sphere with the viewer at its center (see Figure 8.1a). The spheres were distributed with pack.3.20 coordinates [198]. As the user rotates the head, the spheres' location stays fixed with respect to the environment. The target spheres do not coincide with the actuator positions.

There are three feedback conditions in the Guidance Comparison Experiment: visual, auditory, and vibrotactile feedback. We included visual feedback as a baseline because AR and VR applications are usually designed around visual feedback. We also included auditory feedback, as auditory feedback is often used when applications aim not to overload the user's visual sense.




Figure 8.2.: Attention funnels with a tiny red crosshair in the view's center. Visual feedback from the user's perspective.

In the visual feedback condition, guidance towards objects is achieved through the concept of attention funnels as in [19] which are state-of-the-art 3D visual guidance cues utilizing "target goals" to guide a user. We implemented attention funnels using the same green target goals as in [19] and made sure that they also work with targets behind the user (Figure 8.2).

For the auditory feedback condition, we used white noise in combination with Unity 5.3's included spatial sound system (which uses a generic head-related transfer function, g-HRTF) with "spatial blend" set to 1 (full 3D) and Bose QC25 stereo noise cancellation headphones (NC off). We are aware that there are better audio technologies available that utilize personalized head-related transfer functions (p-HRTFs), but these require a complex per-user calibration. HapticHead does not require per-user calibration.

In the vibrotactile feedback condition, we activate the three actuators closest to the target with an interpolated intensity representing closeness. The closest actuator is running at the highest intensity. Actuator intensities are not static and are instead adjusted with head rotation. So as the user turns the head towards the target, the signal travels along the trajectory towards the head's front.

We invited 13 participants (2 female, mean age 23.5, SD 3.2 years). Only 5 had previous experience with VR HMDs.

Participants had to focus a "start box" at a fixed position for half a second to start a trial. Feedback was turned on, and the task was to find the target as quickly and accurately as possible. Once the participants had located and focused on the suspected target, they pressed a hand-held button to end the trial. Upon pressing the button, the sphere was visually highlighted in green or red, depending on whether it was the right

one. Instead of dedicated training trials, we chose this form of active highlighting to measure any learning effects and whether participants would "calibrate themselves" towards this new form of haptic feedback.

Each participant performed 480 trials: 3 feedback conditions (visual, auditory, vibrotactile) x 20 targets x 8 repetitions per target. The three feedback conditions were presented in blocks. Their order was counterbalanced with a Latin square. The order of targets was random. As dependent variables, we measured the head movement trajectory, task completion time, and error rate. The user's focus point was tracked with each frame at a rate of 75 Hz, which was also the update rate for all feedback conditions and data logging. Participants could pause between each trial and had a forced pause when the feedback condition changed. The experiment took around one hour per participant. As a reward, each participant received a bar of chocolate.

8.3.1 Guidance Comparison Experiment – Results

While running the experiment, we observed randomly appearing frame lags of 1 s + frame time due to a graphics driver issue. Because of this, we had to exclude 248 trials (3.97 %). To reduce the influence of outliers on task completion time, we used the median of the 8 repetitions per target that each user performed.

We define movement overhead (m.o. in Figs.) as the ratio of the actual path length to the optimal path length, minus 1, i.e., the movement overhead is a percentage of how much longer the user's trajectory is compared to the optimal path on the sphere towards the target. A value of 0 % means the user strictly follows the optimal path. A value of 100 % means the user's trajectory is twice as long as the optimal path. Movement overhead thus is a measure of how directly the user can localize the target.

Results - task completion times and success rates

	<i>Median trial time [s]</i>	<i>Mean trial time [s] (SD)</i>	<i>Mean movement overhead [%] (SD)</i>	<i>Success rate [%]</i>
Visual (attention funnels)	1.22	1.28 (0.37)	8.8 (3.6)	99.66
Auditory (g-HRTF)	6.28	6.86 (3.55)	67.2 (14.4)	54.22
Vibrotactile (HapticHead)	2.41	2.61 (1.04)	35.1 (10.7)	96.36

Table 8.1.: Task completion times and success rates for different feedback conditions.

Table 8.1 and Figure 8.3 show the measured dependent variables with merged data from all participants and all trials, not just successful ones.

The auditory condition was the slowest at 6.86 s. At 2.61 s vibrotactile only took 38 % and at 1.28 s visual feedback only took 19 % of the time of auditory feedback. A two-way repeated-measures ANOVA shows statistically significant main effects of feedback

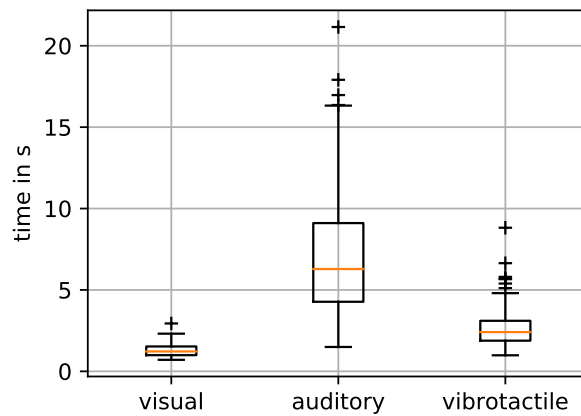


Figure 8.3.: Boxplots of completion times for all conditions with merged data from all participants.

condition ($F_{2, 24} = 50.30, p < 0.001$) and target ($F_{19, 228} = 12.25, p < 0.001$) on task completion time, and an interaction effect of feedback condition and target ($F_{38, 456} = 5.06, p < 0.001$). A Friedman test reveals a significant difference in success rates between conditions ($\chi^2(2) = 25.04, p < 0.001$). Individual success rates (s.r.) are shown in Figures 5 and 6.




Figure 8.4.: Auditory condition, targets: mean movement overhead (m.o.), median trial completion times and success rates (s.r.)

Results for movement overhead are very much as expected for the visual feedback condition (attention funnels), with nearly perfect and immediate localization of the target. Therefore we do not include a detailed report of these results. For comparison, the mean movement overhead for all targets with visual feedback was 8.8 % (SD 3.6 %).




Figure 8.5.: Vibrotactile condition, targets: mean movement overhead (m.o.), median trial completion times and success rates (s.r.)

In Figures 8.4 and 8.5, all targets are projected onto a vertical plane for the initial state of the user looking at the starting cube. The vertical plane’s normal vector is horizontal and points in the user’s frontal direction.

The auditory feedback condition had a mean movement overhead of 67.2 % (SD 14.4 %). Figure 8.4 shows targets in front and back of the user for the auditory condition. Please note that the color scales are different between modalities. For the auditory (g-HRTF + white noise) condition, guidance towards targets near the horizontal plane through the ears works much better than towards targets off the horizontal plane. The dependent variables show a clear trend for targets near the horizontal plane being faster to reach, with a higher success rate and a lower movement overhead. Targets directly above and below the user were tough to find and took users a long time to identify.

The vibrotactile feedback condition had a mean movement overhead of 35.1 % (SD 10.7 %). As Figure 8.5 shows, compared to auditory feedback, targets off the horizontal plane worked much better with vibrotactile feedback. The performance was higher for all the measured variables.

However, targets such as T0, T12, and to a lesser extent, T9, T14, and T18 had higher-than-average movement overheads unexpectedly. We presume that this is due to the chin belt of the initial HapticHead prototype being too inflexible and distributing vibrotactile signals from one actuator along the whole chin (T0, T12), and also due to our guidance algorithm interpolating between the three actuators closest to the target which turned out not to be the best solution for non-uniform actuator distributions. We describe solutions to these issues in the next section.

Comparing the mean movement overheads for front targets on a diagonal guidance path (T0, T7, T12, T13, T14, T15, T18; mean movement overhead: 44.5 %) to those on vertical or horizontal guiding paths (all other targets; mean movement overhead: 30.1 %) suggests that users had more problems locating targets on the diagonals.

Vibrotactile targets on the front had an average completion time of 2.12 s, targets on the back of 2.87 s. This is expected because users first need to turn around to reach targets behind them. Visual targets, in comparison, had an average time of 1.01 s on the front and 1.54 s on the back. The average success rate of front targets in the vibrotactile condition was 98.0 %, whereas it was 94.2 % for back targets.




Figure 8.6.: Selection time by angular distance between the start box and the target center.




Figure 8.7.: Selection time by yaw (horizontal heading) distance between the start box and the target center.

Figure 8.6 shows the selection time by angular distance between the starting orientation of the user and the orientation of each target. For the visual ($R^2 = 0.89$, $p < 0.001$) and vibrotactile ($R^2 = 0.59$, $p < 0.001$) conditions there is a good linear fit between angular distance and selection time. For the auditory condition there is a much weaker relationship ($R^2 = 0.22$, $p < 0.001$). For each target, the visual and vibrotactile conditions outperform the auditory condition.




Figure 8.8.: Selection time by elevation angle between the ground plane and the target center.

Figure 8.7 shows the selection time by yaw (horizontal heading) distance between the starting orientation and each target. Again, there is a good linear fit for the visual ($R^2 = 0.86$, $p < 0.001$) and vibrotactile ($R^2 = 0.72$, $p < 0.001$), but not for the auditory condition ($R^2 = 0.12$, $p = 0.025$).

Figure 8.8 shows the selection time by elevation angle between the starting orientation and each target. There is a weak quadratic fit for the auditory condition ($R^2 = 0.43$, $p < 0.001$) but not for the other conditions. Our observations show that participants had trouble locating targets below or above ear level in the auditory condition. The further targets are off the horizontal plane at ear height, the longer the selection time. This effect is symmetric above and below the plane and roughly has a parabola shape. The visual and vibrotactile conditions do not exhibit this effect.




Figure 8.9.: Average learning effect – time. Merged data, all participants. Curves: Gaussian weighted moving average (width=3, blue=visual, green=vibrotactile, red=auditory).

Figure 8.9 shows the development of completion time across all trials. Note the steep learning curve for vibrotactile feedback that flattens around trial 40 and the auditory learning curve, which initially stays rather constant until around trial 65, when participants started being quite a bit faster.




Figure 8.10.: Average learning effect - success rate. Merged data, all participants. Curves: Gaussian weighted moving average (width=3, blue=visual, green=vibrotactile, red=auditory).

Figure 8.10 shows the success rates over time. Participants had a steep learning curve in the auditory condition, which flattens after trial 40 but still shows a high variance compared to the other conditions. Even though participants learned to be a lot faster after trial 65 in the auditory condition, the success rate did not drop but even increased a bit. Participants needed less than 15 trials in the vibrotactile condition to accommodate themselves with this new form of feedback. After the first few trials, the success rate curve for vibrotactile feedback flattens and stays close to 97 % without much variance or measurable fatigue effects.

8.3.2 Guidance Comparison Experiment – Subjective Results

Subjective results were measured through a post-questionnaire with 5-point Likert scales (Figure 8.11). Participants agreed that the vibrotactile feedback helped find virtual objects, while they disagreed that the auditory feedback was helpful. They agreed that the feedback position around the head was appropriate and that the vibrotactile feedback was comfortable. Participants had mixed opinions about the sound level of the vibrotactile feedback. They agreed that the vibrotactile feedback was unambiguous when looking at the correct target, and most participants could imagine using such feedback regularly.

8.3.3 Guidance Comparison Experiment – Discussion

Our experiment clearly shows that HapticHead vibrotactile feedback guides users towards visible virtual targets around them substantially faster than spatial auditory feedback (directional hearing with a g-HRTF). For vibrotactile feedback, there is a linear relationship between angular distance and selection time. In addition to higher selection times, weak points of g-HRTF auditory feedback include confusion between

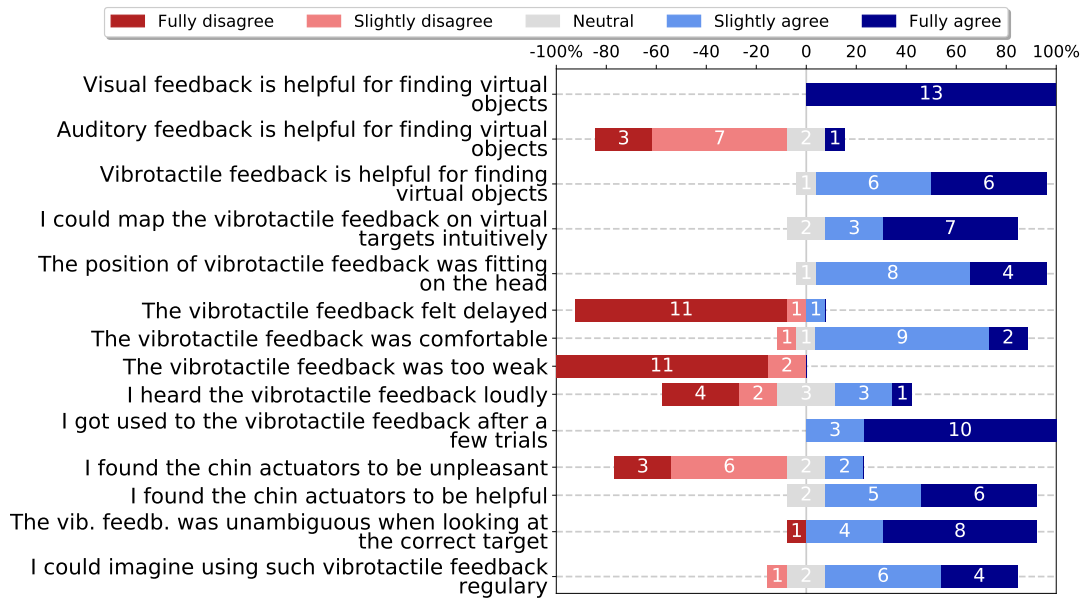


Figure 8.11.: Subjective results of the Guidance Comparison Experiment. Diverging stacked bar chart: scales in percent, and absolute values on bars.

targets directly in front of and behind the user and issues in locating targets above or below the horizontal ear plane, which supports earlier work [141]. This is expected because g-HRTF auditory feedback is known to cause localization difficulties and can be improved by using p-HRTFs [55] or letting users do training on g-HRTF localization with visual feedback on the right target as in work by Klein et al. [112].

We also saw a considerable improvement in the auditory condition’s performance as users adapted to the g-HRTF due to getting visual feedback on the correct target after each trial. Fortunately, vibrotactile feedback does not share these issues and more closely resembles visual feedback’s performance characteristics. Vibrotactile feedback thus suggests itself as an alternative to visual feedback. However, we could not measure the achievable precision due to the arbitrary target size, and we saw opportunities to refine our prototype and guidance algorithm.

The subjective results provide a strong indication that the overlaying hypothesis H2 on the *intuitiveness* of 3D guidance of this thesis (see section 1.1) is accurate as almost all study participants agreed that they could intuitively map the vibrotactile feedback to virtual targets.




Figure 8.12.: Side and front view of modeled actuator positions. Does not fit perfectly due to arbitrary size and asymmetries of the Styrofoam head. The refined guidance algorithm uses triangles between actuators, including a virtual point zero (in green) between the eyes.

8.4 Refining the Algorithm for Precise Guidance

Based on the experiences from the Guidance Comparison Experiment, we built a second main HapticHead prototype in order to improve precision for guidance applications and user comfort (see subsection 4.2.2).

We also refined the guidance algorithm, whose aim is to adapt the actuators' intensity to best guide the user towards a target. The previous version of the algorithm just interpolated between the three actuators closest to the target. This can cause unintuitive behavior, especially when the target is right in front of the user's face.

For the refined guidance algorithm we defined a virtual point zero ("VPZ") exactly between the eyes of the user (marked in green in Figure 8.12). We then tessellated the actuator space in that we placed triangles between each triple of adjacent actuators (including the VPZ) without overlaps as shown in Figure 8.12.

A ray between the center of the head and the virtual target around the head intersects exactly one of the triangles t of the tessellation in a hit point h . Triangle t is defined by its adjacent actuators (v_0, v_1, v_2) . Let point e_i be the intersection of a line through v_i and h and a line through $v_{(i+1) \bmod 3}$ and $v_{(i+2) \bmod 3}$, the other two actuator positions (\bmod is the modulo operation).




Figure 8.13.: Intensity calculation visualization.

The intensities (0 to 1) of the three actuators are calculated as:

$$intensity(v_i) = 1 - \frac{|h-v_i|}{|e_i-v_i|}$$

Special case: If the VPZ is part of the intersected triangle, the intensity of the remaining two actuators is amplified to give the user a sense of direction on the "ring around his face". The user is then drawn a bit more in the indicated direction.

This way of computing the intensities is similar to the 2D linear approach in work by Schneider et al. on "Tactile Animation" [183] whose experiment indicated that users rated a 2D straight-line motion best using a logarithmic approach over a linear one. On the other hand, Seo et al. [192] found that location accuracy in a 1D case between two actuators was higher with a linear approach than with a logarithmic one. Since we focus on localization precision, we chose a linear interpolation approach. Our 3D-sphere linear interpolation approach achieves several objectives: The experienced overall tactile intensity is only weakly dependent on the hit point's position in the triangle, just like in [183]. The algorithm worked well in pilot testing despite non-uniform placements of the actuators. Whether the same algorithm with a logarithmic interpolation would outperform our chosen linear interpolation in a 3D scenario remains an open topic for future work.

We also included a vibrotactile pattern to indicate the angular closeness of the VPZ to the target. In a variant (the Real-World 3D Guidance Experiment), we mapped the pattern to the depth-axis instead. Actuators have a duty cycle of 95 % for a 1 Hz pattern if the user is above a certain distance threshold (either angular or depth). A 95 % duty cycle instead of 100 % was chosen to periodically remind the user that the pattern is still there while the distance is still way off. Once below the threshold, the duty cycle linearly decreases while the frequency increases until the actuators have a 70 % duty cycle for a 50 Hz pattern at zero distance. This is experienced as permanently on with an intensity of 70 %, as the motor's stop time (from full speed to stop) is greater than the signal off-time. Thresholds should be chosen depending on the task. Choosing larger thresholds results in less time needed to complete a task but also less accuracy. In pilot testing, these values and the resulting vibrotactile pattern were found to be effective but remain to be optimized in future work.

8.5 Guidance Precision Experiment

We were interested in how precise users could be with the visual and vibrotactile feedback mechanisms and how our second main HapticHead prototype (see section 4.2.2) and new guidance algorithm would perform compared to the initial prototype.

As in the Guidance Comparison Experiment, users wore an Oculus Rift DK2 in all conditions. The data logging and refresh rate of actuators were the same.

The Guidance Precision Experiment uses a within-subject design with feedback type as the independent variable. There are three levels for feedback condition: *vibrotactile-visible-targets*, *visual-1-pixel*, and *vibrotactile-invisible-targets*. The *vibrotactile-visible-targets* condition is designed to be compared to the vibrotactile condition from the Guidance Comparison Experiment and not to the other conditions in this experiment.

Vibrotactile-visible-targets condition: For comparing the initial prototype and the old guidance algorithm, we ran the vibrotactile feedback condition from the Guidance Comparison Experiment again with the new prototype and guidance algorithm (same target size, 8 repetitions x 20 targets per user). The guidance algorithm used the pattern as explained above with an angular threshold of 40°, which was determined in a pre-experiment.

Visual-1-pixel condition: Since final precision cannot be directly concluded from the Guidance Comparison Experiment due to the arbitrary target size, we decided to use visual feedback as a precision baseline condition. Attention funnels with a target crosshair as in the Guidance Comparison Experiment (shown in Figure 8.2) were used, but here with tiny white 1-pixel targets instead of the large ones (4 repetitions x 20 targets per user).

Vibrotactile-invisible-targets condition: We used invisible targets with the refined prototype and the new guidance algorithm as our vibrotactile precision condition (8 repetitions x 20 targets per user). The guidance algorithm used the vibrotactile pattern described above with an angular threshold of 16°, which was also determined in a pre-experiment. The angular threshold chosen for this condition is smaller than in the *vibrotactile-visible-targets* condition because this condition does not have visual target markers, and we wanted to focus on maximum possible precision within an acceptable timeframe.

The trials were executed as in the Guidance Comparison Experiment. The correct target was indicated by a green sphere ($r = 1.0$ m for the *vibrotactile-visible-targets* condition and $r = 0.2$ m for the others) at the correct position after each trial.

We invited 13 participants (3 female, mean age 22.8, SD 2.6 years). The set of participants for the Guidance Precision Experiment was fully disjoint from the set of participants in the Guidance Comparison Experiment, so no participant had prior experience with HapticHead.

8.5.1 Guidance Precision Experiment – Objective Results

The *vibrotactile-visible-targets* condition had a mean trial completion time of 4.30 s, which is 65 % higher than in the previous study. An increase was expected because participants had to wait and consciously perceive the vibrotactile pattern to confirm a target. However, the success rate increased only marginally to 96.6 %. We suppose that this is because we had one participant (ID 11) who focused on being fast instead of precise. With a mean trial completion time of 2.67 s, this participant was 61 % faster than the average but only reached a success rate of 86.8 %. The other participants had a mean success rate of 97.5 %.

For the other two conditions (*visual-1-pixel* vs. *vibrotactile-invisible-targets*) a two-way repeated-measures ANOVA shows significant main effects for feedback condition ($F_{1,12} = 24.88$, $p < 0.001$) and target ($F_{19,228} = 8.70$, $p < 0.001$) on trial time, and a significant interaction of feedback condition and target ($F_{19,228} = 2.28$, $p < 0.01$) for trial time. Furthermore, feedback condition ($F_{1,12} = 14.69$, $p < 0.01$) and target ($F_{19,228} = 2.08$, $p < 0.01$) have statistically significant main effects on precision. There is also a significant interaction effect ($F_{19,228} = 2.25$, $p < 0.01$) for precision.

The main purpose of this study was to evaluate precision with the refined prototype and algorithm. In the *vibrotactile-invisible-targets* condition, participants reached a mean final deviation from the target of 2.33° (SD 1.80° , 95 % CI [0.59° , 7.13°]) with a mean trial completion time of 8.92 s. The mean final deviation in the *visual-1-pixel* condition was 0.80° (SD 0.43° , 95 % CI [0.24° , 1.80°]) with a mean trial completion time of 3.41 s. In comparison, the vibrotactile condition is less precise but still very close to the target.

8.5.2 Guidance Precision Experiment – Subjective Results

As shown in Figure 8.14, participants agreed that the HapticHead vibrotactile feedback helped find virtual targets, and most of the participants could intuitively map the feedback to the targets. Participants weakly agreed that the vibrotactile feedback was comfortable and disagreed about it being disturbing. They also mostly disagreed that they felt the vibrotactile feedback to be too weak.

8.5.3 Guidance Precision Experiment – Discussion

While the time to find targets increased by more than what we expected in direct comparison to the old prototype and guidance algorithm, we also saw a slight increase in success rates and measured only a tiny error of 2.3° towards targets in the

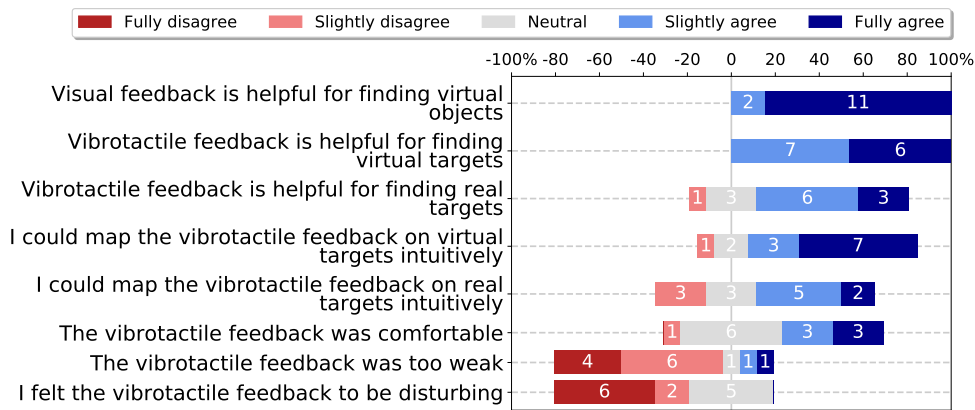


Figure 8.14.: Subjective results of the Guidance Precision and Real-World 3D Guidance Experiments. Diverging stacked bar chart: scales in percent, and absolute values on bars.

vibrotactile-invisible-targets condition. With such a small error, it should be easy to find targets in a real-world scenario, which leads us to our final experiment.

8.6 Real-World 3D Guidance Experiment

Independent of the previous two experiments, we were interested in whether the HapticHead concept can also be used as a sole feedback variant to find real targets around the user, such as keys lost on top of the fridge or guiding visually impaired people into particular directions. *Research question 1:* Can blindfolded users find real objects around them with HapticHead vibrotactile feedback?

We used the OptiTrack tracking system for positions of physical targets and the user's hand position and orientation. This time the user's primary hand was the "center point of attention" because, in pilot testing, it proved to be more intuitive than the center of the head when trying to grab something. This means that the direction of the target object relative to the hand was displayed on the head. *Research question 2:* Is this indirection still intuitive for users? The hand tracker's effective position was manually calibrated for each user to be on the palm of their hand. The guidance algorithm used a vibrotactile pattern as described above with a depth threshold of 1.5 m. This is different from the Guidance Precision Experiment, where an angular threshold was used to trigger the final phase and the vibrotactile pattern. Here, the pattern is used to show target depth instead because when trying to find real targets, it is also essential to give users an impression of the remaining distance to the targets so they do not accidentally grab a knife with too much speed. The threshold distance was determined in a pre-experiment.



Figure 8.15.: 10 items for the Real-World 3D Guidance Experiment. From left to right (height): Three books (0.9 m), a pen (1.8 m), a Lego piece (1.1 m), a ball (1.7 m, 6 cm diameter), a screwdriver (0.8 m), a remote control (1.4 m) and two balls (3 cm and 12 cm diameter, on the ground).

We conducted this experiment right after the Guidance Precision Experiment with the same participants, which allowed us to skip training trials. Ten small items were either hanging from the ceiling on small threads at different heights or were placed on a table or the ground in the lab. The items on the table were three books among ten books at 0.9 m height, so this was a "choose one from many" search task. Of the other items, two were placed on the floor and another two 1.7 and 1.8 m from the ground. The remaining three items were placed at comfortable heights between 0.8 and 1.4 m. This configuration of items was the same for all participants. The items were equipped with OptiTrack markers (Figure 8.15).

The blindfolded participants started a trial in the center of the room within a small marked square, facing in the same direction every time. Their task was to find one randomly chosen item (no repetitions) at a distance of 1.2 to 2.5 m. In each trial, the user was guided to one of the items using HapticHead vibrotactile feedback, and the trial was manually stopped when the participant was sure to have found the right target and said "stop". The experimenter took note of whether the target was the right one. We measured task completion time and whether participants found the right target. Finally, questionnaires gathered subjective feedback on intuitiveness and usefulness.

8.6.1 Real-World 3D Guidance Experiment – Results and Discussion

Six of the seven non-book targets were found by all participants. Two participants (15.4 %) failed to find the last target, a small pen at the height of 1.8 m. They were both relatively short people and remarked that they did not expect something that high up.

The correct book targets were found successfully in 52.6 % of all cases. However, the tracking method (OptiTrack with a marker attached to the user's hand, effective tracking position calibrated below the user's palm) introduced tracking errors for the book targets due to how some participants turned their hand while searching for these targets. Because we logged the hand marker's position and orientation, we were able to manually classify failed trials where the hand marker pointed at the right target as

successful. Excluding errors introduced by the tracking method, two out of three books were found successfully by all participants. The thinnest book was missed 3 out of 13 times. An overall success rate of 96 % was achieved across all targets. On average, it took 42.0 s (SD 45.0 s) to find the correct target. Participants were somewhat cautious and moved slowly because they were not used to being blindfolded and did not want to run into things.

As shown in Figure 8.14, participants found the vibrotactile feedback helpful for finding real targets around them and could intuitively map vibrotactile signals to targets; thus, research question 2 can be answered positively. Both of these measures were less agreed on than for virtual targets, however. We believe this is because participants were used to the angular vibrotactile pattern (the Real-World 3D Guidance Experiment used a depth pattern) and the algorithm reacting to position and orientation changes of their head instead of their hand. Because participants did only ten trials each with these changes (no training trials), they had little time to get used to them.

The Real-World 3D Guidance Experiment shows the HapticHead concept's high potential to be used in a real-world application in conjunction with a suitable tracking system for finding items or simply to orient a user in the right direction.

8.7 Limitations

Using Unity 5.3's included audio system for comparison is a limitation as this system only uses g-HRTFs and is thus not a state-of-the-art audio system. However, p-HRTF systems require a complex per-user calibration, which HapticHead does not. Still, using a p-HRTF system would likely improve auditory results substantially, and a comparison between HapticHead feedback and auditory p-HRTF feedback remains an interesting topic for future work.

We did not investigate the influence of a potential funneling illusion on guidance performance. If users feel one instead of two or three stimuli when two or three actuators are active at a time, this could have an impact on performance and remains to be investigated in future work.

8.8 Conclusion

The experiments' results show that a spherical grid of vibrotactile actuators around the head and our guidance algorithm can effectively guide users towards virtual and real targets in 3D. Vibrotactile feedback turned out to be superior to g-HRTF auditory feedback and almost on par with visual feedback. The main contributions

are the guidance algorithm and the evaluation of speed and precision of a prototype implementing the concept.

The objective results in the Guidance Precision and Real-World 3D Guidance Experiments prove that the overlaying hypothesis H3 on the precision of 3D guidance of this thesis (see section 1.1) is accurate as our objective guidance precision in the Guidance Precision Experiment compared to the other feedback modalities was substantially better than auditory feedback and almost on par with visual feedback. Furthermore, the object finding accuracy in the blindfolded object finding task (Real-World 3D Guidance Experiment) was almost perfect.

Assistance of People with Visual Impairments

Tactile patterns are a means to convey navigation instructions to pedestrians and are especially helpful for people with visual impairments. This chapter presents a concept to provide intuitive and precise micro-navigation instructions through a tactile around-the-head display as another compelling step to reach this thesis's overarching goal to establish the head as a means for tactile communication.

Our system presents four tactile patterns for fundamental navigation instructions in conjunction with continuous directional guidance. We followed an iterative, user-centric approach to design the patterns for the fundamental navigation instructions, combined them with a continuous guidance stimulus, and tested our system with 13 sighted (blindfolded) and two blind participants in an obstacle course, including stairs. We optimized the patterns and validated the final prototype with another five blind participants in a follow-up study. The system steered our participants successfully with a 5.7 cm average absolute deviation from the optimal path. Our guidance is only a little less precise than the usual shoulder wobbling during normal walking and an order of magnitude more precise than previous tactile navigation systems. Our system allows various new use cases of micro-navigation for people with visual impairments, e.g., preventing collisions on a sidewalk or as an anti-veering tool. It also has applications in other areas, such as personnel working in low-vision environments (e.g., firefighters).

This chapter is based on the TOCHI' 2021 paper "Around-the-Head Spatial Tactile System for Supporting Micro Navigation of People with Visual Impairments" [102], written in collaboration with Michael Rohs. Some of the experiments in this chapter were conducted as part of Bachelor and Master theses by Marc Mogalle [139], and Guido Gardlo [53]. Benjamin Simon helped to conduct two of the experiments in this chapter as an assistant researcher.

We used the third main prototype (see subsection 4.2.3) for the first four experiments in this chapter and the fourth and final main prototype (see subsection 4.2.4) for the last experiment (Improved System Validation Experiment).

9.1 Introduction

Blind and visually impaired people (VIPs) face major challenges when navigating in unknown spaces or searching for objects, even in their own homes [127]. Historically, only a few tools for exploring unknown spaces were available such as guide dogs or white canes. Only a low percentage of VIPs even utilize these aids (less than 10 % white cane users [216, 219] and about 2 % guide dog users [61] in the USA). One likely reason for this is the stigma of using a white cane or a guide dog, which shows the user to be visually impaired. There are also issues with adapting to white cane usage, safety concerns [70], and the high cost for a guide dog and the need to care for it [61]. A possible solution to these stigma issues are hidden auditory or tactile guidance systems.

Navigation involves two main components: mobility and orientation as defined by Loomis et al. [126]. Mobility or *micro-navigation* involves sensing the near-field environment and working out a way around static or dynamic obstacles. Orientation or *macro-navigation* involves being oriented (e.g., by detecting landmarks), path-planning on a broader scale, and detecting when a destination has been reached [89]. A navigation system for VIPs typically includes the following three components: *a set of position and orientation sensors* (e.g., magnetometer + GPS, DGPS, or Bluetooth beacons), *a geographic information system* for path-planning (e.g., pre-computed maps), and *a user interface* (e.g., auditory feedback or a tactile belt in conjunction with a white cane).

Examples for the first two components can be found in recent work on detecting open areas in front of a walking person [179], a social distancing assistant for VIPs with real-time semantic segmentation on RGB-D video [131]. One recent work even attempts to use technology intended for autonomous cars to sense the environment around VIPs [130]. Some systems merge the first two components into a single entity, such as simultaneous localization and mapping approaches with LiDARs, or modern Augmented Reality self-localization frameworks such as Google ARCore [58] and Apple ARKit [6].

Occupying the sense of hearing for a navigation system's user interface is a safety concern as walking on the street requires unobstructed hearing, even for sighted people. A solution for this are tactile user interfaces that do not occupy the sense of hearing. Most existing tactile user interfaces for VIP navigation focus on the macro-navigation aspect. These systems aim to roughly guide VIPs through cities by providing turn-by-turn directions towards a destination, using either GPS as a position and orientation component or a remote operator (e.g., [10, 11, 150, 180, 215]). Due to the relatively low accuracy of their position and orientation components, these systems do not attempt to provide micro-navigation instructions but instead require the user to apply other micro-navigation tools (e.g., guide dogs or white canes). There

is one work by Flores et al. [52], who presented an 8-actuator tactile belt system for micro-navigation of VIPs using a *dynamic stimulus vibrotactile pattern*. However, an average deviation from the optimal path of 49 cm prevents usage for a variety of micro-navigation cases requiring high precision, such as walking on a sidewalk while safely moving around obstacles and other people or walking to the counter in a crowded bar.

The human head presents itself as a mostly spherical surface for tactile feedback, which increases the design space of possible tactile patterns and is also intuitively the center of attention for humans, which has various advantages for tactile feedback on the head. For example, a strong tactile stimulus on the back of the head can intuitively be recognized as “something is behind me, might be dangerous” due to natural mapping and instincts [101, 159]. We define *intuitive* tactile patterns as mostly “self-explanatory” and requiring minimal training, like a single presentation of the available tactile patterns, to feel comfortable with the system and to operate it with minimal interpretation errors.

In the previous chapter, we showed that our HapticHead system can be used in 3D guidance and localization scenarios in virtual (VR) and augmented reality (AR) with relatively high precision and low task completion times. We also investigated characteristics of tactile patterns on the head in chapter 5. The results gave us insights into important aspects of intuitive tactile patterns.

These prior works initially inspired us to think about other exciting use cases for precise guidance, such as micro-navigating VIPs. Our initial research question was: “Can we extend the HapticHead system to provide VIPs with precise micro-navigation instructions?” This first notion naturally led to more research questions: “What kind of instructions should we provide to VIPs and how should we represent them?” and finally: “How accurately can our system guide VIPs?”

Our primary goal for this work was to achieve a higher micro-navigation precision for VIPs than the prior state-of-the-art system [52]. Due to our experiences with the blindfolded 3D guidance experiment in chapter 8, and HapticHead’s ability to steer users to an invisible target on a 3D sphere around them at a high median precision of 2.3° to the target, we assumed the HapticHead prototype to be able to steer VIPs alongside an optimal path at a higher precision than reported in [52].

9.1.1 Approach

After reviewing related work (section 9.2), we started with the elicitation of fundamental VIP navigation instructions through an informal interview (section 9.3.1). We then proceeded to design tactile patterns for these navigation instructions through a user-centric design approach, which included two user studies (sections 9.3.2 and

9.3.3). This approach produced four intuitive tactile patterns, which we combined the continuous guidance stimulus from the prior chapter (section 8.4) and an additional ATTENTION pattern to form a highly precise micro-navigation guidance system (section 9.4).

To test the precision of our guidance system, we invited a total of 13 participants with normal vision who were blindfolded and two blind participants for the *Obstacle Course Experiment* (section 9.5). We found that our system was already significantly more precise than related work [52] and that we should further reduce the presentation time of our static patterns as participants sometimes went off-track because the long presentation time of our static patterns overshadowed the directional continuous guidance stimulus.

Consequently, we shortened and refined our static patterns (section 9.6) and finally validated the improved micro-navigation system with another five VIPs in the *Improved System Validation Experiment* (section 9.7).

9.1.2 Contributions

In advancing the fields of tactile micro-navigation and assistive technologies for VIPs, we make the following contributions: (1) a set of fundamental navigation instructions and associated *intuitive* vibrotactile patterns on the head, optimized in three consecutive studies, and (2) a system using these optimized tactile patterns for fundamental navigation instructions combined with a continuous tactile guidance stimulus to provide precise micro-navigation instructions. The system supports navigation around obstacles and on stairs. The precision of our system is substantially better (5.7 cm mean deviation from the optimal path) than related work (49 cm mean deviation from optimal path [52]).

Compared to prior work, our system's substantially higher precision opens up a whole new set of use cases, like steering VIPs around obstacles on a sidewalk or steering VIPs in a mall without running into others. While the system was developed and optimized with an emphasis on high precision for micro-navigation, it can also be used in (combined) macro-navigation use cases and thus presents itself as a complete output solution to micro and macro-navigation for VIPs, given a suitable and precise tracking and obstacle detection system. Even though our system is primarily intended for VIPs, we imagine it to be also applicable in other scenarios where precise tactile guidance is necessary. For example, it could be used in guidance scenarios for firefighters or other personnel operating in low vision environments, jet or drone pilots, or VR/AR scenarios where the visual and auditory channels should not be overtaxed to show navigation instructions or guidance to specific targets.

9.2 Related Work

The fundamental related work on tactile perception around the head (subsection 2.3.8), tactile displays (section 2.6), and indicating direction via tactile feedback (subsection 2.6.1) also applies to this chapter. We further use the vibrotactile pattern terminology as defined in section 2.5.

Specific related work concerning assistive systems for VIPs is presented below.

9.2.1 Assistive Technologies for VIPs

Csapó et al. [31] summarize developments of assistive technologies for VIPs based on audio and tactile feedback.

Recently, there have been attempts to allow VIPs to experience virtual reality and allow easy white cane training through enhanced white canes, which are tracked, actively braked by virtual obstacles, and that even provide vibrotactile feedback about ground properties [196, 229].

Scene sonification is an exciting research direction, which allows VIPs to perceive a scene via auditory cues [75]. Hu et al. [75] investigated three different kinds of scene sonification (depth image sonification, obstacle sonification, and path sonification) in a comparative study and found that preference for specific sonification approaches was highly individual and that sonification of high-level scene information (e.g., the direction of a pathway) is generally easier to learn than sonification of low-level scene information (e.g., raw depth images).

In terms of obstacle detection, Poggi et al. [161] proposed a mobile system that detects objects through deep learning to give speech-based warnings of obstacles to VIPs. Using a tactile 3×3 grid on the abdomen, Van Erp et al. [50] presented a system to indicate obstacle information around the user, including direction (3 levels), distance (4 levels), height (3 levels), and type (4 levels). They found that users had difficulties distinguishing the large number of tactile patterns needed to identify the obstacle information with detection rates between 42 to 76 % for direction and height and 12.8 to 47 % for object distance after training. These results highlight the importance of well-designed patterns that ensure correct identification, especially in critical situations (e.g., crossing a street). Van Erp et al. went for a multimodal pattern presentation approach (tactile+auditory) in their follow-up experiments [50].

9.2.2 VIP Guidance via Tactile Feedback

Besides *verbal* guidance systems (e.g., [10, 11, 218]), there are two main kinds of assistive *tactile* technologies in navigation scenarios:

1. Vision substitution systems that map a depth-image from an RGB-D camera to a high-density tactile grid placed on the tongue [8], back [9, 25], forehead [83], or abdomen [189].
2. Tactile feedback systems that directly map orientation information onto a low-density actuator arrangement (usually a ring or grid configuration) and mostly placed on the wrist [152, 180], head [34, 150], neck [132, 178], feet [215], or waist [30, 52, 71, 206].

Not all of the systems mentioned above are targeted explicitly at VIPs but could conceivably be used in assistive scenarios as well. Specifically for VIPs, Scheggi et al. [180] performed a macro-navigation task using two vibrotactile wristbands and a remote operator for providing simple left or right navigation instructions to the VIP. Micro-navigation was still performed by the VIP using a white cane. Velazquez et al. [215] presented vibrotactile shoes to indicate four directions for macro-navigation (forward, backward, left, and right) and achieved a detection accuracy of > 89 % in their pattern discrimination study. While they also performed a macro-navigation study in a city, participants performed micro-navigation using their white cane.

Finally, Flores et al. [52] tested a vibrotactile belt with eight actuators in a micro-navigation scenario comparable to our work but without obstacles and stairs. They further compared vibrotactile to auditory guidance and found an average absolute distance to the optimal path of 49 cm using the tactile belt, compared to 61 cm using auditory guidance. Their tracking system has an accuracy of < 10 cm (ours < 1 cm). These deviations can be directly compared to the average absolute distance to our final system's optimal path, which is an order of magnitude (8.6 times) smaller at 5.7 cm.

The HapticHead hardware system used in this work (see chapter 4) is most closely related to the waist-belts mentioned above, as it basically takes three actuator ring-configurations and puts them on the head in different orientations to provide 3D guidance instead of the 2D guidance of other systems. While 3D guidance is not always necessary, the increased actuator count compared to the usual waist belts and their spatial distribution around the head allows for more detailed vibrotactile patterns that feel more “present” to the user and are easier to interpret due to the given spatial relations and natural mappings [101, 159]. Therefore, we expect less of a chance to misinterpret a tactile pattern in a stressful or dangerous situation, such as when approaching a down-leading staircase at a train station.

9.3 Initial Design of Tactile Patterns for Special Navigation Instructions

This section introduces our steps to elicit four fundamental navigation instructions through an informal interview with VIP navigation experts. Two subsequent user studies allow us to select suitable static tactile patterns from a large set and finally to optimize and validate the selected patterns for intuitiveness and practical use.

9.3.1 Informal Interview at an Educational Center for VIPs

Early on in our research, we made an appointment at an educational center for VIPs in Hanover, Germany, which teaches blind children and teenagers usage of the white cane and additional skills and navigation tools. We were greeted by a teacher, who is herself fully blind, and a caretaker. Because of their jobs in the educational center for VIPs and personal experience, both the teacher and the caretaker can be considered experts in VIP navigation and tools for supporting blind navigation. We presented them an early version of our continuous guidance stimulus (see section 8.4), implemented on a Google project Tango device (Lenovo Phab 2 Pro) with the third main HapticHead prototype (see subsection 4.2.3) as user interface. This version of the guidance system had the advantage of being mobile, but the tracking system was very slow (5 fps), and it did not yet include any static patterns (e.g., stairs down). Despite the advantage of offering fully mobile guidance, this prototype system was not published, as we deemed it was not safe enough due to the somewhat inaccurate and slow obstacle tracking provided by the Lenovo Phab 2 Pro.

We demonstrated the system to the blind teacher by navigating her successfully around a corridor and through doors while the caretaker supervised the experiment. We then conducted an informal interview with both the teacher and the caretaker and asked how they would improve the system and what other information VIPs would need to navigate, potentially without a white cane. Through our conversation, we came up with the four navigational instructions that were deemed necessary by the teacher and the caretaker. Apart from the general direction, the following four navigation instructions: GO / START, STOP, UP, and DOWN. The START instruction indicates that the user should start following navigation instructions. The general directional guidance may then be used to navigate around obstacles on a safe path. Elevation changes (e.g., single sidewalk steps or stairs) can be indicated by the UP and DOWN instructions, and the STOP instruction may be used in potentially dangerous situations, if the user diverges too far from the safe path, or has to wait for other reasons (e.g., at a streetlight). Thus, these four identified navigation instructions can be used in all kinds of micro-navigation scenarios alongside general directional guidance.

9.3.2 Elicitation Study: Designing the Initial Tactile Patterns

In a pre-experiment, two of the authors iteratively designed a total of 16 spatial vibrotactile patterns for the head. The design rationale for these tactile patterns was defined by our experiences in prior work (see chapter 5), where we conducted experiments exploring tactile patterns for general use cases on the head. We found that (a) stimulus location was significantly more straightforward to identify than pattern rhythm or intensity and (b) static patterns with a static stimulus location were easier to identify yet more uncomfortable than static patterns with dynamic stimulus locations. Thus, we chose strong static patterns with static stimulus locations for the navigational instruction STOP as we wanted these patterns to be clearly recognizable and feel very strong and uncomfortable. For the other navigational instructions, we chose static patterns with dynamic stimulus locations as we wanted these to feel comfortable yet intuitive and recognizable through their movement over extended areas of the head and natural mapping (e.g., a movement from the chin over the sides of the head to the top for the UP navigational instruction) [159].

We created many more patterns than needed to represent the required instructions and then discarded those that we could not identify correctly. Through this pre-experiment, we singled out four patterns for each of the instructions above (START, STOP, UP, and DOWN), which we assumed to intuitively represent those instructions (4 instructions \times 4 patterns per instruction = 16 patterns). To find the best four possible patterns, we conducted a short study with 11 participants (9 male, 2 female, mean age 28.2, standard deviation 10.2 years). Even though all participants had normal vision, this is no confound as it did not help them solve the task.

For this short user study on the patterns' intuitive meaning, we implemented a simple Android application running on a Lenovo Phab 2 Pro device and used the third main HapticHead prototype as user interface (see subsection 4.2.3). This application can play predefined spatial tactile patterns around the head by sending actuator control commands to the Raspberry Pi 3. For this study, we implemented a mode to play a random pattern (no repetition allowed) and then asked the participant for that pattern's meaning. The participant could answer by pressing one of four buttons (START, STOP, UP, and DOWN). The choice was forced as there was no neutral button. During the study, participants did not receive any feedback about whether their intuition for the pattern meaning was correct. Each of the 16 possible patterns was repeated 10 times. The patterns were randomly shuffled (seeded by participant id) to counterbalance possible learning effects. There was no training phase. Participants were unaware of which specific patterns were in the set of possible patterns and also unaware of how many different patterns there were in total. The whole study, including filling out questionnaires, took around 45 minutes per participant.

As a result of this first user study, every pattern, except for one, scored more than 50 % accuracy, so our initial iterative design approach for finding intuitive patterns did




Figure 9.1.: Visualization of patterns. Blue: start of pattern; green: end of pattern. The UP pattern uses two actuators as end points, because the top of the head is less sensitive to tactile stimuli (see chapter 3 and [144]) and needs a stronger tactile impulse to emphasize the direction at the end of the pattern.

yield acceptable results. The best patterns scored accuracy scores of 69 % for START, 84 % for STOP, 94 % for UP, and 83 % for DOWN. After the study, we optimized the winning patterns by adding a 100 ms directional stimulus after the initial patterns (for all but STOP), based on user feedback and our experiences. The optimized patterns are defined below (see also Fig. 9.1):

- START – starting at the back of the head, simultaneously moving across both sides and ending at the forehead. Signal 800 ms + pause 300 ms + signal 800 ms + pause 100 ms + direction signal 100 ms.
- STOP – all actuators at the back of the head at the same time. Signal 100 ms + pause 150 ms + signal 100 ms + pause 150 ms + signal 100 ms.
- UP – starting at the chin, simultaneously moving up on both sides and ending at the top of the head. Signal 800 ms + pause 300 ms + signal 800 ms + pause 100 ms + direction signal 100 ms.
- DOWN – like UP but starting at the top of the head and ending at the chin.

We hypothesized that if people receive minimal training on these four patterns, a recognition rate of close to 100 % can be achieved. This hypothesis was tested in a second short user study.

9.3.3 Pattern Recognition Performance Study

For the pattern recognition performance follow-up study, we invited 10 participants (9 male, 1 female, mean age 24.3, standard deviation 2.8 years) with normal or corrected to normal vision, which again did not help them solve the task. We used the same

hardware and software setup as in the previous study. Compared to the prior study, we introduced a 5 minute training phase with visual feedback on correctness before the study's test part began, which was identical to the first study (no feedback on correctness). We logged recognition rates for both the training and the test part.

The results showed a nearly perfect recognition of the optimized patterns. Overall, 391 of 400 trials or **97.8 %** in the **training part** and 397 of 400 trials or **99.3 %** in the **test part** were successful. One participant had some difficulties understanding the experimental task due to language barriers and scored a little lower than the other participants. This participant failed 8 trials (20 %) at the beginning of the training part before we intervened and explained the task again, and another two (5 %) in the study part.

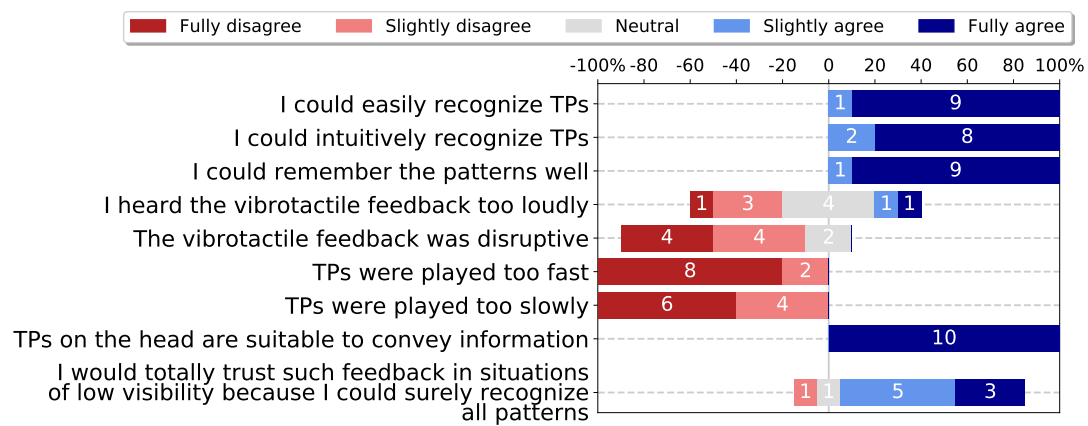


Figure 9.2.: Results of the post-questionnaire of the Pattern Recognition Performance Study about the suitability of the chosen four tactile patterns (TPs) for conveying the four instructions (N = 10).

Fig. 9.2 shows the subjective results of the post-questionnaire. Participants strongly agreed that they could easily and intuitively recognize the four tactile patterns and had no trouble memorizing them. Furthermore, most of them did not feel the feedback to be disruptive and would even trust their ability to recognize the meanings of those patterns in low-visibility situations correctly.

The recognition results of the pattern recognition performance study are close to perfect, just like we expected. We did not expect the results in the initial training part to be that good, however. These results lead to the conclusion that these four new, slightly modified patterns were more intuitive than the winning patterns in the prior study. Only a single participant (P8) failed more than one trial due to language barriers. Together with the encouraging subjective results – participants stated that they could easily recognize these patterns – we believe these patterns may be used in mission-critical applications, such as VIP guidance near traffic lights or stairs after a brief familiarization step.

These results provide the basis for the implementation of our micro-navigation system. We use the developed patterns in conjunction with a continuous tactile navigation pattern for VIP guidance in the implemented system. This combination poses the particular challenge of how to interleave the continuous directional tactile guidance signal and the specific command patterns.

9.4 Implementing the Micro-Navigation System

When we first started experimenting with micro-navigation, we noticed that playing a STOP pattern every time the user diverges a little too far from the optimal path is disturbing, as the pattern is very strong and meant to convey immediate danger. We felt the need for an instruction to focus back on the navigation task once the user diverges too much, but non-critically, from the optimal path. A friendly reminder to focus is needed, as users in our preliminary tests were often distracted by random thoughts, noises, or talking while navigating.

In addition to the four patterns that resulted from the previous studies, we introduced an additional ATTENTION pattern as a reminder for the user to focus. It is identical to the STOP pattern but without repetitions. The ATTENTION pattern is used when a participant deviates too far from the route or gets close to an obstacle (other than the stairs). On the other hand, the STOP pattern is used when the participant deviates too much from the path and needs to realign while stopping. Moreover, the STOP pattern is used before playing the UP or DOWN patterns, so the participant stops and recognizes the pattern before climbing stairs.

We combined our patterns with a continuous tactile navigation stimulus that indicates the next waypoint's direction while no other pattern is active. We used the same continuous guidance stimulus as in the prior chapter (see section 8.4). The following patterns were implemented in the first version of our micro-navigation system in addition to the patterns validated in the Pattern Recognition Performance Study, also visualized in Fig. 9.1:

- ATTENTION – same as the STOP pattern in the Pattern Recognition Performance Study but without repetitions.
- CONTINUOUS-GUIDANCE – while no other pattern is active, actuate the three actuators closest to the next waypoint to guide towards the next waypoint. Details are given in the next subsection.

Because related work has shown that the number of simultaneously active actuators should be low to avoid confusion [43], we decided only to present one pattern at a time, while prioritizing some patterns over others in case the system needed to playback multiple patterns at the same time (e.g., in front of the stairs if the participant was

moving fast). The STOP pattern had the highest priority and could thus overwrite (and push back) any other pattern, followed by the UP/DOWN and the ATTENTION patterns. The CONTINUOUS-GUIDANCE pattern had the lowest priority and was only active while no other pattern was being played back at the time.

9.5 Obstacle Course Experiment

To test the designed micro-navigation system in a realistic environment, we designed an obstacle course and invited 15 participants (11 male, 4 female, mean age 28.3, standard deviation 9.7 years; range 19-57 years) for the experiment. Of those 15 participants, 2 had a severe limitation of sight with 0 % and 2 % vision, respectively. To ensure that these two participants did not use any remaining sight, they also wore blindfolds. The two VIP participants were a couple who were in the experiment room together, doing the experiment one at a time. The other participants did the experiment alone. This is a confound in the experiment as the two VIP participants kept talking to each other and the experimenters and thus were less concentrated on the task at hand. We felt that it would have been inappropriate to separate them for the experiment as they would likely have felt uneasy about being alone. Another confound is that the two VIP participants were older (45 and 57 years) than the other participants.

In the sections below, the participants with normal vision are called “sighted participants” while the VIP participants are called “blind participants.” All of the participants wore blindfolds.

9.5.1 Obstacle Course Experiment – Design and Implementation

To test our navigation approach, an obstacle course (see Fig. 9.3) was created with cardboard boxes and a specially constructed wooden platform. The platform consisted of plywood and was planned and built in accordance with the German DIN 18065 [40]. The norm limits the ascent to 14-20 cm, the tread to 23-37 cm, and the platform’s width to at least 80 cm. These limits are mandatory for essential stairs in residential buildings. The platform consists of a step on each side (ascent: 17 cm; tread: 29 cm) and an 80 cm wide level on top (Fig. 9.4).

Specifically for this study, we also implemented a simple path-finding system where paths through the obstacle course were hard-coded as a set of waypoints. The next waypoint in a path was selected once the participant was within a 10 cm range of the current waypoint. The participant’s orientation was determined by an OptiTrack [149] head tracker. Simultaneously, the body position was determined by the average position of two shoulder-trackers instead of the head position as the head may move

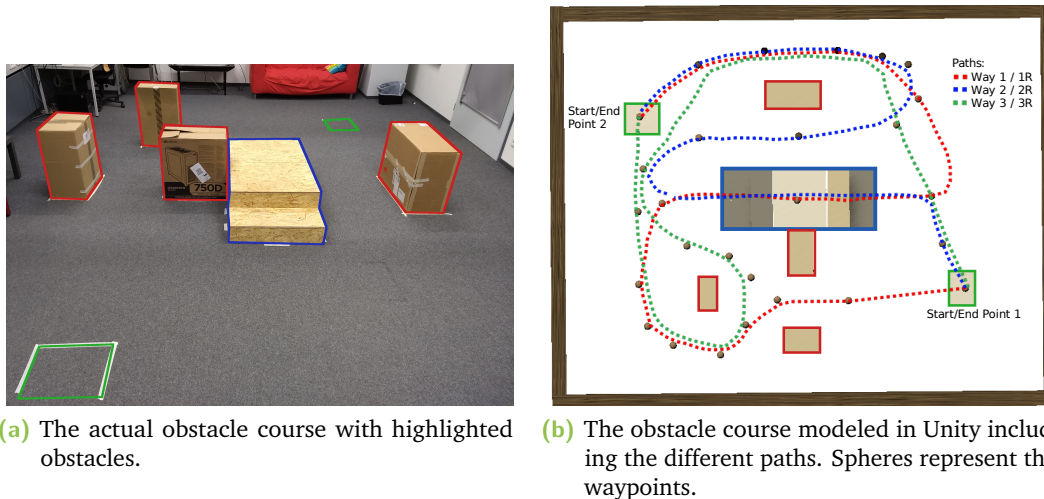


Figure 9.3.: Photo and model of obstacle course. Color coding: green – start/end zones; red – obstacle; blue – stairs.




Figure 9.4.: Model of the staircase with measurements. Top view in center of Fig. 9.3b. The stairs are 80 cm wide.

around too much, which may cause the signal to change rapidly, causing confusion. OptiTrack reported an accuracy of better than 1 mm.

We used the third main HapticHead prototype as user interface in this study (see subsection 4.2.3).

9.5.2 Obstacle Course Experiment – Procedure

After a short greeting, the participants were asked to fill out an introductory questionnaire, a mandatory informed consent form, and optionally a photographic release form. Then they were introduced to the HapticHead prototype. After that, the HapticHead was adjusted to fit over the balaclava that participants wore for hygienic purposes. A vest with additional shoulder trackers was fitted, and two more trackers were attached to the participant’s shoes with painter’s tape. Lastly, the participants were blindfolded.

The blindfolded participant was then guided into the experiment room with the obstacle course. At the first starting position (see Fig. 9.3b), the five static patterns were played once and explained to the participants. Following the pattern introduction, the



Figure 9.5.: Participant climbing the stairs while being accompanied by one of the experimenters for safety reasons.

OptiTrack tracking markers for the HapticHead were attached and roughly calibrated into the correct orientation by asking the participant to turn around until they felt a stimulus precisely between the eyes.

After testing whether the participants remember the static patterns correctly, the patterns were either explained again, or the experiment was started. In 24 runs, the participants were guided on all six routes (three paths, in both directions, Fig. 9.3b) four times in random order (participant number as seed). If the starting point of the following path was different from the endpoint of the last path, the participants were guided on a randomly chosen route through the obstacle course by the experimenter. At least one attendant was walking next to the participants at all times to prevent any injuries from possible falls (see Fig. 9.5).

In case the tracking system failed (lost tracking for more than 1 s), we repeated a trial. This happened a total of 12 times. A post-study questionnaire gathered further subjective data on the performance of the prototype. The whole study, including setup and questionnaires, took about one hour per participant.

9.5.3 Obstacle Course Experiment – Objective Results




Figure 9.6.: Number of times participants needed assistance when on the stairs and number of obstacles hit (including participants just brushing the obstacles with their clothes). Participants 12 and 13 are VIPs, the others are sighted, but all of them were blindfolded. There were 24 runs per participant, and there could be multiple hits or needs for assistance in a single run.




Figure 9.7.: Mean time needed per run (across all participants). The means are smoothed by a weighted moving average with a window size of 3. The weights are 0.25, 0.50, and 0.25.

For all runs, we counted every time a participant touched an obstacle or needed assistance. We were conservative in counting obstacle hits. Even just brushing the side of an obstacle with their clothing was counted as an obstacle hit. Being too close to the staircase’s borders (shoe rim on or over the side-ledge of the staircase) was counted as needing assistance on the stairs as we did not want to risk participants falling.

The results are shown in Fig. 9.6, where they are graphically presented for sighted and blind participants in addition to the overall data. In 33.6 % of all runs, an obstacle was hit, or assistance was needed. Of those, 17.4 % were obstacle hits, and 16.2 % were assistance needed while using the stairs.




Figure 9.8.: Boxplots showing the median absolute deviations of each participant from the optimal path for all trials. Participants 12 and 13 are VIPs.

Our blind participants hit an obstacle on 27.7% of their runs while needing assistance on the stairs or re-positioning in 36.2 % of all runs. In contrast, the sighted participants hit an obstacle on 15.8 % of all runs on average while needing assistance with the stairs in 13.2 % of the cases, totaling 29 %. Of course, these results cannot be generalized to the blind population as we only had two blind participants in our experiment.

An additional measurement was the duration of every run. Fig. 9.7 shows that the average duration in each case reduces with the number of runs, indicating a learning effect. Across all participants, the average run duration starts at 103 s for the first run and reduces to 61 s for the last run, which amounts to a decrease of 41 %.

Furthermore, the deviation from the predefined optimal path was measured overall and for left and right deviations (see Figures 9.8 and 9.9). We found an overall average absolute deviation to the optimal path of 9.3 cm (SD=2.0 cm) across all participants. The group of blindfolded participants with normal vision scored an average deviation to the optimal path of 9.0 cm (SD=1.8 cm), while the group of the two blind participants scored a marginally higher average deviation to the optimal path of 11.1 cm (SD=2.5 cm).

9.5.4 Obstacle Course Experiment – Subjective Results

During the study, we asked the participants whether they had a sense of where they were in the obstacle course and how many paths they thought there were during trials while they were not at risk of bumping into an obstacle. We were interested in whether participants were able to create a mental map of the obstacle course. While they figured out that there was only one staircase with two steps after around half the




Figure 9.9.: The median (left=red, right=blue) deviation per trial for all participants (N=15). The numbers above the boxplots show the percentage of time a participant spent on the left of the optimal path and the average deviation to the left. Ideally, these numbers should be 50% and 0 cm; deviations from them indicate a systematic error. Participants 12 and 13 are VIPs.

study was done, not even the VIP participants could tell how many paths or obstacles there were.

The evaluation of the post-study questionnaire is shown in Fig. 9.10. All 15 participants were pleased with the navigation method they had experienced. Furthermore, 12 of all 15 felt safe while being navigated through the obstacle course. Fourteen of the fifteen participants stated that the way of navigating was intuitive. Even stair recognition was deemed easy by 13 participants, while 10 participants stated that climbing stairs with the given tactile support was easier than climbing stairs without support. Lastly, all participants judged the vibration feedback placement as appropriate and the patterns themselves as not disruptive. Overall, 93 % of the participants stated that they would use this tactile navigation method in everyday life if integrated into a beanie.

Sighted participants All 13 of our sighted participants liked the way of navigating, with 10 of those feeling safe while doing so. For 12 participants, the navigation was intuitive and the patterns easy to remember, while all 13 stated that they did not confuse the meaning of the five static patterns. 11 of 13 participants stated that it was easy to recognize the stairs, with only eight claiming ease of use for stair climbing. Nevertheless, 12 of 13 participants (92.3%) expressed an interest in using a system like HapticHead if integrated into a beanie.

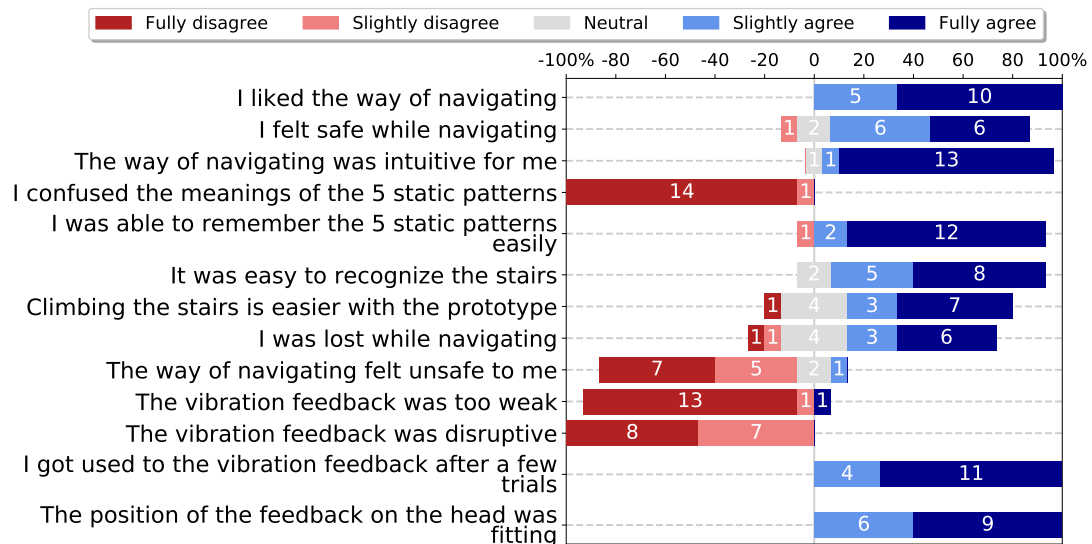


Figure 9.10.: Subjective results of the post-study questionnaire.

Blind participants For the blind participants taking part in the study, the ratings were predominantly positive. The way of navigating was experienced as pleasing, the navigation was perceived as safe, and the usage seemed intuitive. The stair recognition and climbing were judged as easy. Additionally, the blind participants got used to the vibration feedback after a few trials, and the vibration motors did not feel disruptive to them. Both blind participants stated that they would use a system like HapticHead integrated into a beanie and were quite excited about the prospect.

9.5.5 Obstacle Course Experiment – Discussion

The above results suggest that our system worked rather well as a guidance method in unknown environments. While the users felt safe and were pleased with navigating, the vibration feedback was not too strong or disruptive. It was rated as unobtrusive, which surprised us, especially since both our blind participants specifically stated that the vibration noise on the head did not bother them at all. We expected them to complain about the noise, as their sense of hearing is usually one of their primary senses to navigate an unknown environment. We were also surprised at how safe our users rated the system, given the number of times obstacles were hit. This is likely because we were rather inclusive on what to count as an obstacle hit. Furthermore, the obstacles were quite close together (smallest gap: 71 cm), so a slight deviation from the path already led to an obstacle hit.

We did not have any incident of a participant falling, and the most dangerous situations arose when participants got too fast and hit the lower part of the stairs with their feet before reacting to the STOP pattern. This aspect could be improved by implementing

the system to consider the current movement speed and play the vibration patterns earlier if users get too fast.

The graphs in Fig. 9.7 imply a learning effect: On average, the run's duration decreases with practice. This was not a result of learning a path, as almost all participants stated that they felt completely lost and did not know where they were or even how many paths there were.

While blind participant 1 (P12) walked unhurried, blind participant 2 (P13) nearly jogged through the obstacle course. Except for the unhurried pace, blind participant 1 was not as careful while walking and hit seven obstacles. Most obstacle hits of blind participant 1 occurred because of talking while moving. We suppose that the attention faded from the navigation task. Blind participant 2 only hit six obstacles but needed regular assistance on the stairs because of the walking curve's large radius. With the increased walking speed in the small obstacle course, sharp turns would have been necessary to stay on the optimal path but were rarely done. Therefore, the participant often walked into different directions and diverged from the path, not responding fast enough to the tactile guidance signal changes. The measurements clearly show this increased deviation: See P13 in Fig. 9.8.

The increased curve radius is an effect that was observable for most participants after getting used to the prototype, feeling safe, and starting to walk faster. With the reduced run-duration, the effect of increased walking speed dominates, meaning that users with more training will take bigger steps and widen their curves. Consequently, problems occurred with our prototype's software at specific bottlenecks of the path, like the stairs' starting point. The participants often stood on the very edge of the first step on the stairs because their wide curve from the last navigation point led them outwards. While climbing the stairs, sometimes no directional feedback was offered due to the relatively long durations of the static patterns and to keep the number of currently active signals to one, both for static or dynamic patterns. On top of the platform, the continuous guidance pattern was activated again, and the participants corrected their alignment back into the middle of the path. We improve this aspect of the prototype in the Improved System Validation Experiment by reducing the static patterns' playback time to give continuous guidance feedback in between the other patterns while climbing the stairs.

Fig. 9.9 shows the median left and right deviations from the optimal path for all participants. Ideally, the participants should diverge as much to the left as to the right while wobbling around the optimal path with each step. However, we did expect a systematic error here consisting of (a) tracker placement errors on the HapticHead, (b) HapticHead placement errors on the participant's forehead, and (c) the individual participant's feeling of where the frontal vibration direction actually is as this does not necessarily have to be exactly between the eyes on the forehead. Since we are aware of and can quantify the systematic error in this experiment, we can consider it in the

conclusions drawn. Almost all of the participants experienced a systematic error to the left except for P1. While we do not have an explanation for this phenomenon, the average absolute systematic error for this experiment is still rather low at an average of 4.4 cm (SD=3.2 cm) across all participants. This systematic error could be reduced by performing a more sophisticated per-participant calibration in future experiments. On the other hand, reducing this systematic error would also mean that the experiment would be less realistic as users in the real world would likely not want to perform a calibration step each time they put on their navigation aid.

Fig. 9.8 shows the median absolute deviations of each participant from the optimal path for all trials. Our participants tended to deviate an average of 9.3 cm from the optimal path. In light of the systematic directional error discussed above, which also influences the absolute deviation, the absolute deviation from the optimal path still seems relatively low. It may be further lowered by improving the guidance algorithm, e.g., by taking into account the current participant speed rather than directing to the next waypoint once within 10 cm of the current waypoint. Compared to related work by Flores et al. [52], our average deviation from the optimal path was already significantly lower in this experiment (9.3 cm vs. 49 cm).

Our OptiTrack tracking system turned out to be not as resilient as expected. The foot tracking precision was significantly reduced in the corners of the course, primarily due to limitations of the OptiTrack [149] camera perspective and masking by obstacles. This limited the obstacle detection in front of the feet because the feet' virtual representation may point in a different direction shortly before an obstacle was hit. Furthermore, the trackers were sometimes confused with each other even though we made sure they were configured correctly.

9.6 Static Patterns Refinement Study

This study is concerned with evaluating further improvements of the static patterns. While the static patterns were generated in an iterative process, we still felt the need to optimize these further as the total playtime, especially for the START, UP, and DOWN patterns was too long and sometimes led to the system playing just static patterns (e.g., on the stairs) without letting the participant know the correct direction by playing the dynamic guidance pattern. Furthermore, we saw room for improving the guidance algorithm in terms of collision prevention performance (e.g., steer users depending on current speed, steer towards the inside of a curve).

In order to optimize the static patterns, we took the following measures:

- We decided to omit the START pattern altogether as it is pretty obvious the system is guiding the user while the dynamic guidance pattern or any of the static patterns is currently on.

- We decided that the two repetitions of the UP and DOWN patterns took too much time, and thus we only play the pattern once instead of twice now.
- As our prior work and related work shows that the localization precision and perceived stimulation strength is worse on the top of the head compared to the center of the chin (see chapter 3 and [144]), we included two more actuators to represent the “top of the head” for both, UP and DOWN patterns and let that part of the pattern run for 20 % longer for the UP pattern.
- We further worked on the UP and DOWN patterns so that the feeling of something moving up or down the head feels smoother than before.
- We worked on distinguishing the STOP and ATTENTION patterns more clearly by adding another seven full intensity actuators to the STOP pattern so that the resulting actuation is perceived as much stronger than the ATTENTION pattern.

9.6.1 Static Patterns Refinement Study – Implementation

We implemented the four static patterns as follows:

- ATTENTION – 4 actuators at the back of the head at the same time. Signal 100 ms.
- STOP – 11 total actuators (7 in the front and 4 in the back of the head) at the same time. Signal 100 ms + pause 50 ms + signal 100 ms + pause 50 ms + signal 100 ms (total 400 ms).
- UP – starting at the chin, simultaneously moving up on both sides and ending at the top of the head. Signal 800 ms + pause 100 ms + direction signal (four actuators on the top of the head) 250 ms (total 1150 ms).
- DOWN – starting at the top of the head, simultaneously moving down on both sides and ending at the chin. Signal 800 ms + pause 100 ms + direction signal (single center actuator on the chin) 250 ms (total 1150 ms).

With these improvements implemented, we had to re-validate these patterns for their intuitiveness and recognition accuracy. We also decided to test different pattern playback durations in the same study as we still felt the patterns were too long. Even without UP and DOWN’s repetitions, we wanted to test how fast we could play these patterns without losing intuitiveness and recognition accuracy. Thus, we defined five different duration factor conditions: 0.4, 0.5, 0.6, 0.75 and 1.0, where 1.0 represents patterns as defined above. We modified the Android app that was used in the Pattern Recognition Performance Study in the following ways:

- In addition to STOP, ATTENTION, UP, and DOWN, a fifth option was given so that participants could indicate that they were unsure.

- In the first half of the trials, we did not give participants feedback on whether their guess was correct, while in the second half of the trials, participants received visual feedback on the correctness of their guess. The visual feedback was implemented as an additional view indicating CORRECT or FALSE with the correct pattern after choosing. This design allows assessing the patterns' intuitiveness and how well participants can learn the patterns with feedback training in the second half of the study. However, we do not expect perfect accuracy in the second half, as the feedback was only provided *after* selecting an answer.
- Thus, the participants experienced (in randomized order within the feedback condition):
 - 4 patterns (ATTENTION, STOP, UP, DOWN)
 - × 5 duration factors (0.4, 0.5, 0.6, 0.75, 1.0)
 - × 5 repetitions per pattern
 - × 2 feedback conditions (no feedback, visual feedback)
 - = 200 trials per participant.

We used the third main HapticHead prototype as user interface in this study (see subsection 4.2.3).

9.6.2 Static Patterns Refinement Study – Procedure

The procedure of this study is very similar to the Elicitation Study and the Pattern Recognition Performance Study. After a short greeting and filling out an introductory questionnaire, the participants were handed a Galaxy S8 with the aforementioned Android app and started the trials. Each trial consisted of the participant feeling a particular pattern, then a choice between “attention,” “stop,” “up,” “down,” “not sure,” and “repeat.” After their choice and only in the second part of the study (trials 101-200), the participants got visual feedback on whether they were correct or not.

The repeat option was added as we planned to conduct the study in a crowded computer room with other students, so some distractions were inevitable and desired to create a more realistic environment and increase external validity. If a participant was distracted at a particular trial, they had the option to repeat the pattern.

Finally, the participants filled out another questionnaire on their experience and were compensated with a bar of chocolate. The entire study, including filling out questionnaires, took around 15 minutes per participant.

9.6.3 Static Patterns Refinement Study – Participants

We invited 20 participants (19 male, 1 female, mean age 21.5 y, SD = 3.4 y) for this study, which took around 15 minutes on average, including the questionnaires. None of the participants took part in the prior experiments.

9.6.4 Static Patterns Refinement Study – Results




Figure 9.11.: Confusion matrices of the most interesting results from the Static Patterns Refinement Study. Cell values are absolute with $N = 100$ per row.

The participants had the option of repeating a pattern before deciding on its meaning if they were distracted. Patterns were repeated in 9.1 % of the cases. In these repetition cases, a pattern was repeated 1.3 times on average.

We generated 10 confusion matrices out of the 2 feedback conditions and 5 duration factors. Fig. 9.11 shows confusion matrices for the best-scoring combinations of

feedback conditions and duration factors. The numbers in the confusion matrices are absolute ($N = 100$ for each row of the matrix).

In terms of intuitiveness (no-feedback condition), the UP and DOWN patterns scored high accuracies of 98 % and 97 % at a duration factor of 0.5 and do not improve by much at duration factors higher than 0.5 (see Fig. 9.11a). Even at a duration factor of 0.4, UP and DOWN's accuracies are still rather high at 98 % and 91 %. HOWEVER, the STOP and ATTENTION patterns were mixed up quite often with intuitiveness accuracies of only 72 % and 83 % (Fig. 9.11a). This improves to 79 % and 80 % at a duration factor of 0.75 (Fig. 9.11b). In the feedback condition (while receiving visual feedback), the STOP and ATTENTION patterns improve to 95 % and 97 % at most with a duration factor of 0.75 (Fig. 9.11c).

9.6.5 Static Patterns Refinement Study – Discussion

The STOP and ATTENTION patterns received mixed intuitiveness results of around 80 % correctness. We estimate this is due to the feeling of urgency of, e.g., an ATTENTION pattern with a duration factor of 1.0, which can be similar to a STOP pattern with a duration factor of 0.4 as this translates into one 100 ms actuation vs. three short actuations of 40 ms each (total actuation time 120ms). The reader should keep in mind that with the five different duration factors, we were essentially testing 20 different patterns for four different meanings, leaving more room for error for the participants compared to the Elicitation Study (16 patterns, no training) and the Pattern Recognition Performance Study (4 patterns, training, and evaluation phase, close to perfect accuracies). Therefore, the patterns are not as counter-intuitive as they seem, but this is instead caused by this study's design with different duration factors. Thus, if we were to repeat this study with a set duration factor of 0.75, we would likely get much better results for STOP (total actuation time 225 ms) and ATTENTION (total actuation time 75 ms), as they would differ more clearly.

Even an intuitiveness result of around 80 % is sufficient, especially if participants receive a short training period *before* using these patterns (Fig. 9.11c, shows results *while* conducting training). With these results in mind, we set the final patterns to be used in the Improved System Validation Experiment as follows:

- ATTENTION – 4 actuators at the back of the head at the same time. Signal 75 ms.
- STOP – 11 total actuators (7 in the front and 4 in the back of the head) at the same time. Signal 75 ms + pause 37.5 ms + signal 75 ms + pause 37.5 ms + signal 75 ms (total 300 ms).
- UP – starting at the chin, simultaneously moving up on both sides and ending at the top of the head. Signal 400 ms + pause 50 ms + direction signal (four actuators on the top of the head) 125 ms (total 575 ms).

- DOWN – starting at the top of the head, simultaneously moving down on both sides and ending at the chin. Signal 400 ms + pause 50 ms + direction signal (single center actuator on the chin) 125 ms (total 575 ms).

The very short duration of the ATTENTION pattern (75 ms) combined with the relatively high recognition rate provides further evidence that the overlaying hypothesis H1 of this thesis, that tactile patterns on the head are perceived as *strongly present* by users, is valid. In comparison, most users failed to feel a 250 ms stimulus played back through a tactile belt in related work [206].

In conclusion, these patterns are much shorter than those found in the Elicitation Study and used in the Obstacle Course Experiment while offering comparable recognition performance and are thus a better fit for real-time micro-navigation.

9.7 Improved System Validation Experiment with VIPs




Figure 9.12.: Top view of the obstacle course for the Improved System Validation Experiment, including dimensions. Cardboard obstacle (C1-C5) dimensions are 53.5 cm x 33 cm (height 61.5 cm). Starting and ending zones are marked in blue. The dimensions of the stairs are given in Fig. 9.4.

Since the Obstacle Course Experiment only included two VIPs, we still had to validate our system with a larger number of VIPs. Thus, this experiment has a very similar setup as the Obstacle Course Experiment. However, we improved the implementation in various regards, added more possible paths through the obstacle course, and slightly

altered the obstacle course to include an additional obstacle and an additional starting point (see Fig. 9.12).

9.7.1 Improved System Validation Experiment – Implementation

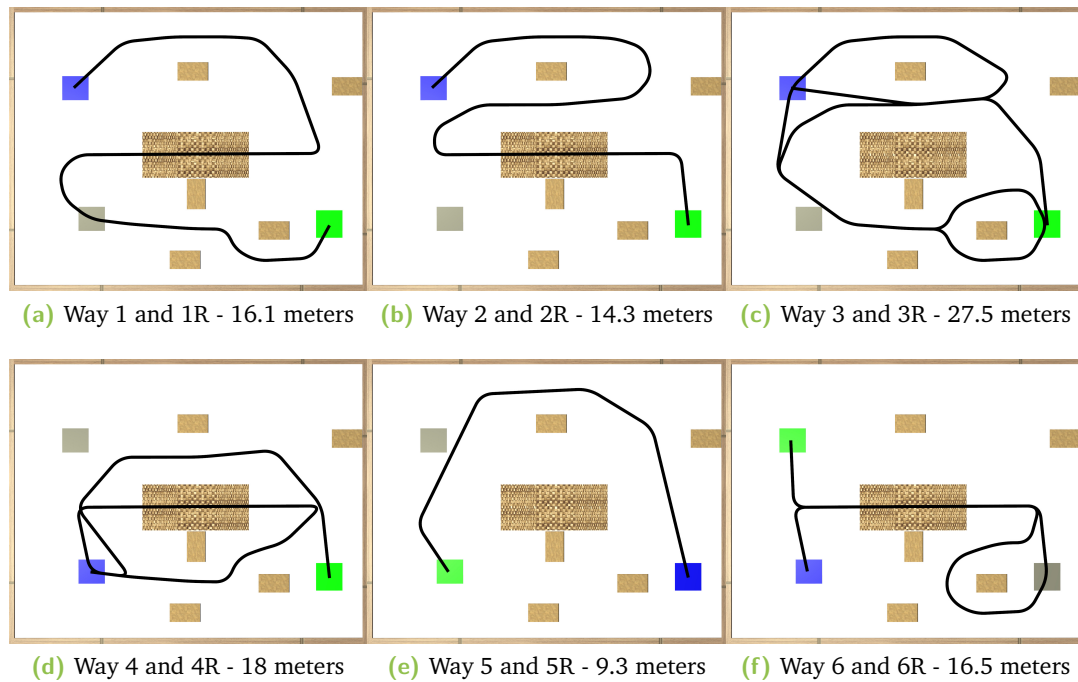


Figure 9.13.: The six possible ways (including reverses) for the Improved System Validation Experiment. Start and end zones are shown in green and blue. The black line shows the smoothed optimal way pre-calculated by our navigation algorithm.

Compared to the Obstacle Course Experiment, we changed the following implementation-specific aspects in the improved version of the system:

1. The users' average speed over the last second and reaction time is now taken into account for the guidance algorithm, leading to more adaptive guidance across participants.
2. Instead of just taking defined waypoints for navigation, the guidance algorithm interpolates smoothly between defined waypoints using thousands of intermediate waypoints generated by a Gaussian moving average on the position (shown as the black line in Fig. 9.13). This avoids instant direction changes and an overall smoother feel of the directional guidance.
3. We added an additional starting/ending position (three in total), moved one of the obstacles, added a fifth one, and doubled the path count to 12 (6 + 6 reverse) as shown in Fig. 9.13. Three of the paths include the stairs once, and one path includes it twice. This was done so that participants could not be sure

of encountering another staircase after the first one. The paths were randomized with the participant ID as a seed.

4. The static patterns used were exchanged with those resulting from the Static Patterns Refinement Study. The ATTENTION pattern is played once the user deviates more than 22 cm from the optimal path, and the STOP pattern is played at greater than 40 cm deviation from the optimal path. Both of these are repeated once every second until the user returns to an acceptable distance to the optimal path. The STOP pattern is also played once before playing the UP pattern when the user approaches the stairs.
5. We changed the tracking system from OptiTrack to an HTC Vive Tracker setup, using four HTC Vive base stations v2.0. These proved to be more reliable and precise than our OptiTrack setup consisting of 24 720p IR cameras (reported error < 1 mm), even though it should, in theory, be less precise as related work suggests an average tracking error of 7.5 mm for the Vive tracking system [128]. However, this related work used a setup of just two Vive base stations v1.0 compared to our four base stations v2.0.
6. The previously used 2.4 GHz Wi-Fi and Bluetooth operating in the same frequency range were sometimes subject to massive distortions for 500 ms every 60 seconds in our lab environment. Therefore, we decided to upgrade the Raspberry Pi 3 to a Raspberry Pi 4 which is able to utilize 5 GHz Wi-Fi. Using 5 GHz Wi-Fi fixed all connection issues we sometimes encountered before (see also subsection 4.2.4).

We mounted two Vive trackers on each of the participants' shoes with Velcro straps, two on each of the shoulders (mounted on a vest) and another one on the head, attached to the HapticHead via Velcro as well. The foot tracking mostly serves for obstacle hit verification. The shoulder trackers' average position and the head tracker's orientation served as input for the directional guidance algorithm. We take the shoulders' average position instead of the head position because of possible head wobbling, just like in the Obstacle Course Experiment.

We used the fourth and final main HapticHead prototype as user interface in this study (see subsection 4.2.4).

9.7.2 Improved System Validation Experiment – Procedure

The procedure of this experiment was like in the Obstacle Course Experiment except for the following changes:

- The number of trials was doubled to 48, but a limit of 2.5 hours, including pauses, was set to the study time. Not all participants completed all trials within this time frame. This is a limitation of this experiment as the trials were still

randomized, and so not every participant conducted the same number of certain routes.

- We only blindfolded participants if they had no residual vision left (including sensing brightness) as it is otherwise not necessary.
- Our participants did not wear a balaclava below the prototype for this experiment because the balaclavas attenuate the vibration intensity, which may lead to feeling only a weak stimulus in certain positions. Instead, we treated the prototype with disinfectant after each participant. This change was implemented due to our own experimentation and observations with the prototype.

In case the tracking system lost tracking for more than 1 s, we repeated a trial. This happened a total of 3 times. The questionnaires used in this experiment are attached in the appendix section A.1 as an example of questionnaires used during user studies of this thesis. All other questionnaires were very similar to this one.

9.7.3 Improved System Validation Experiment – Participants

We talked to several associations for VIPs and recruited 5 participants (all male, mean age 55.4 y, SD = 14.5 y) through word of mouth, public news on websites, and announcements through a local radio station for VIPs. Ideally, the study would have been conducted with more participants, but this was not possible due to the COVID-19 outbreak in 2020. The smaller experiment still provided a wealth of valuable insights. This study took around 1-2.5 hours per participant, including the questionnaires. None of the participants took part in the Obstacle Course Experiment, and all of them are regular white cane users. Participant details are given in Table 9.1.

9.7.4 Improved System Validation Experiment – Objective Results

Table 9.1 shows the five participants' general performance and their demographics. Like in the Obstacle Course Experiment, we were quite conservative in counting obstacle hits and stairs assistance. If the edge of the shoe rim was on or over the side-ledge of the staircase, this event was counted as needing assistance, and obstacle hits were counted as soon as a participant brushed an obstacle with their clothing.

Unsurprisingly, almost all obstacle collisions happened between C3, C4, and C5 (see Fig. 9.12). Only a single collision happened at C1 when one of the participants was distracted by talking. No collisions happened at the stairs as the STOP pattern was played early enough so that the participants were careful, and they also got the UP pattern right on time to lift their foot to the first stair. The six times in total where we had to assist the participants on the stairs were usually because the participants accidentally got too close to the edges of the stairs or when their first step onto the

Participant ID	Sex	Age [years]	Visual impairment diagnosis	Residual vision	Total trials completed	Median time per trial [s]	Average time per trial [s]	SD of time per trial [s]	Stairs assistance count	Obstacle hit count	Issues in % of conducted trials	Walking speed [m/s]
1	m	48	Glaucoma, fully blind since age 16	None	48	51.5	49.5	19.4	3	4	14.6	0.34
2	m	68	Glaucoma, fully blind since age 65	None	14	124.7	148	83	2	7	57.1	0.11
3	m	56	Retinal detachment, fully blind since age 20	None	48	63.7	68.1	32.4	0	1	2.1	0.25
4	m	70	Glaucoma, impairment since age 64	30 pct	24	187.2	191.9	101.2	0	1	4.2	0.09
5	m	35	Retinal detachment, fully blind since age 25	None	48	74.4	76	33.1	1	8	18.8	0.22
Average		55.4				94.2	96.4	46.5	1.0	3.5	9.9	0.23

Table 9.1.: Improved System Validation Experiment – participants overview. Note that P2 was excluded from the average results due to broken cables of actuators (see text).

lower stair was too close to the edge, and we wanted to make sure that they would not fall off with their next step.

When comparing the counts for obstacle hits and stairs assistance in Figures 9.6 and 9.1, the reader has to keep in mind the different number of trials of the Obstacle Course Experiment (24 trials per participant) and the Improved System Validation Experiment (varying number of trials, up to 48).

The results of participant P2 have to be viewed with caution. P2 was curious about the prototype on his head, kept touching it, and ended up breaking a total of three cables of actuators during the experiment, including one on the forehead, which is the most important one for navigation. Thus, P2 had significant disadvantages compared to the other participants regarding navigation, as some actuators stopped working during the experiment. This made P2 more likely to bump into obstacles and increased the deviation from the optimal path, as seen in Fig. 9.14. For this reason, we excluded P2 from all evaluations involving aggregated data (e.g., the averages in table 9.1).

Fig. 9.14 depicts the mean left (red) and right (blue) deviations from the optimal path for all trials conducted by the respective participant. We found an overall average absolute deviation to the optimal path of 5.7 cm (SD=1.4 cm) across all participants. We also report the percentage of time on the left side of the optimal path and the mean deviation towards the left or right here as a measure of a systematic error that the participants experienced. With the prototype perfectly fit to participants, they should ideally have a mean deviation of 0 cm, which would indicate a symmetric movement to the left as to the right side of the optimal path. However, even when the prototype was perfectly placed on the participants (e.g., central actuator on the forehead was perfectly above the nose), they might still experience a slight systematic error, as the center above the nose might not *feel* central to the participant (see discussion about this phenomenon in subsection 9.5.5). This is especially true for individuals who have been fully blind for a long time and develop a condition where they tend to tilt their




Figure 9.14.: Improved System Validation Experiment – mean left and right median deviations from the path. The red boxplot shows mean left deviations, the blue boxplot mean right deviations. Note that P2’s results have to be viewed with caution due to broken cables of actuators (see text).

head a bit off-center permanently (e.g., to compensate for hearing loss in one of the ears) [136].

Fig. 9.15 shows the different wobbling patterns of the five participants around the optimal paths for an easy path (W5) and a more complicated path (W1). Each wobble from left to right and back represents two full steps.

9.7.5 Improved System Validation Experiment – Subjective results

Fig. 9.16 shows the results of our closing questionnaire from this experiment. Generally, these results are similar to the subjective results in the Obstacle Course Experiment. The participants generally agreed that they felt safe and that they liked the way of navigating, which felt intuitive for most of them. While they thought that the tactile feedback on the head was suitable, some had concerns about using the system regularly concerning well-being as the prototype is still somewhat clunky and uncomfortable to wear over more extended periods. A majority would worry about the system’s aesthetics, even if it were hidden under a beanie. They were especially concerned about certain social situations, such as public presentations, dance parties, saunas, or swimming pools. Using the system feels more appropriate to the participants in other social situations such as closed or public rooms, walking through a city, cafes, restaurants, or at home.

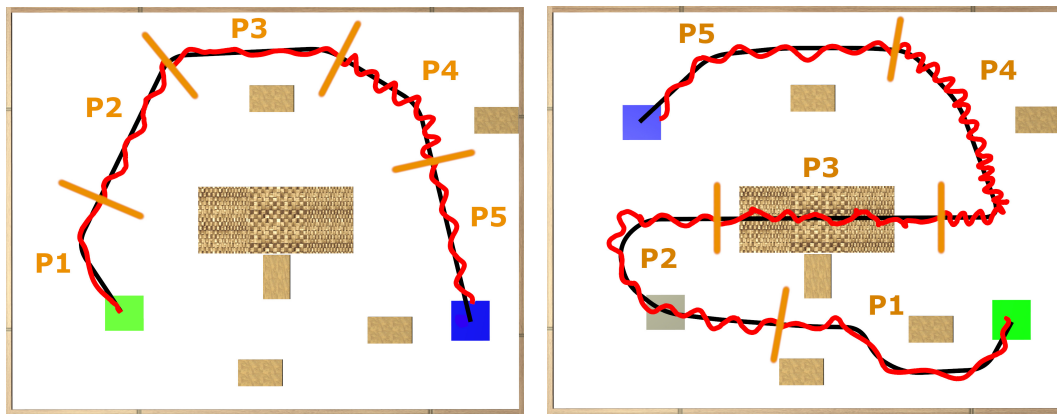


Figure 9.15.: Improved System Validation Experiment – wobble of different participants on way 5 (left) and way 1 (right). The red line shows the shoulder tracker average of the respective participant.

P2 rated our system lower than the other participants regarding certain properties (intuitiveness, confusion of static patterns, safety, strength of vibration feedback). This was likely caused by the broken actuator cables and the resulting less intuitive continuous guidance signal, which also had a much lower intensity for P2 because of the broken actuator on the forehead.

Some of the participants stated that they could imagine using this kind of navigation in everyday life as well (presuming a suitable input tracking system, e.g., Apple ARKit [6] or Google ARCore [58]), and most judged that it could replace the white cane in micro navigation. However, two participants had concerns, primarily because they missed the white cane feedback about ground properties, which the system does not provide. Others stated that this was no issue as they judged that they get sufficient feedback on ground properties from their shoes.

In other verbal and written comments, participants stated that they could imagine using a further developed system based on our prototype in most social contexts if suitably disguised or hidden. All of them encouraged us to continue research in this direction, and P3 even mentioned feeling a bit bored at the end of the study because it was “too easy” to follow the directions and no total concentration, e.g., on the stairs, was necessary.

9.7.6 Improved System Validation Experiment – Discussion

Due to our experiment’s nature, the low number of participants, and the considerable differences between participants, we first discuss each participant’s results individually and highlight important characteristics of each participant that might influence results.

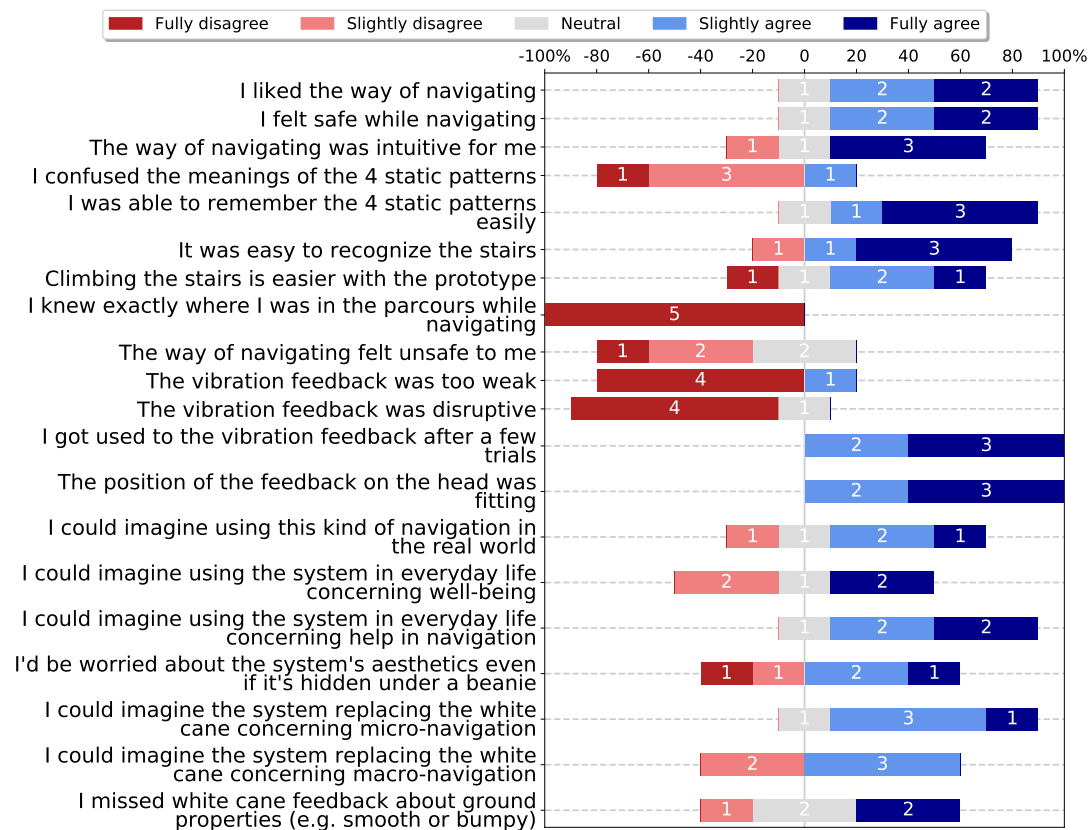


Figure 9.16.: Questionnaire results of the Improved System Validation Experiment.

P1 was amongst the younger participants, quite physically fit, and had been fully blind since a young age. He was the fastest of the participants at an average duration of just 49.5 seconds per trial. Due to his duration and the narrow pathway between the obstacles C3, C4, and C5 (see 9.12), he hit a total of four obstacles and required assistance on the stairs three times due to the sharp turns right before the stairs, which in combination with his short duration caused him sometimes to be misaligned to the stairs in the first third of the study.

While P2 was physically fit, he had only been fully blind for three years before the experiment. He had difficulties keeping his concentration on the tactile signals since he liked to chat. He was very curious about the prototype on his head and kept touching it, resulting in a total of three broken actuator cables. Even though he was slower than most of the other participants, he required assistance on the stairs twice and hit seven obstacles during the experiment. We suppose his lower than average duration per trial is not just the result of broken actuator cables, but also of him only having been blind for three years in total, so he was somewhat cautious. We had to stop the experiment after 14 trials because of the broken actuator cables, and the participant had other appointments to attend to, so he could not wait for us to resolder the actuator cables.

Due to the broken cables, his results have to be viewed with caution and have been excluded from the overall quantitative analysis.

P3 was quite physically fit and had been fully blind since a young age, just like P1. His average duration per trial was a bit lower than P1's, but he managed to react perfectly to the tactile signals as apparent from his record low mean absolute deviation from the optimal path of just 4.1 cm, including a 1.5 cm systematic error (see Fig. 9.14). He required no assistance on the stairs and hit only a single obstacle along the way.

At age 70, P4 was the oldest of the participants. When he arrived, he was not aware that the obstacle course included stairs but still wanted to participate in the study when we told him the specifics. He usually climbs stairs using the railings and walks much more slowly than the other participants in the real world. This can be attributed to his age and visual impairment, which had only manifested six years ago. Furthermore, he also had some residual vision blocked by the sleeping mask that participants with residual vision had to wear during the experiment. His trials took much longer on average as he was very cautious, especially on the stairs. Still, he was able to follow the tactile signals very well, required no assistance on the stairs, and only hit a single obstacle. We stopped the experiment after half the trials were completed because P4 seemed exhausted. He asked to continue, but we did not want to risk him falling due to fatigue.

P5 was the youngest participant at 35 years old, technophile, and had been fully blind for ten years. He has a condition where his head is tilted slightly sideways and forwards. Due to this, we had to re-adjust the Vive tracker on his head after observing the first three trials, where he constantly walked to the right of the optimal path by a large margin. After the readjustment, he performed pretty well even though the data in Fig. 9.14 suggests that he now experienced a systematic error to the left side, as we slightly over-corrected his pose. He brushed more obstacles than the other participants due to his talkativeness and resulting fading concentration on the task.

Our improved system performed more accurately than the system used in the Obstacle Course Experiment. For example, in the Obstacle Course Experiment, the average absolute deviation from the optimal path was 9.3 cm on average (compare Figures 9.8, 9.9 and 9.14). With the improved system, the participants scored an average absolute deviation from the optimal path of only 5.7 cm, which is considerably lower than in the Obstacle Course Experiment (excluding P2 due to broken cables as discussed above) and only a little more than the usual shoulder wobble while walking, as shown in Fig. 9.15. This is not only the result of the improved system but also because the systematic directional error decreased to an average absolute of 1.9 cm ($SD=0.8$ cm). The systematic error decreased compared to the Obstacle Course Experiment, even though we still did not perform a sophisticated user-specific calibration, as explained in the Obstacle Course Experiment's discussion. However, we were aware of the systematic error and ensured to place the prototype perfectly on our participants'

heads this time, ruling out one of the three influences towards the error. We also pre-calibrated the Vive tracker position on the prototype and only had to adjust it once for P5, as discussed above.

In terms of obstacle hits, we can conclude that even with some distraction through occasional talking while conducting the study, none of the participants hit obstacle C1 or the wall when passing through the 115 cm wide gap between the wall and C1. Only a single participant hit C1 when passing through the 89 cm wide gap between C1 and the stairs as he was distracted by talking. Regarding the gap between C3, C4, and C5, when coming from the left, some participants hit C4, and when coming from the right, they either hit C3 or C5. These obstacle hits were usually the cause of a short distraction period and resulting slower reactions to the tactile signal, causing the participants to miss the rather sharp turn or execute it too late. We explicitly expected some obstacle hits as we deliberately designed this study to measure the maximum precision we could achieve with the system. With this in mind, we were positively surprised about the low number of obstacle hits and that there occurred only two obstacle brushes at C3 or C5 when coming from the left. The gap is 71 cm wide, while the hip widths of the participants were between 40 and 50 cm, leaving only very little room for error.

We found the minimal safety distance between obstacles without distraction to be about 89 cm (as one of the participants hit C1 while being distracted), and even with distraction through talking, every single participant managed to always safely pass through the 115 cm gap between the wall and C1, so we can assume this to be a safe distance between obstacles with distraction through talking. Our system worked better for participants who were (a) fully blind for a long time and (b) younger than the average. Age negatively influences reaction time and perceived strength of the tactile signals. Our only participant with residual vision (P4) was extra cautious at blocking the residual vision. P4 felt uncomfortable wearing a sleeping mask in the experiment. Furthermore, we observed that the longer participants had been blind in their lifetime, the more courageous and faster they were in terms of walking speed, even without the white cane.

9.8 Overall Discussion

In summary, we identified four suitable static tactile patterns around the head for four fundamental navigation instructions (START, STOP, UP, DOWN). In a series of experiments, the patterns were shown to be easily recognizable, intuitive, and easy to remember. In preparation for the Obstacle Course Experiment, we implemented these patterns in conjunction with a continuous guidance pattern and an additional ATTENTION pattern (similar to STOP). Our initial tests for the Obstacle Course Experiment revealed that an additional ATTENTION pattern is essential to keep users focused on

the task. This was not obvious from our informal interview and only surfaced after we conducted several preliminary tests in the Obstacle Course Experiment with other researchers, who kept losing concentration when they tried to give us feedback while navigating.

In the Obstacle Course Experiment, we guided sighted but blindfolded and real VIPs through an obstacle course, including stairs. Due to our observations in the Obstacle Course Experiment, we discarded the initially designed START pattern in later experiments because users were aware that the navigation was working when any kind of tactile stimulus was presented. We gained valuable insights into improving our system and further optimized our static patterns in the Static Patterns Refinement Study. The precise definitions of the final static pattern parameters are given in Section 9.6.5.

Finally, we incorporated all identified improvements into our system and validated it with another set of five VIPs in the Improved System Validation Experiment. We showed that the designed static patterns for essential navigation instructions, in combination with the continuous guidance feedback to find objects around the user from section 8.4, work very well to keep VIPs on a predefined path in a complex micro-navigation task. With an average absolute deviation of 5.7 cm, the Improved System Validation Experiment participants diverged only slightly from the optimal path than the usual shoulder wobbling. Even when the user is distracted by talking, the safe distance between obstacles was determined to be 115 cm, which is less than the typical width of sidewalks [210].

When comparing our final system with related work, it can be stated that it performs better than any macro-navigation vibrotactile head, belt, or foot systems (e.g., [71, 150, 180, 215]). This is to be expected because of the low tracking accuracy of GPS. The only other work we are aware of that used an indoor tracking system with decent accuracy (< 10 cm) for a micro-navigation task is Flores et al. [52]. They did not include any obstacles or stairs in their study, and their paths only included two turns, each in an empty space. They measured an average absolute distance to the optimal path using their tactile belt of 49 cm and compared it to auditory guidance in the same study, which resulted in an average distance of 61 cm. Our system performed an order of magnitude better at 5.7 cm average absolute distance to the optimal path, including a 1.9 cm average systematic error. A sophisticated per-user calibration may reduce this systematic error. The increased precision comes at the price of a slower walking speed (0.23 m/s vs. 0.46 m/s in [52]). This is likely a result of our course, including stairs and obstacles. Moreover, the participants in our final experiment had a relatively high average age of 55.4 years.

Due to our system's high precision, it allows use cases in new scenarios that were not possible with prior systems [52] (e.g., navigating through a crowded bar, through a festival, or on a sidewalk while avoiding obstacles and other people). Our approach

also presents itself as a highly effective *anti-veering tool* for VIPs. Veering off an intended path is a common issue for VIPs [36]. Even with all the turns and the stairs in our obstacle course, the average absolute deviation from the optimal path at 5.7 cm is still better than other anti-veering tools that just use a straight line as an experimental condition (e.g., [36], 9 cm deviation, using a 6-actuator vibrotactile strap on the forehead).

The objective results in subsection 9.5.3 and subsection 9.7.4 indicate that the overlaying hypothesis **H3** on the precision of micro-navigation of this thesis (see section 1.1) is accurate as the experiments prove that the precision of the HapticHead tactile user interface is almost an order of magnitude more precise than a tactile belt used for the same purpose in prior work [52].

9.8.1 Limitations

Noise is still an important factor when working with vibration actuators around the head as noise is transferred via bone conduction and especially loud and possibly disruptive near the ear openings. Even though most of our VIP participants were very optimistic about the noise and told us that it did not bother them, we suppose that this is probably not the case for everyone in a larger sample of the VIP population. When designing a prototype to be used by VIPs, we have to consider this limitation more than with sighted people because VIPs rely on hearing a lot more and would probably be disoriented by loud noises next to their ears. An alternative are LRAs or specialized voice coil actuators [125] that allow controlling amplitude and frequency independently. These can reduce reaction times and audible noise by working at lower frequencies and higher amplitudes at the cost of a larger form-factor. Working at higher amplitudes is required because the vibration perception threshold decreases dramatically as vibration frequency decreases [82].

Our prototype's current aesthetics should undoubtedly be improved before being used in any real-world social situation. A miniaturized version of HapticHead integrated into a beanie (similar to [41]) would mostly solve the aesthetics issue and reduce the concerns of our study participants.

This chapter does not evaluate the effect of different hair densities on tactile pattern recognition or guidance precision. Related work found no significant effect of hair density on localization performance on the head [41]. Furthermore, while [146] did find a significant effect of hair density on vibration perception threshold, as long as a tactile display operates above that threshold, as our prototype did, this should have no substantial influence on localization performance [41]. Our experiment participants had very different hair densities, except for the last validation experiment, where all participants were male and had relatively light hair. We chose not to evaluate a possible effect of different hair densities as we did not have a systematic sample of

different hair densities and thus not enough data to draw reliable conclusions on this aspect. While we expect a small effect of strong or curly hair on our four static patterns' recognition accuracy, we do not expect a significant effect on the continuous guidance stimulus, as this stimulus is present on the forehead if the user is following the path.

9.9 Conclusion and Future Work

We presented a tactile micro-navigation system with substantially higher precision than prior work. Aside from the limiting factors mentioned earlier, our study participants were quite content with the system as is, stating a feeling of safety and intuitiveness with the navigational method. With a further refined and adequately integrated system, more training, and combined with a proper self-contained indoor and outdoor tracking method, our tactile guidance approach may eventually improve the lives of many VIPs. It provides a viable information channel about the best path around obstacles in addition to information obtained by the white cane. A robust, commercial variant of the system may potentially fully replace the white cane, but users should still carry a foldable backup white cane with them in case of, e.g., a battery failure.

While our system is primarily intended for VIPs, other use cases that require precise tactile guidance are feasible with the same system. Possible use cases include guidance for firefighters or other personnel operating in low-vision environments, guidance in virtual and augmented reality scenarios where navigation instructions should not occupy the visual and auditory channels, and guidance and warnings for jet or drone pilots. Navigation in 3D spaces with the same system is achievable, as hinted at in the prior chapter, but this needs to be verified in future work.

Future work may investigate a variety of exciting research directions. Our system may be combined with a system that detects moving obstacles and *prevent future dynamic collisions* by warning the user through the STOP pattern or dynamically guide the user around the moving obstacle. Our system could also be improved by using specialized voice coil actuator types such as [125] for lowering the audible noise through lower frequencies or using regular linear resonant actuators (LRAs), as both of these options offer significantly faster reaction times to changes in voltage, compared to the ERM actuators we used. Switching to a voice coil based actuator type (e.g., [125] or LRAs) will allow users to react slightly faster (~ 100 ms) to changes of the guidance stimulus, which should further improve the guidance precision.

Finally, research should be conducted on how accurate a navigation system has to be for different micro-, macro-, and combined navigation scenarios. We concentrated on developing a tactile system with maximum precision. However, this maximum precision may not always be necessary or desirable, and users may still trust micro navigation systems with less precision depending on the scenario. A significant advantage of a

system with less precision may be that it requires less user attention. We imagine an adaptive system, which decreases or increases the strength of the continuous guidance stimulus. The adaptive system may even fully turn it off if the user follows the optimal path within a certain margin of error. This adaptation may depend on the current environment and possible obstacles (e.g., 5 m in a large open area, 25 cm on a sidewalk, and 10 cm in a bar).

Summary and Conclusion

10

The following paragraphs summarize this thesis and describe how we were able to reach the overarching goal of establishing the head as a means for tactile communication laid out for this thesis (see section 1.1).

The first subgoal G1 of this thesis was to research tactile sensitivities and relevant tactile illusions on the head to gain an understanding of human tactile perception on the head with implications for tactile user interface design. We achieved this goal by reviewing literature and conducting our own experiments on these subjects in chapter 3. As a result, we calculated a detailed heatmap of tactile localization accuracies around the head from our available experiment data and enriched this information by additional data on the funneling illusion occurrence. Notably, we validated an experiment from Kerdegari et al. [105] and corrected their problematic midline bias conclusion for the forehead, caused by an improper study design. With our fundamental research in mind, future head-mounted tactile user interfaces may benefit in terms of performance and user comfort by designing their systems and algorithms according to our research results. For instance, according to our data, it makes sense to place a higher density of actuators on head regions with a relatively higher localization accuracy (e.g., forehead, top of the head, chin) to optimize a tactile system for guidance precision and user comfort. The latter is achieved through users perceiving less abrupt switches between individual actuators when they are relatively closer together on head regions with a relatively lower funneling illusion threshold distance.

Reviewing the literature and our research in chapter 3 also helped us fulfill the second subgoal G2, to construct a tactile interface around the head, based on what we learned. The design approach, its limitations, and the HapticHead tactile user interface evolution through multiple iterations are shown in chapter 4. While not all prototypes ended up being evaluated in experiments, we report valuable lessons learned while iterating the design of HapticHead and trying different actuator types. In particular, future head-based tactile user interfaces should not increase vibration amplitude at all costs as this leads to noisy, bulky, and power-hungry designs which cannot easily be deployed or integrated into a potential commercial product.

The third subgoal, G3, to research how many and what kinds of tactile patterns on the head users can distinguish and which kinds of patterns are accepted among the

user base, was fulfilled through our experiments presented in chapter 5. We were able to show that the addition of spatial location dramatically increases the total distinguishable number of patterns available to the user compared to sole intensity and rhythm modulation and which kind of patterns are more accepted amongst the experiment participants in terms of recognition accuracy and comfort. These results are especially relevant to all head-based tactile user interfaces working with many tactile instructions. These systems should design their tactile patterns such that they are easily distinguishable (utilizing mostly pattern location as a modifier) and comfortable to the users by interpolating and thus providing smooth transitions between actuators.

Our final subgoal G4 of this thesis was to build a pattern design software that can deal with a large number of actuators in complex configurations. G4 was fulfilled by chapter 6 where we describe the process of the design and implementation of our tactile design software VRTactileDraw and its evaluation. We found that even novice users with non-technical backgrounds can work with VRTactileDraw to design relatively complex tactile patterns for various use cases. Because our system is easily extensible and released alongside this thesis as open-source software, the designers of future tactile systems all around the human body may use VRTactileDraw to prototype tactile patterns for their respective systems rapidly.

Finally, we achieved our overarching goal of establishing the head as a means for tactile communication through evaluating head-based tactile feedback in several use cases.

In chapter 7, we showed that adding vibrotactile feedback around the head can significantly increase the level of presence users experience in certain VR scenes. This increase in presence through tactile stimuli can be applied in immersive games and specific VR training simulations where the level of presence or spatial awareness (including collision prevention) is essential such as in complex maintenance jobs, anxiety therapy, or flight training.

Our experiments in chapter 8 indicate that HapticHead tactile guidance is significantly faster and more precise than spatial audio for finding visible and invisible virtual objects in 3D space around the user. We further successfully navigated blindfolded users to actual household items at different heights using HapticHead as a 3D guidance Augmented Reality device. Our 3D guidance is substantially faster and more precise than other related work approaches, which used less complex tactile user interfaces. It is transferrable to a large variety of other use cases, as hinted at in the introduction (see chapter 1).

Interestingly, Matsuda et al. [132] replicated our 3D target search experiments from section 8.3 with a less complex tactile user interface featuring 16 actuators around the neck (“*HapticPointer*”). While their system has the advantage of being hidden under a t-shirt, it performed substantially worse in all measures, providing evidence that the slightly more complex HapticHead system is far superior in 3D guidance performance

compared to simpler tactile user interfaces. Specifically, they reported 2.5 times higher median task completion time, 8.8 times higher movement overhead, and a lower success rate of 90.7 % compared to 96.4 % for HapticHead.

In chapter 9, we tested explicitly designed tactile patterns in two micro-navigation experiments with blindfolded and visually impaired users. We found that our final implemented micro-navigation guidance system performs almost an order of magnitude more precise than a state-of-the-art tactile belt from related work. In addition, our system can signify important navigation instructions such as “stairs up” or “STOP” to the user through highly intuitive tactile patterns using natural mapping, which was not possible with prior systems. Our study participants were quite content with the prototype system, stating a feeling of safety and intuitiveness with the navigational method. Apart from precise micro-navigation guidance, our system may even replace state-of-the-art anti-veering systems designed to simply keep the user on a straight line due to its outstanding precision.

With the large variety of use cases mentioned above, we successfully established the head as a means for tactile communication and thus fulfilled this thesis’s overarching goal. However, there is still an abundance of possible use cases for tactile feedback around the head, which could be evaluated in future work (see next chapter).

In terms of hypotheses (see section 1.1), the subjective results in subsection 5.4.5 indicate that the overlaying hypothesis H1 of this thesis on the perception of tactile patterns feeling strongly *present* to the user is accurate. Most study participants agreed that they could easily recognize TPs on the head, intuitively recognized them, and strongly disagreed that the vibration feedback on the head was too weak. All of these measures were substantially weaker for the smartphone condition in the same experiment. Furthermore, the objectively very short duration of our ATTENTION TP in sections 9.6 and 9.7 combined with a high recognition rate provides further evidence that H1 is valid. In comparison, all users failed to feel a 250 ms stimulus played back through a tactile belt in related work [206]. Thus, we can accept H1.

All three experiments of chapter 8 indicate that the overlaying hypothesis H2 on the *intuitiveness* of 3D guidance of this thesis is accurate. Almost all study participants agreed that they could intuitively map the vibrotactile feedback to virtual targets, and a majority could also intuitively map them to real targets in the third experiment. Therefore, H2 can be accepted.

The objective results in chapters 8 and 9 prove that the overlaying hypothesis H3 on the *precision* of 3D guidance of this thesis is accurate as our objective guidance precision in the experiment in section 8.5 compared to the other feedback modalities was substantially better than auditory feedback and almost on par with visual feedback. Furthermore, the object finding accuracy in the blindfolded object finding task (see section 8.6) was almost perfect. In terms of micro-navigation precision, the experiments in chapter 9 prove that the precision of the HapticHead tactile user interface

is almost an order of magnitude more precise than a tactile belt used for the same purpose in prior work. Consequently, we can accept H3.

This thesis makes several significant contributions to the field of tactile research:

- An in-depth investigation into tactile localization accuracy and the occurrence of the funneling illusion around the head,
- HapticHead, an advanced tactile display around the head,
- fundamental research on tactile patterns around the head,
- a Virtual Reality tactile pattern design software,
- an investigation on improving the feeling of presence in Virtual Reality through head-based tactile feedback,
- an investigation on HapticHead's 3D guidance performance in Virtual and Augmented Reality, and
- a tactile system to micro-navigate people with visual impairments.

The details of these contributions are discussed in section 1.2. Major parts of this thesis have previously been published in journals, conferences, and workshops [91, 93, 94, 95, 97, 98, 99, 100, 101, 102, 103].

10.1 Limitations

The HapticHead concept features some inherent limitations described in subsection 4.1.1 and the individual discussions of several chapters on use cases.

The most severe limitations include the concept's aesthetics in public scenarios and the vibration noise. It is impossible to entirely hide the system under a beanie due to the chin strap, which is vital for intuitive 3D guidance. This aesthetics issue prevents the use of the system in several use cases involving public situations. Secondly, the vibration noise caused by actuators vibrating close to ear openings is quite significant even though our user studies' participants rarely complained about it.

The vibration noise may be decreased by using lower frequency actuators as discussed in subsection 4.3.2. However, these will also have to provide higher vibration amplitudes as the vibration sensitivity of the Pacinian corpuscles and Pacinian type receptors in hairy skin decrease with lower frequencies below 200 Hz. The issue with high-amplitude actuators is generally that they proportionally increase in size, leading to power-hungry, clunky actuators (see subsections 4.3.1 and 4.3.2).

The following chapter on future work will highlight a vision to improve the system's aesthetics to make it more socially acceptable for 2D micro-navigating people with visual impairments around obstacles in public scenarios.

Future Work

11

An abundance of future work opportunities emerges from generating tactile feedback around the head. This thesis provides an in-depth investigation into the fundamentals of head-based tactile feedback and several application scenarios. However, there are many more scenarios, mostly in the domain of precise 3D tactile guidance, that could be further explored in future work (some already mentioned in the introduction):

- indicating plane positions and distances around air traffic controllers,
- displaying enemy positions and distances around fighter jet pilots,
- providing obstacle position warnings around drone pilots,
- signifying fish and obstacle positions around scuba divers,
- helping to find sought constellation or star positions above stargazers,
- indicating approaching traffic towards skiers (collision feed-forward),
- displaying object positions around crane operators,
- enriching 360° videos with haptic feedback,
- using HapticHead combined with a 360° camera or LIDAR scanner for spider-sense, similar to [16, 26], and
- using HapticHead for VR sickness prevention, improving upon [157].

Apart from concentrating on novel use cases, future work may also improve upon existing limitations of the HapticHead concept (see subsection 4.1.1) and use cases presented in this thesis. For example, our study on increasing presence in VR scenarios in chapter 7 indicated possible improvements of the tactile feedback modalities in future work.

One of the most compelling use cases of head-based tactile feedback is the guidance and micro-navigation of people with visual impairments. The extensive research presented in chapter 9 highlighted several possible improvements of the HapticHead prototype that should be researched in future work. The following subsection will present a vision of a head-based tactile micro-navigation system based on HapticHead that improves upon the two most severe issues with the current design: aesthetics and noise.

11.1 Vision of a Hidden and Precise 2D Tactile Micro-Navigation System




Figure 11.1.: Vision of a future 2D micro-navigation system for people with visual impairments, based on the original HapticHead concept.

With all that was learned from working with tactile feedback on the head and guiding blindfolded and visually impaired people, this thesis closes with a final vision of a future mobile 2D micro-navigation system for people with visual impairments. Figure 11.1 shows the tactile user interface of this vision, based on the original HapticHead concept. Some actuators were moved, and the 2D actuator ring that included the chin strap in the original model is missing. Furthermore, the actuators' spacing on the forehead more closely resembles the localization accuracy and funneling illusion occurrence (see section 3.5). This new spacing should further improve comfort and guidance precision, but this hypothesis remains to be validated in future work.

Compared to the original HapticHead concept, this visionary tactile user interface has the advantage that this system can be fully integrated into a beanie, similar to [41]. Thus, it is invisible to the public and can be used as a discrete guidance tool, fixing the stigma issue of people with visual impairments as they can no longer be distinguished from people with regular sight because their assistive system is invisible.

The visionary model has some disadvantages compared to the original HapticHead concept: It can no longer offer 3D guidance but is limited to 2D guidance due to the

missing chin strap, and the intuitiveness of the UP and DOWN patterns is slightly decreased as these patterns would now move between the back of the head and the top. The head's back would not be clearly identifiable as "below the user" as the natural mapping no longer fully applies.

In terms of actuator type, this future system should use LRAs or other future voice coil actuators with a low frequency as mentioned in subsection 9.8.1 to lower the perceived amplitude of the actuator noise. The Lofelt actuators we tested in subsection 4.3.2 were too power-hungry and clunky for our purpose, but a future, miniaturized version of the same actuator could produce an exciting actuator type for head-based tactile feedback.

This visionary model of a tactile output system should be combined with a real-time AR position and orientation system such as Google ARCore [58] or Apple ARKit [6] in future work. This thesis mentions that we implemented a similar fully mobile obstacle avoidance and micro-navigation system based on a prototype Google Project Tango device combined with HapticHead (see subsection 9.3.1). However, when we tested this prototype system, it was apparent that the available technology was not sufficient to provide real-time obstacle detection with acceptable accuracy. One of the main reasons for this was the low framerate of only five frames per second of the depth sensor integrated with the Lenovo Phab 2 Pro device that we used.

The latest Apple iPad 2020 Pro and iPhone 12 Pro, released shortly before this thesis was submitted, already have included LiDAR scanners in addition to the other cameras. This additional LiDAR sensor will likely significantly improve the AR environment recognition and accuracy in future versions of ARKit. Thus, a position and orientation system based on this or a similar system will eventually detect static and dynamic obstacles in the user's environment with acceptable accuracy in real-time. Combining this future AR position and orientation with the visionary model of the tactile user interface above will result in a micro-navigation system capable of guiding the user around static and dynamic obstacles in the real world, inside and outside, with similar high precision as observed in our final experiment.

A robust, commercial variant of the combined system may eventually fully replace white canes and guide dogs, improving the lives of many affected by visual impairments by providing them with warnings of elevation changes and high-precision guidance to targets as well as around static and dynamic obstacles. It would also help to resolve the stigma issue of using a visible assistive tool.

Bibliography

- [1] Victoria E. Abraira and David D. Ginty. “The Sensory Neurons of Touch”. In: *Neuron* 79.4 (Aug. 2013), pp. 618–639. ISSN: 08966273. DOI: 10.1016/j.neuron.2013.07.051 (cit. on pp. 14, 15, 17–19, 21).
- [2] Teemu Tuomas Ahmaniemi and Vuokko Tuulikki Lantz. “Augmented reality target finding based on tactile cues”. In: *Proceedings of the 2009 international conference on Multimodal interfaces - ICMI-MLMI '09*. New York, New York, USA: ACM Press, Nov. 2009, p. 335. ISBN: 9781605587721. DOI: 10.1145/1647314.1647383 (cit. on p. 26).
- [3] Mike Alger. *Visual Design Methods for Virtual Reality*. 2015. URL: <https://api.semanticscholar.org/CorpusID:31705053> (cit. on pp. 93, 94).
- [4] David Alles. “Information Transmission by Phantom Sensations”. In: *IEEE Transactions on Man Machine Systems* 11.1 (Mar. 1970), pp. 85–91. ISSN: 0536-1540. DOI: 10.1109/TMMS.1970.299967 (cit. on p. 37).
- [5] Anonymous. *Pigpio library*. 2020. URL: <http://abyz.me.uk/rpi/pigpio/> (visited on Mar. 28, 2019) (cit. on pp. 40, 66).
- [6] Apple Inc. *Apple ARKit*. 2020. URL: <https://developer.apple.com/augmented-reality/arkit/> (visited on July 2, 2020) (cit. on pp. 138, 167, 181).
- [7] Mojtaba Azadi and Lynette A. Jones. “Vibrotactile actuators: Effect of load and body site on performance”. In: *2014 IEEE Haptics Symposium (HAPTICS)*. IEEE, Feb. 2014, pp. 351–356. ISBN: 978-1-4799-3131-6. DOI: 10.1109/HAPTICS.2014.6775480 (cit. on p. 79).
- [8] P Bach-y-Rita, K A Kaczmarek, M E Tyler, and J Garcia-Lara. “Form perception with a 49-point electrotactile stimulus array on the tongue: a technical note.” In: *Journal of rehabilitation research and development* 35.4 (1998), pp. 427–430. ISSN: 0748-7711 (cit. on p. 142).
- [9] Paul Bach-Y-Rita, Carter C. Collins, Frank A. Saunders, Benjamin White, and Lawrence Scadden. *Vision substitution by tactile image projection [18]*. Mar. 1969. DOI: 10.1038/221963a0 (cit. on pp. 26, 142).
- [10] Jan Balata, Jakub Franc, Zdenek Mikovec, and Pavel Slavik. “Collaborative navigation of visually impaired”. In: *Journal on Multimodal User Interfaces* 8.2 (June 2014), pp. 175–185. ISSN: 1783-7677. DOI: 10.1007/s12193-013-0137-9 (cit. on pp. 138, 142).
- [11] Jan Balata, Zdenek Mikovec, and Ivo Maly. “Navigation Problems in Blind-to-Blind Pedestrians Tele-assistance Navigation”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 9296. Springer Verlag, 2015, pp. 89–109. ISBN: 9783319227009. DOI: 10.1007/978-3-319-22701-6_8 (cit. on pp. 138, 142).

- [12] Aaron Bangor, Philip Kortum, and James Miller. “Determining what individual SUS scores mean”. In: *Journal of Usability Studies* 4.3 (2009), pp. 114–123. URL: <https://dl.acm.org/citation.cfm?id=2835589> (cit. on p. 100).
- [13] Ahmad Barghout, Jongeun Cha, Abdulmotaleb El Saddik, Julius Kammerl, and Eckehard Steinbach. “Spatial resolution of vibrotactile perception on the human forearm when exploiting funneling illusion”. In: *2009 IEEE International Workshop on Haptic Audio visual Environments and Games*. IEEE, Nov. 2009, pp. 19–23. ISBN: 978-1-4244-4217-1. DOI: 10.1109/HAVE.2009.5356122 (cit. on pp. 36, 38).
- [14] Kersten Behrens. “Combining a Neural Network with Scene Recognition to Provide Mobile Audio Guidance for the Blind”. Master thesis. Gottfried Wilhelm Leibniz University Hannover, 2020 (cit. on p. 5).
- [15] Christopher C. Berger and Mar Gonzalez-Franco. “Expanding the sense of touch outside the body”. In: *Proceedings - SAP 2018: ACM Symposium on Applied Perception* (2018). DOI: 10.1145/3225153.3225172 (cit. on p. 39).
- [16] Matthias Berning, Florian Braun, Till Riedel, and Michael Beigl. “ProximityHat: a head-worn system for subtle sensory augmentation with tactile stimulation”. In: *Proceedings of the 2015 ACM International Symposium on Wearable Computers - ISWC '15*. New York, New York, USA: ACM Press, 2015, pp. 31–38. ISBN: 9781450335782. DOI: 10.1145/2802083.2802088 (cit. on pp. 26, 30, 31, 36, 179).
- [17] BHaptics. *bHaptics Tactot*. 2020. URL: <https://www.bhaptics.com/tactsuit/> (visited on Apr. 11, 2021) (cit. on pp. 27, 88, 90, 91, 106).
- [18] Dominik Bial, Dagmar Kern, and F Alt. “Enhancing Outdoor Navigation Systems through Vibrotactile Feedback”. In: *Human factors in computing systems* (2011), pp. 1–6. DOI: 10.1145/1979742.1979760 (cit. on pp. 26, 28).
- [19] Frank Biocca, Arthur Tang, Charles Owen, and Fan Xiao. “Attention funnel: omnidirectional 3D cursor for mobile augmented reality platforms”. In: *Proceedings of the SIGCHI conference on Human Factors in computing systems* 06 (2006), pp. 1115–1122. DOI: 10.1145/1124772.1124939 (cit. on p. 120).
- [20] Aaron Bloomfield and Norman I Badler. “A Low Cost Tactor Suit for Vibrotactile Feedback”. In: *Technical Reports (CIS)* January (2003), p. 7 (cit. on pp. 26, 88, 90, 106).
- [21] Jonas Bock. “Implementierung und Evaluierung eines geräuschlosen am Kopf getragenen taktilen Ausgabegeräts”. Master thesis. Gottfried Wilhelm Leibniz University Hannover, 2018 (cit. on p. 5).
- [22] John Brooke. “SUS : A Retrospective”. In: *Journal of Usability Studies* 8.2 (2013), pp. 29–40. URL: <https://uxpajournal.org/sus-a-retrospective/> (cit. on pp. 99, 100).
- [23] A. G. Brown and A. Iggo. “A quantitative study of cutaneous receptors and afferent fibres in the cat and rabbit”. In: *The Journal of Physiology* 193.3 (Dec. 1967), pp. 707–733. ISSN: 00223751. DOI: 10.1113/jphysiol.1967.sp008390 (cit. on p. 18).

- [35] Victor Adriel De Jesus Oliveira, Luciana Nedel, Anderson Maciel, and Luca Brayda. “Spatial discrimination of vibrotactile stimuli around the head”. In: *IEEE Haptics Symposium, HAPTICS 2016-April* (2016), pp. 1–6. ISSN: 23247355. DOI: 10.1109/HAPTICS.2016.7463147 (cit. on pp. 36, 51, 52).
- [36] Victor Adriel de Jesus Oliveira, Luciana Nedel, Anderson Maciel, and Luca Brayda. “Anti-veering Vibrotactile HMD for Assistance of Blind Pedestrians”. In: *Proc. EuroHaptics 2018*. Ed. by Domenico Prattichizzo, Hiroyuki Shinoda, Hong Z. Tan, Emanuele Ruffaldi, and Antonio Frisoli. Vol. 10894. Lecture Notes in Computer Science June. Cham: Springer International Publishing, 2018, pp. 500–512. ISBN: 978-3-319-93398-6. DOI: 10.1007/978-3-319-93399-3_43 (cit. on p. 172).
- [37] Michael F. Deering. “HoloSketch: A Virtual Reality Sketching / Animation Tool”. In: *ACM Transactions on Computer-Human Interaction (TOCHI) 2.3* (1995), pp. 220–238. ISSN: 15577325. DOI: 10.1145/210079.210087 (cit. on p. 93).
- [38] Alexandra Delazio, Ken Nakagaki, Jill Fain Lehman, et al. “Force jacket: Pneumatically-actuated jacket for embodied haptic experiences”. In: *Conference on Human Factors in Computing Systems - Proceedings 2018-April* (2018), pp. 1–12. DOI: 10.1145/3173574.3173894 (cit. on pp. 27, 88, 90, 106).
- [39] Kerem Can Demir. “Auswirkungen der Funneling Illusion auf die taktile Zielführung auf dem Kopf”. Bachelor thesis. Gottfried Wilhelm Leibniz University Hannover, 2019 (cit. on pp. 5, 35).
- [40] Deutsches Institut für Normung. *DIN 18065*. 2015. (Visited on May 4, 2018) (cit. on p. 148).
- [41] Vincent Diener, Michael Beigl, Matthias Budde, and Erik Pescara. “VibrationCap: Studying vibrotactile localization on the human head with an unobtrusive wearable tactile display”. In: *Proceedings - International Symposium on Wearable Computers, ISWC Part F1305* (2017), pp. 82–89. ISSN: 15504816. DOI: 10.1145/3123021.3123047 (cit. on pp. 30, 31, 51, 52, 55, 71, 76, 172, 180).
- [42] H.Q. Dinh, N. Walker, L.F. Hodges, Chang Song, and A. Kobayashi. “Evaluating the importance of multi-sensory input on memory and the sense of presence in virtual environments”. In: *Proceedings IEEE Virtual Reality (Cat. No. 99CB36316)*. IEEE Comput. Soc, 1999, pp. 222–228. ISBN: 0-7695-0093-5. DOI: 10.1109/VR.1999.756955 (cit. on pp. 109, 110).
- [43] Michal Karol Dobrzynski, Seifeddine Mejri, Steffen Wischmann, and Dario Floreano. “Quantifying Information Transfer Through a Head-Attached Vibrotactile Display: Principles for Design and Control”. In: *IEEE Transactions on Biomedical Engineering* 59.7 (July 2012), pp. 2011–2018. ISSN: 0018-9294. DOI: 10.1109/TBME.2012.2196433 (cit. on pp. 20, 26, 30, 31, 36, 56, 147).
- [44] Andreas Domin. “Entwicklung einer Virtual-Reality-Anwendung für den Entwurf vibrotaktile Muster für ein am Kopf getragenes Ausgabegerät”. Bachelor thesis. Gottfried Wilhelm Leibniz University Hannover, 2019 (cit. on pp. 5, 87).

- [45] Andreas Dünser, Karin Steinbügl, Hannes Kaufmann, and Judith Glück. “Virtual and augmented reality as spatial ability training tools”. In: *ACM International Conference Proceeding Series* 158 (2006), pp. 125–132. DOI: 10.1145/1152760.1152776 (cit. on p. 108).
- [46] M. J. Enriquez and K. E. MacLean. “The hapticon editor: A tool in support of haptic communication research”. In: *Proceedings - 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, HAPTICS 2003*. IEEE Comput. Soc, 2003, pp. 356–362. ISBN: 0769518907. DOI: 10.1109/HAPTIC.2003.1191310 (cit. on pp. 88, 91).
- [47] Jan B F Van Erp and Hendrik A H C Van Veen. “A Multi-purpose Tactile Vest for Astronauts in the International Space Station”. In: *Proceedings of eurohaptics* (2003), pp. 405–408 (cit. on pp. 27, 88, 90, 106).
- [48] Jan B. F. Van Erp, Hendrik a. H. C. Van Veen, Chris Jansen, and Trevor Dobbins. “Waypoint navigation with a vibrotactile waist belt”. In: *ACM Transactions on Applied Perception* 2.2 (2005), pp. 106–117. DOI: 10.1145/1060581.1060585 (cit. on p. 26).
- [49] Jan B.F. van Erp. “Tactile navigation display”. In: *Haptic HCI 2000*. 2001, pp. 165–173. ISBN: 978-3-540-42356-0. DOI: 10.1007/3-540-44589-7_18 (cit. on pp. 26, 29).
- [50] Jan B.F. van Erp, Liselotte C.M. Kroon, Tina Mioch, and Katja I. Paul. “Obstacle detection display for visually impaired: Coding of direction, distance, and height on a vibrotactile waist band”. In: *Frontiers in ICT* 4.SEP (2017), pp. 1–19. ISSN: 2297198X. DOI: 10.3389/fict.2017.00023 (cit. on p. 141).
- [51] Jonas Fabisiak. “Extraktion immersiver Effekte aus Filmen und ihre taktile Darstellung auf dem Kopf”. Bachelor thesis. Gottfried Wilhelm Leibniz University Hannover, 2019 (cit. on p. 5).
- [52] German Flores, Sri Kurniawan, Roberto Manduchi, et al. “Vibrotactile Guidance for Wayfinding of Blind Walkers”. In: *IEEE Transactions on Haptics* 8.3 (July 2015), pp. 306–317. ISSN: 1939-1412. DOI: 10.1109/TOH.2015.2409980 (cit. on pp. 8, 26, 139, 140, 142, 156, 171, 172).
- [53] Guido Gardlo. “Real-Time Navigation of Visually Impaired Users via a Self-Positioning Device and Vibrotactile Feedback”. Master thesis. Gottfried Wilhelm Leibniz University Hannover, 2017 (cit. on pp. 5, 137).
- [54] Frank A. Geldard. “Adventures in tactile literacy.” In: *American Psychologist* 12.3 (1957), pp. 115–124. ISSN: 0003-066X. DOI: 10.1037/h0040416 (cit. on pp. 1, 26).
- [55] Michele Geronazzo, Alberto Bedin, Luca Brayda, Claudio Campus, and Federico Avanzini. *Geronazzo et al. - 2016 - Interactive spatial sonification for non-visual exploration of virtual maps.pdf*. 2016. DOI: 10.1016/j.ijhcs.2015.08.004 (cit. on p. 127).
- [56] G. A. Gescheider. “Evidence in support of the duplex theory of mechanoreception”. In: *SENS.PROCESSES* 1.1 (June 1976), pp. 68–76. ISSN: 0363-3799. URL: <https://europepmc.org/article/med/1029078> (cit. on p. 16).

- [57] Kirby Gilliland and Robert E. Schlegel. “Tactile stimulation of the human head for information display”. In: *Human Factors* 36.4 (1994), pp. 700–717. ISSN: 00187208. DOI: 10.1177/001872089403600410 (cit. on pp. 2, 36, 56, 118).
- [58] Google. *Google ARCore*. 2020. URL: <https://developers.google.com/ar> (visited on July 2, 2020) (cit. on pp. 138, 167, 181).
- [59] Google LLC. *Tilt Brush by Google*. 2016. URL: <https://www.tiltbrush.com/> (visited on Aug. 23, 2019) (cit. on p. 94).
- [60] Green-Armytage and Paul. “A Colour Alphabet and the Limits of Colour Coding”. In: *JAIC - Journal of the International Colour Association* 5.0 (2010). ISSN: 22271309 (cit. on p. 98).
- [61] Guiding Eyes for the Blind. *How many people use guide dogs?* 2020. URL: <https://www.guidingeyes.org/about/faqs/> (visited on Sept. 17, 2020) (cit. on p. 138).
- [62] Sebastian Günther, Jan Kirchner, Florian Müller, et al. “TactileGlove”. In: *Proceedings of the 11th Pervasive Technologies Related to Assistive Environments Conference on - PETRA '18* (2018), pp. 273–280. DOI: 10.1145/3197768.3197785 (cit. on pp. 26, 28).
- [63] Shameem Hameed, Thomas Ferris, Swapnaa Jayaraman, and Nadine Sarter. “Supporting Interruption Management Through Informative Tactile and Peripheral Visual Cues”. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50.3 (Oct. 2006), pp. 376–380. ISSN: 1541-9312. DOI: 10.1177/154193120605000335 (cit. on pp. 72, 78, 84, 85).
- [64] Leonard Hansing. “Entwicklung einer Benutzerschnittstelle für den Entwurf vibrotaktiler Muster für ein am Kopf getragenes Ausgabegerät”. Bachelor thesis. Gottfried Wilhelm Leibniz University Hannover, 2018 (cit. on pp. 5, 87).
- [65] Moritz Niklas Hasemann. “A Survey of Recent Advances in Tactile Feedback”. Bachelor thesis. Gottfried Wilhelm Leibniz University Hannover, 2019 (cit. on p. 5).
- [66] Marc Hassenzahl, Michael Burmester, and Franz Koller. “AttrakDiff: Ein Fragebogen zur Messung wahrgenommener hedonischer und pragmatischer Qualität”. In: Vieweg+Teubner Verlag, 2003, pp. 187–196. DOI: 10.1007/978-3-322-80058-9_19 (cit. on p. 99).
- [67] Vincent Hayward. “A brief taxonomy of tactile illusions and demonstrations that can be done in a hardware store”. In: *Brain Research Bulletin* 75.6 (Apr. 2008), pp. 742–752. ISSN: 03619230. DOI: 10.1016/j.brainresbull.2008.01.008 (cit. on pp. 32, 33).
- [68] Walter Hendelman M.D. *Atlas of Functional Neuroanatomy*. CRC Press, June 2015. ISBN: 9780429169021. DOI: 10.1201/b18595 (cit. on pp. 9–11).
- [69] Ernst Friedrich Gustav Herbst. *Die Pacinischen Körper und ihre Bedeutung: Ein Beitrag zur Kenntniß der Nervenprimitivfasern*. Göttingen: Vandenhoeck & Ruprecht, 1848 (cit. on p. 17).
- [70] Marion Hersh. “Cane use and late onset visual impairment”. In: *Technology and Disability* 27.3 (2015), pp. 103–116. ISSN: 1878643X. DOI: 10.3233/TAD-150432 (cit. on p. 138).

- [71] Wilko Heuten, Niels Henze, Susanne Boll, and Martin Pielot. “Tactile wayfinder: A Non-Visual Support System for Wayfindin”. In: *Proceedings of the 5th Nordic conference on Human-computer interaction building bridges - NordiCHI '08*. New York, New York, USA: ACM Press, 2008, p. 172. ISBN: 9781595937049. DOI: 10.1145/1463160.1463179 (cit. on pp. 26, 29, 142, 171).
- [72] K. A. Holbrook and G. F. Odland. “Regional differences in the thickness (cell layers) of the human stratum corneum: an ultrastructural analysis”. In: *Journal of Investigative Dermatology* 62.4 (Apr. 1974), pp. 415–422. ISSN: 0022202X. DOI: 10.1111/1523-1747.ep12701670 (cit. on p. 13).
- [73] Mark Hollins and S. Ryan Risner. “Evidence for the duplex theory of tactile texture perception”. In: *Perception & Psychophysics* 62.4 (Jan. 2000), pp. 695–705. ISSN: 0031-5117. DOI: 10.3758/BF03206916 (cit. on p. 16).
- [74] Kyungpyo Hong, Jaebong Lee, and Seungmoon Choi. “Demonstration-based vibrotactile pattern authoring”. In: *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction - TEI '13*. New York, New York, USA: ACM Press, 2013, p. 219. ISBN: 9781450318983. DOI: 10.1145/2460625.2460660 (cit. on pp. 88, 91).
- [75] Weijian Hu, Kaiwei Wang, Kailun Yang, et al. “A comparative study in real-time scene sonification for visually impaired people”. In: *Sensors (Switzerland)* 20.11 (2020), pp. 1–17. ISSN: 14248220. DOI: 10.3390/s20113222 (cit. on p. 141).
- [76] Da Yuan Huang, Liwei Chan, Xiao Feng Jian, et al. “VibroPlay: Aorind three-dimensional spatial-temporal tactile effects with direct manipulation”. In: *SA 2016 - SIGGRAPH ASIA 2016 Emerging Technologies* (2016). DOI: 10.1145/2988240.2988250 (cit. on p. 92).
- [77] Gi-Hun Yang, Ki-Uk Kyung, M.A. Srinivasan, and Dong-Soo Kwon. “Quantitative tactile display device with pin-array type tactile feedback and thermal feedback”. In: *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. Vol. 2006. IEEE, 2006, pp. 3917–3922. ISBN: 0-7803-9505-0. DOI: 10.1109/ROBOT.2006.1642302 (cit. on p. 21).
- [78] S. Ino, S. Shimizu, T. Odagawa, et al. “A tactile display for presenting quality of materials by changing the temperature of skin surface”. In: *Proceedings of 1993 2nd IEEE International Workshop on Robot and Human Communication*. IEEE, 1993, pp. 220–224. ISBN: 0-7803-1407-7. DOI: 10.1109/ROMAN.1993.367718 (cit. on p. 21).
- [79] Ali Israr and Ivan Poupyrev. “Control space of apparent haptic motion”. In: *2011 IEEE World Haptics Conference, WHC 2011* (2011), pp. 457–462. DOI: 10.1109/WHC.2011.5945529 (cit. on pp. 26, 38).
- [80] Ali Israr and Ivan Poupyrev. “Tactile brush”. In: *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*. New York, New York, USA: ACM Press, May 2011, p. 2019. ISBN: 9781450302289. DOI: 10.1145/1978942.1979235 (cit. on pp. 26, 33, 38, 56, 92).
- [81] Lynette A. Jones, Brett Lockyer, and Erin Piatieski. “Tactile display and vibrotactile pattern recognition on the torso”. In: *Advanced Robotics* 20.12 (Jan. 2006), pp. 1359–1374. ISSN: 0169-1864. DOI: 10.1163/156855306778960563 (cit. on p. 26).

- [82] Lynette A. Jones and Nadine B. Sarter. “Tactile displays: Guidance for their design and application”. In: *Human Factors* 50.1 (Feb. 2008), pp. 90–111. ISSN: 00187208. DOI: 10.1518/001872008X250638 (cit. on pp. 26, 27, 31, 70–72, 78, 83–86, 172).
- [83] H Kajimoto, Y Kanno, and S Tachi. “Forehead electro-tactile display for vision substitution”. In: *Proc EuroHaptics*. 2006, p. 11 (cit. on pp. 22, 59, 142).
- [84] Hiroyuki Kajimoto. “Electrotactile Display with Real-Time Impedance Feedback Using Pulse Width Modulation”. In: *IEEE Transactions on Haptics* 5.2 (Apr. 2012), pp. 184–188. ISSN: 1939-1412. DOI: 10.1109/TOH.2011.39 (cit. on pp. 21, 59).
- [85] Hiroyuki Kajimoto. “Skeletouch: transparent electro-tactile display for mobile surfaces”. In: *SIGGRAPH Asia 2012 Emerging Technologies*. New York, NY, USA: ACM, Nov. 2012, pp. 1–3. ISBN: 9781450319126. DOI: 10.1145/2407707.2407728 (cit. on p. 22).
- [86] Hiroyuki Kajimoto, Masaki Suzuki, and Yonezo Kanno. “HamsaTouch”. In: *Proceedings of the extended abstracts of the 32nd annual ACM conference on Human factors in computing systems - CHI EA '14* (Apr. 2014), pp. 1273–1278. DOI: 10.1145/2559206.2581164 (cit. on p. 22).
- [87] Jeonggoo Kang, Junghwan Kook, Kwangsu Cho, Semyung Wang, and Jeha Ryu. “Effects of amplitude modulation on vibrotactile flow displays on piezo-actuated thin touch screen”. In: *International Journal of Control, Automation and Systems* 10.3 (2012), pp. 582–588. ISSN: 15986446. DOI: 10.1007/s12555-012-0315-7 (cit. on p. 39).
- [88] Jeonggoo Kang, Jongsuh Lee, Heewon Kim, et al. “Smooth vibrotactile flow generation using two piezoelectric actuators”. In: *IEEE Transactions on Haptics* 5.1 (2012), pp. 21–32. ISSN: 19391412. DOI: 10.1109/TOH.2012.1 (cit. on p. 39).
- [89] Brian F. G. Katz, Slim Kammoun, Gaëtan Parseihian, et al. “NAVIG: augmented reality guidance system for the visually impaired”. In: *Virtual Reality* 16.4 (Nov. 2012), pp. 253–269. ISSN: 1359-4338. DOI: 10.1007/s10055-012-0213-6 (cit. on p. 138).
- [90] David Katz. *Der Aufbau der Tastwelt*. Leipzig: Verlag von Johann Ambrosius Barth, 1925, p. 131 (cit. on p. 16).
- [91] O.B. Kaul and M. Rohs. “Wearable head-mounted 3D tactile display application scenarios”. In: *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct, MobileHCI 2016*. 2016. ISBN: 9781450344135. DOI: 10.1145/2957265.2965022 (cit. on pp. 178, 227).
- [92] Oliver Beren Kaul, Kersten Behrens, and Michael Rohs. “Mobile Recognition and Tracking of Objects in the Environment through Augmented Reality and 3D Audio Cues for People with Visual Impairments”. In: *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2021. DOI: 10.1145/3411763.3451611 (cit. on p. 227).
- [93] Oliver Beren Kaul, Andreas Domin, Michael Rohs, et al. “VRTactileDraw: A Virtual Reality Tactile Pattern Designer for Complex Spatial Arrangements of Actuators”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (2021) (cit. on pp. 5, 87, 178, 226).

- [94] Oliver Beren Kaul, Leonard Hansing, and Michael Rohs. “3DTactileDraw: A Tactile Pattern Design Interface for Complex Arrangements of Actuators”. In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2019, pp. 1–6. ISBN: 9781450359719. DOI: 10.1145/3290607.3313030 (cit. on pp. 87, 89, 91, 92, 94, 95, 178, 227).
- [95] Oliver Beren Kaul, Kevin Meier, and Michael Rohs. “Increasing presence in virtual reality with a vibrotactile grid around the head”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 10516 LNCS (2017). Ed. by Regina Bernhaupt, Girish Dalvi, Anirudha Joshi, et al., pp. 289–298. ISSN: 16113349. DOI: 10.1007/978-3-319-68059-0_19 (cit. on pp. 5, 109, 178, 226).
- [96] Oliver Beren Kaul, Max Pfeiffer, and Michael Rohs. “Follow the Force: Steering the Index Finger towards Targets using EMS”. In: *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2016, pp. 2526–2532. ISBN: 9781450340823. DOI: 10.1145/2851581.2892352 (cit. on pp. 62, 227).
- [97] Oliver Beren Kaul and Michael Rohs. “Concept for Navigating the Visually Impaired using a Tactile Interface around the Head”. In: *Proceedings of the CHI 2019 Workshop on Hacking Blind Navigation*. 2019. URL: <https://blindnavigationchi19.wordpress.com/> (cit. on pp. 178, 227).
- [98] Oliver Beren Kaul and Michael Rohs. “HapticHead: 3D Guidance and Target Acquisition through a Vibrotactile Grid”. In: *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2016, pp. 2533–2539. ISBN: 9781450340823. DOI: 10.1145/2851581.2892355 (cit. on pp. 4, 5, 57, 61, 62, 92, 117, 118, 178, 227).
- [99] Oliver Beren Kaul and Michael Rohs. “HapticHead: A Spherical Vibrotactile Grid around the Head for 3D Guidance in Virtual and Augmented Reality”. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. New York, New York, USA: ACM Press, 2017, pp. 3729–3740. ISBN: 9781450346559. DOI: 10.1145/3025453.3025684 (cit. on pp. 4, 5, 26, 36, 56, 57, 61, 88, 90, 92, 95, 106, 117, 178, 226).
- [100] Oliver Beren Kaul and Michael Rohs. “Requirements of Navigation Support Systems for People with Visual Impairments”. In: *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. New York, NY, USA: ACM, Oct. 2018, pp. 680–683. ISBN: 9781450359665. DOI: 10.1145/3267305.3267685 (cit. on pp. 178, 227).
- [101] Oliver Beren Kaul, Michael Rohs, and Marc Mogalle. “Design and Evaluation of On-the-Head Spatial Tactile Patterns”. In: *19th International Conference on Mobile and Ubiquitous Multimedia*. New York, NY, USA: ACM, Nov. 2020, pp. 229–239. ISBN: 9781450388702. DOI: 10.1145/3428361.3428407 (cit. on pp. 4, 69, 139, 142, 178, 226).

- [102] Oliver Beren Kaul, Michael Rohs, Marc Mogalle, and Benjamin Simon. “Around-the-Head Spatial Tactile System for Supporting Micro Navigation of People with Visual Impairments”. In: *ACM Transactions on Computer-Human Interaction* 28.4 (Aug. 2021), p. 34 (cit. on pp. 5, 137, 178, 226).
- [103] Oliver Beren Kaul, Michael Rohs, Benjamin Simon, Kerem Can Demir, and Kamillo Ferry. “Vibrotactile Funneling Illusion and Localization Performance on the Head”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, Apr. 2020, pp. 1–13. ISBN: 9781450367080. DOI: 10.1145/3313831.3376335 (cit. on pp. 4, 35, 62, 178, 226).
- [104] Hamideh Kerdegari, Yeongmi Kim, and Tony J. Prescott. “Head-Mounted Sensory Augmentation Device: Comparing Haptic and Audio Modality”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 9793. Springer International Publishing, 2016, pp. 107–118. ISBN: 9783319424163. DOI: 10.1007/978-3-319-42417-0_11 (cit. on pp. 26, 30, 31, 36, 56).
- [105] Hamideh Kerdegari, Yeongmi Kim, Tom Stafford, and Tony J. Prescott. “Centralizing bias and the vibrotactile funneling illusion on the forehead”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 8619. Springer Verlag, 2014, pp. 55–62 (cit. on pp. 4, 26, 36, 37, 39–41, 43, 46, 47, 51, 53, 54, 56, 61, 62, 175).
- [106] M Kim, A Abdulali, and S Jeon. “Rendering Vibrotactile Flow on Backside of the Head: Initial Study”. In: *2018 IEEE Games, Entertainment, Media Conference (GEM)*. Aug. 2018, pp. 1–250. DOI: 10.1109/GEM.2018.8516545 (cit. on pp. 36, 56).
- [107] Sang Youn Kim, Jae Oh Kim, and Kyu Yong Kim. “Traveling vibrotactile wave - A new vibrotactile rendering method for mobile devices”. In: *IEEE Transactions on Consumer Electronics* 55.3 (2009), pp. 1032–1038. ISSN: 00983063. DOI: 10.1109/TCE.2009.5277952 (cit. on p. 39).
- [108] Sang Youn Kim and Jeong Cheol Kim. “Vibrotactile rendering for a traveling vibrotactile wave based on a haptic processor”. In: *IEEE Transactions on Haptics* 5.1 (2012), pp. 14–20. ISSN: 19391412. DOI: 10.1109/TOH.2011.72 (cit. on p. 39).
- [109] Youngsun Kim, Jaedong Lee, and Gerard J. Kim. “Designing of 2D illusory tactile feedback for hand-held tablets”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 9299 (2015), pp. 10–17. ISSN: 16113349. DOI: 10.1007/978-3-319-22723-8_2 (cit. on p. 39).
- [110] Youngsun Kim, Jaedong Lee, and Gerard J. Kim. “Extending “out of the body” tactile phantom sensations to 2D and applying it to mobile interaction”. In: *Personal and Ubiquitous Computing* 19.8 (2015), pp. 1295–1311. ISSN: 16174909. DOI: 10.1007/s00779-015-0894-4 (cit. on p. 39).
- [111] Youngsun Kim, Jaedong Lee, and Gerard Jounghyun Kim. “Design and application of 2D illusory vibrotactile feedback for hand-held tablets”. In: *Journal on Multimodal User Interfaces* 11.2 (2017), pp. 133–148. ISSN: 17838738. DOI: 10.1007/s12193-016-0234-7 (cit. on p. 39).

- [112] Florian Klein and Stephan Werner. “Auditory Adaptation to Non-Individual HRTF Cues in Binaural Audio Reproduction”. In: *Journal of the Audio Engineering Society* 64.1/2 (Feb. 2016), pp. 45–54. ISSN: 15494950. DOI: 10.17743/jaes.2015.0092 (cit. on p. 127).
- [113] Michael Koch, Kai Von Luck, Jan Schwarzer, and Susanne Draheim. “The Novelty Effect in Large Display Deployments-Experiences and Lessons-Learned for Evaluating Prototypes”. In: *ECSCW 2018 - Proceedings of the 16th European Conference on Computer Supported Cooperative Work*. European Society for Socially Embedded Technologies (EUSSET), 2018. URL: <https://dl.eusset.eu/handle/20.500.12015/3115> (cit. on p. 106).
- [114] Yukari Konishi, Nobuhisa Hanamitsu, Benjamin Outram, et al. “Synesthesia Suit”. In: Springer, Singapore, Nov. 2018, pp. 499–503. DOI: 10.1007/978-981-10-4157-0_84 (cit. on pp. 27, 88, 90, 106).
- [115] Bastian Krefeld. “Improving VR Presence using Head-Based Vibrotactile Feedback and Hand Tracking”. Master thesis. Gottfried Wilhelm Leibniz University Hannover, 2017 (cit. on p. 5).
- [116] Lester E. Krueger. “David Katz’s Der Aufbau der Tastwelt (The world of touch): A synopsis”. In: *Perception & Psychophysics* 7.6 (Nov. 1970), pp. 337–341. ISSN: 0031-5117. DOI: 10.3758/BF03208659 (cit. on p. 16).
- [117] J.P. LACOUR, D. DUBOIS, A. PISANI, and J.P. ORTONNE. “Anatomical mapping of Merkel cells in normal human adult epidermis”. In: *British Journal of Dermatology* 125.6 (Dec. 1991), pp. 535–542. ISSN: 0007-0963. DOI: 10.1111/j.1365-2133.1991.tb14790.x (cit. on p. 17).
- [118] Joseph J. LaViola and Joseph J. “A discussion of cybersickness in virtual environments”. In: *ACM SIGCHI Bulletin* 32.1 (Jan. 2000), pp. 47–56. ISSN: 07366906. DOI: 10.1145/333329.333344 (cit. on p. 93).
- [119] Jaebong Lee, Jonghyun Ryu, and Seungmoon Choi. “Vibrotactile score: A score metaphor for designing vibrotactile patterns”. In: *Proceedings - 3rd Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, World Haptics 2009* March 2009 (Mar. 2009), pp. 302–307. ISSN: 1098-6596. DOI: 10.1109/WHC.2009.4810816 (cit. on p. 91).
- [120] Jaedong Lee, Youngsun Kim, and Gerard Kim. “Funneling and saltation effects for tactile interaction with virtual objects”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '12* (2012), pp. 3141–3148. DOI: 10.1145/2207676.2208729 (cit. on p. 39).
- [121] Jaedong Lee, Youngsun Kim, and Gerard J. Kim. “Applying “Out of Body” Vibrotactile illusion to Two-Finger interaction for perception of object Dynamics”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 9299 (2015), pp. 506–509. ISSN: 16113349. DOI: 10.1007/978-3-319-22723-8_49 (cit. on p. 39).

- [122] Ville Lehtinen, Antti Oulasvirta, Antti Salovaara, and Petteri Nurmi. “Dynamic tactile guidance for visual search tasks”. In: *Proceedings of the 25th annual ACM symposium on User interface software and technology - UIST '12*. New York, New York, USA: ACM Press, 2012, p. 445. ISBN: 9781450315807. DOI: 10.1145/2380116.2380173 (cit. on pp. 26, 28).
- [123] James Jeng-Weei Lin, Habib Abi-Rached, Do-Hoe Kim, Donald E. Parker, and Thomas A. Furness. “A “Natural” Independent Visual Background Reduced Simulator Sickness”. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 46.26 (Sept. 2002), pp. 2124–2128. ISSN: 1541-9312. DOI: 10.1177/154193120204602605 (cit. on p. 93).
- [124] Robert W. Lindeman, Robert Page, Yasuyuki Yanagida, and John L. Sibert. “Towards full-body haptic feedback: The design and deployment of a spatialized vibrotactile feedback system”. In: *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST* (2004), pp. 146–149 (cit. on pp. 27, 88, 90, 106).
- [125] Lofelt GmbH. *Elevating Haptic Technology with Lofelt Wave*. 2019. URL: <https://lofelt.com/white-paper> (cit. on pp. 22, 24, 27, 67, 68, 172, 173).
- [126] Jack M Loomis, Reginald G Golledge, Roberta L Klatzky, James R Marston, and Gary L Ed Allen. “Assisting wayfinding in visually impaired travelers”. In: *Applied spatial Cognition From Research to Cognitive Technology* 1.60587 (2007), pp. 179–202 (cit. on p. 138).
- [127] Jack M. Loomis, Roberta L. Klatzky, Reginald G. Golledge, et al. “Nonvisual navigation by blind and sighted: Assessment of path integration ability.” In: *Journal of Experimental Psychology: General* 122.1 (1993), pp. 73–91. ISSN: 1939-2222. DOI: 10.1037/0096-3445.122.1.73 (cit. on p. 138).
- [128] Ethan Luckett. “A Quantitative Evaluation of the HTC Vive for Virtual Reality Research”. Bachelor thesis. University of Mississippi, 2018 (cit. on p. 163).
- [129] Alexander Marquardt, Christina Trepkowski, Tom David Eibich, Jens Maiero, and Ernst Kruijff. “Non-Visual Cues for View Management in Narrow Field of View Augmented Reality Displays”. In: *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* September (2019), pp. 190–201. DOI: 10.1109/ISMAR.2019.000-3 (cit. on pp. 26, 30, 31).
- [130] Manuel Martinez, Alina Roitberg, Daniel Koester, Rainer Stiefelhagen, and Boris Schauerte. “Using Technology Developed for Autonomous Cars to Help Navigate Blind People”. In: *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. Vol. 2018-Janua. IEEE, Oct. 2017, pp. 1424–1432. ISBN: 978-1-5386-1034-3. DOI: 10.1109/ICCVW.2017.169 (cit. on p. 138).
- [131] Manuel Martinez, Kailun Yang, Angela Constantinescu, and Rainer Stiefelhagen. “Helping the blind to get through covid-19: Social distancing assistant using real-time semantic segmentation on rgb-d video”. In: *Sensors (Switzerland)* 20.18 (2020), pp. 1–17. ISSN: 14248220. DOI: 10.3390/s20185202 (cit. on p. 138).

- [132] Akira Matsuda, Kazunori Nozawa, Kazuki Takata, Atsushi Izumihara, and Jun Rekimoto. “HapticPointer: A Neck-worn Device that Presents Direction by Vibrotactile Feedback for Remote Collaboration Tasks”. In: *Proceedings of the Augmented Humans International Conference*. New York, NY, USA: ACM, Mar. 2020, pp. 1–10. ISBN: 9781450376037. DOI: 10.1145/3384657.3384777 (cit. on pp. 26, 29, 142, 176).
- [133] Anita Meier, Denys J. C. Matthies, Bodo Urban, and Reto Wettach. “Exploring vibrotactile feedback on the body and foot for the purpose of pedestrian navigation”. In: *Proceedings of the 2nd international Workshop on Sensor-based Activity Recognition and Interaction - WOAR '15*. New York, New York, USA: ACM Press, June 2015, pp. 1–11. ISBN: 9781450334549. DOI: 10.1145/2790044.2790051 (cit. on pp. 27, 28).
- [134] Kevin Meier. “Evaluation of a Head-Mounted Tactile Display for Guidance and Immersion Applications”. Bachelor thesis. Gottfried Wilhelm Leibniz University Hannover, 2016 (cit. on pp. 5, 109, 117).
- [135] Georg Meissner. “Beiträge zur Anatomie und Physiologie der Haut”. In: (1853) (cit. on p. 17).
- [136] Renato Melo, Polyanna Amorim da Silva, Robson Souza, Maria Raposo, and Karla Ferraz. “Head Position Comparison between Students with Normal Hearing and Students with Sensorineural Hearing Loss”. In: *International Archives of Otorhinolaryngology* 17.04 (Sept. 2013), pp. 363–369. ISSN: 1809-9777. DOI: 10.1055/s-0033-1351685 (cit. on p. 166).
- [137] Fr Merkel. “Tastzellen und Tastkörperchen bei den Hausthieren und beim Menschen”. In: *Archiv für Mikroskopische Anatomie* 11.S1 (Dec. 1875), pp. 636–652. ISSN: 0176-7364. DOI: 10.1007/BF02933819 (cit. on p. 16).
- [138] M. Miyazaki, M. Hirashima, and D. Nozaki. “The "Cutaneous Rabbit" Hopping out of the Body”. In: *Journal of Neuroscience* 30.5 (2010), pp. 1856–1860. ISSN: 0270-6474. DOI: 10.1523/jneurosci.3887-09.2010 (cit. on p. 39).
- [139] Marc Mogalle. “Spatial Patterns for a Head-Worn Vibrotactile Display”. Master thesis. Gottfried Wilhelm Leibniz University Hannover, 2017 (cit. on pp. 5, 69, 137).
- [140] H. Møller and C. S. Pedersen. *Hearing at low and infrasonic frequencies*. Apr. 2004. URL: <https://www.noiseandhealth.org/text.asp?2004/6/23/37/31664> (cit. on pp. 59, 60).
- [141] Henrik Møller, Michael Friis Sørensen, Clemen Boje Jensen, and Dorte Hammershøi. “Binaural Technique: Do We Need Individual Recordings?” In: *Journal of the Audio Engineering Society* 44.6 (1996), pp. 451–469 (cit. on p. 127).
- [142] William Montagna. *The structure and function of skin*. New York: Academic Press Inc., 1974. ISBN: 0-12-505263-4 (cit. on p. 13).
- [143] K. Myles and J. T. Kalb. “Head Tactile Communication: Promising Technology With the Design of a Head-Mounted Tactile Display”. In: *Ergonomics in Design: The Quarterly of Human Factors Applications* 21.2 (Apr. 2013), pp. 4–8. ISSN: 1064-8046. DOI: 10.1177/1064804613477861 (cit. on pp. 3, 20, 31, 36, 41, 54–56).

- [144] Kimberly Myles and Joel T. Kalb. “Guidelines for Head Tactile Communication”. In: *Army Research Laboratory* 1.March (2010), p. 26 (cit. on pp. 3, 20, 31, 41, 61, 76, 145, 157).
- [145] Kimberly Myles and Joel T. Kalb. *Vibrotactile Sensitivity of the Head*. 2009 (cit. on pp. 2, 3, 20, 31, 41, 54, 76, 118).
- [146] Kimberly Myles, Joel T. Kalb, Janea Lowery, and Bheem P. Kattel. “The effect of hair density on the coupling between the tactor and the skin of the human head”. In: *Applied Ergonomics* 48 (2015), pp. 177–185. ISSN: 00036870. DOI: 10.1016/j.apergo.2014.11.007 (cit. on pp. 20, 31, 55, 60, 85, 172).
- [147] Miguel A. Nacenta, Yemliha Kamber, Yizhou Qiang, and Per Ola Kristensson. “Memorability of memorability of pre -designed & user -defined gesture sets”. In: *Conference on Human Factors in Computing Systems - Proceedings* (2013), pp. 1099–1108. DOI: 10.1145/2470654.2466142 (cit. on p. 85).
- [148] Lennart Nacke and Craig A. Lindley. “Flow and immersion in first-person shooters”. In: *Proceedings of the 2008 Conference on Future Play Research, Play, Share - Future Play '08* (2008), p. 81. DOI: 10.1145/1496984.1496998 (cit. on p. 111).
- [149] NaturalPoint Inc. *OptiTrack Tracking System*. 2018. URL: <https://optitrack.com/> (cit. on pp. 148, 156).
- [150] Tomi Nukarinen, Jussi Rantala, Ahmed Farooq, and Roope Raisamo. *Nukarinen et al. - 2015 - Delivering directional haptic cues through eyeglasses and a seat.pdf*. June 2015. DOI: 10.1109/WHC.2015.7177736 (cit. on pp. 26, 30, 36, 56, 138, 142, 171).
- [151] Qiangqiang Ouyang, Juan Wu, Zhiyu Shao, and Dapeng Chen. “A vibrotactile belt to display precise directional information for visually impaired”. In: *IEICE Electronics Express* 15.20 (2018), pp. 20180615–20180615. DOI: 10.1587/elex.15.20180615 (cit. on pp. 26, 29).
- [152] Sabrina Paneels, Margarita Anastassova, Steven Strachan, et al. “What’s around me? Multi-actuator haptic feedback on the wrist”. In: *2013 World Haptics Conference (WHC)*. IEEE, Apr. 2013, pp. 407–412. ISBN: 978-1-4799-0088-6. DOI: 10.1109/WHC.2013.6548443 (cit. on pp. 26, 28, 88, 91, 92, 142).
- [153] Gunhyuk Park and Seungmoon Choi. “Tactile Information Transmission by 2D Stationary Phantom Sensations”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*. 2018. ISBN: 9781450356206. DOI: 10.1145/3173574.3173832 (cit. on pp. 37, 38).
- [154] Jaeyoung Park, Jaeha Kim, Yonghwan Oh, and Hong Z. Tan. “Rendering moving tactile stroke on the palm using a sparse 2D array”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 9774 (2016), pp. 47–56. ISSN: 16113349. DOI: 10.1007/978-3-319-42321-0_5 (cit. on p. 38).
- [155] HAL Pashler and Steven Yantis. *Stevens' Handbook of Experimental Psychology*. Ed. by Hal Pashler. Hoboken, NJ, USA: John Wiley & Sons, Inc., July 2002. ISBN: 0471214426. DOI: 10.1002/0471214426 (cit. on pp. 9, 12, 15–19).

- [156] Randy Pausch, Dennis Proffitt, and George Williams. “Quantifying immersion in virtual reality”. In: *Proceedings of the 24th annual conference on Computer graphics and interactive techniques - SIGGRAPH '97* (1997), pp. 13–18. ISSN: 0097-8930. DOI: 10.1145/258734.258744 (cit. on p. 110).
- [157] Yi-Hao Peng, Carolyn Yu, Shi-Hong Liu, et al. “WalkingVibe: Reducing Virtual Reality Sickness and Improving Realism while Walking in VR using Unobtrusive Head-mounted Vibrotactile Feedback”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, Apr. 2020, pp. 1–12. ISBN: 9781450367080. DOI: 10.1145/3313831.3376847 (cit. on p. 179).
- [158] Martin Pielot, Oliver Krull, and Susanne Boll. “Where is my team? Supporting Situation Awareness with Tactile Displays”. In: *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*. New York, New York, USA: ACM Press, Apr. 2010, p. 1705. ISBN: 9781605589299. DOI: 10.1145/1753326.1753581 (cit. on p. 26).
- [159] Daniel Pietschmann and Peter Ohler. “Spatial Mapping of Physical and Virtual Spaces as an Extension of Natural Mapping: Relevance for Interaction Design and User Experience”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 9179. Springer Verlag, 2015, pp. 49–59. ISBN: 9783319210667. DOI: 10.1007/978-3-319-21067-4_6 (cit. on pp. 2, 4, 57, 118, 139, 142, 144).
- [160] Dario Pittera, Marianna Obrist, and Ali Israr. “Hand-to-hand: an intermanual illusion of movement”. In: *Proceedings of the 19th ACM International Conference on Multimodal Interaction* (2017), pp. 73–81. DOI: 10.1145/3136755.3136777 (cit. on pp. 26, 39).
- [161] Matteo Poggi and Stefano Mattocchia. “A wearable mobility aid for the visually impaired based on embedded 3D vision and deep learning”. In: *2016 IEEE Symposium on Computers and Communication (ISCC)*. Vol. 2016-Augus. IEEE, June 2016, pp. 208–213. ISBN: 978-1-5090-0679-3. DOI: 10.1109/ISCC.2016.7543741 (cit. on p. 141).
- [162] Precision Microdrives Limited. *Precision Microdrives Actuator 307-103*. 2016. URL: <https://www.precisionmicrodrives.com/product/307-103-9mm-vibration-motor-25mm-type> (visited on Mar. 28, 2021) (cit. on pp. 66, 67).
- [163] Precision Microdrives Limited. *Precision Microdrives Actuator 310-117*. 2017. URL: <https://www.precisionmicrodrives.com/product/310-117-10mm-vibration-motor-3mm-type> (visited on Mar. 28, 2021) (cit. on p. 63).
- [164] Precision Microdrives Limited. *Precision Microdrives Actuator 312-101*. 2017. URL: <https://www.precisionmicrodrives.com/product/312-101-12mm-vibration-motor-3mm-type> (visited on Mar. 28, 2021) (cit. on pp. 25, 59, 65, 66).
- [165] Tony Prescott, Ehud Ahissar, and Eugene Izhikevich. *Scholarpedia of Touch*. Ed. by Tony Prescott, Ehud Ahissar, and Eugene Izhikevich. Paris: Atlantis Press, 2016. ISBN: 978-94-6239-132-1. DOI: 10.2991/978-94-6239-133-8 (cit. on pp. 9, 32).

- [166] Jukka Raisamo, Roope Raisamo, and V. Surakka. “Comparison of Saltation, Amplitude Modulation, and a Hybrid Method of Vibrotactile Stimulation”. In: *IEEE Transactions on Haptics* 6.4 (2013), pp. 517–521. ISSN: 19391412. DOI: 10.1109/TOH.2013.25 (cit. on p. 38).
- [167] Jukka Raisamo, Roope Raisamo, and Veikko Surakka. “Evaluating the effect of temporal parameters for vibrotactile saltatory patterns”. In: (2009), p. 319. DOI: 10.1145/1647314.1647381 (cit. on p. 38).
- [168] Raspberry Pi Foundation. *Raspberry Pi 3 Model B*. 2016. URL: www.raspberrypi.org/ (visited on Apr. 11, 2021) (cit. on p. 66).
- [169] Charlotte M. Reed, Hong Z. Tan, Zachary D. Perez, et al. “A Phonemic-Based Tactile Display for Speech Communication”. In: *IEEE Transactions on Haptics* 12.1 (2019), pp. 2–17. ISSN: 23294051. DOI: 10.1109/TOH.2018.2861010 (cit. on p. 1).
- [170] Angelo Ruffini. “Di un nuovo organo nervoso terminale e sulla presenza dei corpuscoli Golgi-Mazzoni nel connettivo sottocutaneo dei polpastrelli delle dita dell’uomo”. PhD thesis. Royal College of Surgeons of England, 1894 (cit. on p. 17).
- [171] Dongseok Ryu, Gi Hun Yang, and Sungchul Kang. “T-hive : Vibrotactile Interface Presenting Spatial Information on Handle Surface”. In: *Proceedings - IEEE International Conference on Robotics and Automation* (2009), pp. 683–688. ISSN: 10504729. DOI: 10.1109/ROBOT.2009.5152740 (cit. on p. 38).
- [172] Jonghyun Ryu and Seungmoon Choi. “posVibEditor: Graphical authoring tool of vibrotactile patterns”. In: *2008 IEEE International Workshop on Haptic Audio visual Environments and Games*. IEEE, Oct. 2008, pp. 120–125. ISBN: 978-1-4244-2668-3. DOI: 10.1109/HAVE.2008.4685310 (cit. on pp. 88, 91).
- [173] Jonghyun Ryu and Gerard Jounghyun Kim. “Using a vibro-tactile display for enhanced collision perception and presence”. In: *Proceedings of the ACM symposium on Virtual reality software and technology - VRST '04*. New York, New York, USA: ACM Press, 2004, p. 89. ISBN: 1581139071. DOI: 10.1145/1077534.1077551 (cit. on p. 110).
- [174] Eva-Lotta Lotta Sallnäs, Kirsten Rasmus-Gröhn, and Calle Sjöström. “Supporting Presence in Collaborative Environments by Haptic Force Feedback”. In: *ACM Transactions on Computer-Human Interaction* 7.4 (Dec. 2000), pp. 461–476. ISSN: 15577325. DOI: 10.1145/365058.365086 (cit. on p. 109).
- [175] Maria V. Sanchez-Vives and Mel Slater. “Opinion: From presence to consciousness through virtual reality”. In: *Nature Reviews Neuroscience* 6.4 (Apr. 2005), pp. 332–339. ISSN: 1471-003X. DOI: 10.1038/nrn1651 (cit. on p. 109).
- [176] Timothy Sanders and Paul Cairns. *Time perception, immersion and music in videogames*. British Informatics Society, 2010, pp. 160–167. ISBN: 9781780171302. URL: <http://dl.acm.org/citation.cfm?id=2146327> (cit. on p. 113).
- [177] Jeff Sauro. *A practical guide to the system usability scale: Background, benchmarks & best practices*. Measuring Usability LLC, 2011 (cit. on p. 100).

- [178] Stefanie Schaack, George Chernyshov, Kirill Ragozin, et al. “Haptic Collar - vibrotactile feedback around the neck for guidance applications”. In: *Proceedings of the 10th Augmented Human International Conference 2019*. New York, NY, USA: ACM, Mar. 2019, pp. 1–4. ISBN: 9781450365475. DOI: 10.1145/3311823.3311840 (cit. on pp. 26, 29, 142).
- [179] Boris Schauerte, Daniel Koester, Manel Martinez, and Rainer Stiefelhagen. “Way to go! detecting open areas ahead of a walking person”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 8927 (2015), pp. 349–360. ISSN: 16113349. DOI: 10.1007/978-3-319-16199-0_25 (cit. on p. 138).
- [180] S. Scheggi, A. Talarico, and D. Prattichizzo. “A remote guidance system for blind and visually impaired people via vibrotactile haptic feedback”. In: *22nd Mediterranean Conference on Control and Automation*. IEEE, June 2014, pp. 20–23. ISBN: 978-1-4799-5901-3. DOI: 10.1109/MED.2014.6961320 (cit. on pp. 138, 142, 171).
- [181] Maximilian Schirmer, Johannes Hartmann, Sven Bertel, and Florian Echtler. “Shoe me the Way : A Shoe-Based Tactile Interface for Eyes-Free Urban Navigation”. In: *Proceedings of the 17th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI’15)* (2015), pp. 327–336 (cit. on p. 27).
- [182] Robert F. Schmidt and Hans-Georg Schaible. *Neuro- und Sinnesphysiologie*. 5th. Heidelberg: Springer Medizin Verlag Heidelberg, 2006. ISBN: 978-3-540-25700-4 (cit. on pp. 9, 12, 21).
- [183] Oliver S. Schneider, Ali Israr, and Karon E. MacLean. “Tactile Animation by Direct Manipulation of Grid Displays”. In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST ’15*. New York, New York, USA: ACM Press, 2015, pp. 21–30. ISBN: 9781450337793. DOI: 10.1145/2807442.2807470 (cit. on pp. 26, 92, 102, 105, 129).
- [184] Oliver S. Schneider and Karon E. MacLean. “Studying design process and example use with Macaron, a web-based vibrotactile effect editor”. In: *IEEE Haptics Symposium, HAPTICS 2016-April* (2016), pp. 52–58. ISSN: 23247355. DOI: 10.1109/HAPTICS.2016.7463155 (cit. on p. 91).
- [185] Oliver S. Schneider and Karon E. Maclean. “Improvising design with a Haptic Instrument”. In: *IEEE Haptics Symposium, HAPTICS* (2014), pp. 327–332. ISSN: 23247355. DOI: 10.1109/HAPTICS.2014.6775476 (cit. on p. 105).
- [186] Maximilian Schrapel. *InnovateFPGA | EMEA | EM032 - MultiWave: A Fully Customizable Mobile Function Generator For Haptic Feedback In VR*. 2019. URL: <http://www.innovatefpga.com/cgi-bin/innovate/teams.pl?Id=EM032> (visited on Feb. 19, 2021) (cit. on pp. 67, 68).
- [187] Maximilian Schrapel, Sergej Loewen, and Michael Rohs. *MultiWave: A Fully Customizable Mobile Function Generator For Haptic Feedback In VR*. Contest. 2019 (cit. on pp. 88, 90, 106, 107).

- [188] J J Scott and Robert Gray. “A comparison of tactile, visual, and auditory warnings for rear-end collision prevention in simulated driving.” In: *Human factors* 50.2 (2008), pp. 264–275. ISSN: 00187208. DOI: 10.1518/001872008X250674 (cit. on pp. 2, 118).
- [189] Hervé Segond, Déborah Weiss, and Eliana Sampaio. “Human Spatial Navigation via a Visuo-Tactile Sensory Substitution System”. In: *Perception* 34.10 (Oct. 2005), pp. 1231–1249. ISSN: 0301-0066. DOI: 10.1068/p3409 (cit. on p. 142).
- [190] Hasti Seifi, Chamila Anthonypillai, and Karon E. MacLean. “End-user customization of affective tactile messages: A qualitative examination of tool parameters”. In: *2014 IEEE Haptics Symposium (HAPTICS)*. IEEE, Feb. 2014, pp. 251–256. ISBN: 978-1-4799-3131-6. DOI: 10.1109/HAPTICS.2014.6775463 (cit. on pp. 88, 90).
- [191] Jongman Seo and Seungmoon Choi. “Edge flows: Improving information transmission in mobile devices using two-dimensional vibrotactile flows”. In: *IEEE World Haptics Conference, WHC 2015 (2015)*, pp. 25–30. DOI: 10.1109/WHC.2015.7177686 (cit. on p. 39).
- [192] Jongman Seo and Seungmoon Choi. “Initial study for creating linearly moving vibrotactile sensation on mobile device”. In: *2010 IEEE Haptics Symposium*. IEEE, Mar. 2010, pp. 67–70. ISBN: 978-1-4244-6821-8. DOI: 10.1109/HAPTIC.2010.5444677 (cit. on p. 129).
- [193] Jongman Seo and Seungmoon Choi. “Perceptual analysis of vibrotactile flows on a mobile device”. In: *IEEE Transactions on Haptics* 6.4 (2013), pp. 522–527. ISSN: 19391412. DOI: 10.1109/TOH.2013.24 (cit. on p. 39).
- [194] William R. Sherman and Alan B. Craig. “Understanding Virtual Reality—Interface, Application, and Design”. In: *Presence: Teleoperators and Virtual Environments* 12.4 (Aug. 2003), pp. 441–442. ISSN: 1054-7460. DOI: 10.1162/105474603322391668 (cit. on p. 93).
- [195] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, et al. “Large-scale assessment of mobile notifications”. In: *Conference on Human Factors in Computing Systems - Proceedings (2014)*, pp. 3055–3064. DOI: 10.1145/2556288.2557189 (cit. on p. 71).
- [196] Alexa F. Siu, Mike Sinclair, Robert Kovacs, et al. “Virtual Reality Without Vision: A Haptic and Auditory White Cane to Navigate Complex Virtual Worlds”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, Apr. 2020, pp. 1–13. ISBN: 9781450367080. DOI: 10.1145/3313831.3376353 (cit. on p. 141).
- [197] Mel Slater. “A Note on Presence Terminology”. In: *Emotion* 3 (2003), pp. 1–5. ISSN: 12938505 (cit. on p. 110).
- [198] N J A Sloane, R H Hardin, W D Smith, et al. *Spherical codes*. 2000. URL: <http://neilsloane.com/packings/> (cit. on p. 119).
- [199] Mark M. Smith. *Sensing the Past. Seeing, Hearing, Smelling, Tasting, and Touching in History*. University of California Press, 2007, p. 180. ISBN: 978-0-520-25495-4 (cit. on p. 1).

- [200] Katherine O. Sofia and Lynette Jones. “Mechanical and psychophysical studies of surface wave propagation during vibrotactile stimulation”. In: *IEEE Transactions on Haptics* 6.3 (2013), pp. 320–329. ISSN: 19391412. DOI: 10.1109/TOH.2013.1 (cit. on p. 39).
- [201] Blausen.com staff. “Medical gallery of Blausen Medical 2014”. In: *WikiJournal of Medicine* 1.2 (2014). DOI: 10.15347/wjm/2014.010 (cit. on pp. 14, 16).
- [202] H. Christiaan Stronks, Janine Walker, Daniel J. Parker, and Nick Barnes. “Training Improves Vibrotactile Spatial Acuity and Intensity Discrimination on the Lower Back Using Coin Motors”. In: *Artificial Organs* (2017). ISSN: 15251594. DOI: 10.1111/aor.12882 (cit. on p. 26).
- [203] Colin Swindells, Seppo Pietarinen, and Arto Viitanen. “Medium fidelity rapid prototyping of vibrotactile haptic, audio and video effects”. In: *IEEE Haptics Symposium, HAPTICS* (2014), pp. 515–521. ISSN: 23247355. DOI: 10.1109/HAPTICS.2014.6775509 (cit. on p. 91).
- [204] Hong Z. Tan, Charlotte M. Reed, Yang Jiao, et al. “Acquisition of 500 English Words through a TActile Phonemic Sleeve (TAPS)”. In: *IEEE Transactions on Haptics* 13.4 (2020), pp. 745–760. ISSN: 1939-1412. DOI: 10.1109/toh.2020.2973135 (cit. on p. 1).
- [205] Andreas Tarnowsky. “Modelle zur Synthese taktiler Reize”. PhD thesis. Gottfried Wilhelm Leibniz Universität Hannover, 2018. DOI: 10.15488/3737 (cit. on pp. 9–13, 16).
- [206] Koji Tsukada and Michiaki Yasumura. “ActiveBelt: Belt-Type Wearable Tactile Display for Directional Navigation”. In: *Ubiquitous Computing*. Vol. 3205. Berlin, Heidelberg: Springer, 2004, pp. 384–399. ISBN: 978-3-540-30119-6. DOI: 10.1007/978-3-540-30119-6_23 (cit. on pp. 1, 26, 28, 57, 142, 161, 177).
- [207] Raymond K. Tucker. “Reliability of semantic differential scales: The role of factor analysis”. In: *Western Speech* 35.3 (Dec. 1971), pp. 185–190. ISSN: 0043-4205. DOI: 10.1080/10570317109373702 (cit. on p. 111).
- [208] Unity Technologies. *Unity - Game Engine*. 2019. URL: <https://unity3d.com/> (visited on Sept. 12, 2016) (cit. on pp. 40, 44, 57, 61).
- [209] Unity Technologies. *Unity - Manual: VR overview*. 2019. URL: <https://docs.unity3d.com/Manual/VROverview.html> (visited on Aug. 13, 2019) (cit. on p. 93).
- [210] US Department of Transportation. “Lesson 9: Walkways, Sidewalks, and Public Spaces”. In: *Federal Highway Administration University Course on Bicycle and Pedestrian Transportation* 1.July (2006), p. 452 (cit. on p. 171).
- [211] A. B. Vallbo, H. Olausson, J. Wessberg, and N. Kakuda. “Receptive field characteristics of tactile units with myelinated afferents in hairy skin of human subjects.” In: *The Journal of Physiology* 483.3 (Mar. 1995), pp. 783–795. ISSN: 00223751. DOI: 10.1113/jphysiol.1995.sp020622 (cit. on p. 18).

- [212] Jan B.F. Van Erp and Hendrik A.H.C. Van Veen. “Vibrotactile in-vehicle navigation system”. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 7.4-5 (2004), pp. 247–256. ISSN: 13698478. DOI: 10.1016/j.trf.2004.09.003 (cit. on pp. 2, 118).
- [213] M W Van Someren, Y F Barnard, and J A C Sandberg. *The think aloud method: a practical approach to modelling cognitive processes*. 1994 (cit. on pp. 95, 96).
- [214] Abraham Vater. “De consensu partium corporis humani”. PhD thesis. 1741, pp. 953–973 (cit. on p. 17).
- [215] Ramiro Velazquez, Edwige Pissaloux, Carolina Del-Valle-Soto, Aime Lay-Ekuakille, and Bruno Ando. “Usability evaluation of foot-based interfaces for blind travelers”. In: *IEEE Instrumentation & Measurement Magazine* 23.4 (2020), pp. 4–13. ISSN: 1094-6969. DOI: 10.1109/mim.2020.9126045 (cit. on pp. 138, 142, 171).
- [216] Washington State. *Dispelling Myths, Department of services for the blind*. 2020. URL: <https://dsb.wa.gov/resources/blind-awareness/dispelling-myths> (visited on Sept. 17, 2020) (cit. on p. 138).
- [217] T.P. P. Way and K.E. E. Barner. “Automatic visual to tactile translation - Part I: Human factors, access methods, and image manipulation”. In: *IEEE Transactions on Rehabilitation Engineering* 5.1 (Mar. 1997), pp. 81–94. ISSN: 10636528. DOI: 10.1109/86.559353 (cit. on pp. 2, 118).
- [218] Bruce B. Blasch; William R. Wiener; Richard L. Welsh. *Foundations of Orientation and Mobility*. 2nd ed. AFB Press, 1997, p. 775. ISBN: 9780891289463 (cit. on p. 142).
- [219] White Cane Day. *White Cane Day FAQ*. 2020. URL: <http://whitecane.org/canes/> (visited on Sept. 17, 2020) (cit. on p. 138).
- [220] Bob G. Witmer and Michael J. Singer. “Measuring Presence in Virtual Environments: A Presence Questionnaire”. In: *Presence: Teleoperators and Virtual Environments* 7.3 (1998), pp. 225–240. DOI: 10.1162/105474698565686 (cit. on pp. 111, 113).
- [221] Flynn Wolf and Ravi Kuber. “Developing a head-mounted tactile prototype to support situational awareness”. In: *International Journal of Human Computer Studies* 109.July 2017 (2018), pp. 54–67. ISSN: 10959300. DOI: 10.1016/j.ijhcs.2017.08.002 (cit. on pp. 26, 30).
- [222] Zhen Ya-Xian, Takaki Suetake, and Hachiro Tagami. “Number of cell layers of the stratum corneum in normal skin - relationship to the anatomical location on the body, age, sex and physical parameters”. In: *Archives of Dermatological Research* 291.10 (Nov. 1999), pp. 555–559. ISSN: 0340-3696. DOI: 10.1007/s004030050453 (cit. on p. 13).
- [223] Lee Y. and Hwang K. “Skin thickness of Korean adults”. In: *Surgical and Radiologic Anatomy* 24.3-4 (Jan. 2002), pp. 183–189. ISSN: 0930-1038. DOI: 10.1007/s00276-002-0034-5 (cit. on p. 13).

- [224] Akio Yamamoto, Benjamin Cros, Hironori Hashimoto, and Toshiro Higuchi. “Control of thermal tactile display based on prediction of contact temperature”. In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*. Vol. 2004. 2. IEEE, 2004, 1536–1541 Vol.2. ISBN: 0-7803-8232-3. DOI: 10.1109/ROBOT.2004.1308042 (cit. on p. 21).
- [225] Gi Hun Yang, Moon Sub Jin, Yeonsub Jin, and Sungchul Kang. “T-mobile: Vibrotactile display pad with spatial and directional information for hand-held device”. In: *IEEE/RSJ 2010 International Conference on Intelligent Robots and Systems, IROS 2010 - Conference Proceedings (2010)*, pp. 5245–5250. DOI: 10.1109/IROS.2010.5651759 (cit. on p. 38).
- [226] Gi Hun Yang, Dongseok Ryu, and Sungchul Kang. “Vibrotactile display for hand-held input device providing spatial and directional information”. In: *Proceedings - 3rd Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, World Haptics 2009 (2009)*, pp. 79–84. DOI: 10.1109/WHC.2009.4810831 (cit. on p. 38).
- [227] Koji Yatani and KN Truong. “SemFeel: a user interface with semantic tactile feedback for mobile touch-screen devices”. In: *22nd annual ACM symposium on User interface (2009)*, pp. 111–120. URL: <http://portal.acm.org/citation.cfm?id=1622198> (cit. on p. 38).
- [228] Siyan Zhao, Ali Israr, and Roberta Klatzky. “Intermanual apparent tactile motion on handheld tablets”. In: *IEEE World Haptics Conference, WHC 2015 (2015)*, pp. 241–247. DOI: 10.1109/WHC.2015.7177720 (cit. on p. 39).
- [229] Yuhang Zhao, Cynthia L. Bennett, Hrvoje Benko, et al. “Demonstration of enabling people with visual impairments to navigate virtual reality with a haptic and auditory cane simulation”. In: *Conference on Human Factors in Computing Systems - Proceedings 2018-April (2018)*, pp. 1–14. DOI: 10.1145/3170427.3186485 (cit. on p. 141).
- [230] Karl Zilles and Gerd Rehkämper. *Funktionelle Neuroanatomie : Lehrbuch und Atlas*. 3rd. Springer Berlin Heidelberg, 1994, pp. 228–241. ISBN: 3-540-57855-2 (cit. on pp. 9–11).

Acronyms

AR – Augmented Reality
CNS – Central nervous system
DC – Direct current
DGPS – Differential Global Positioning System
ERM – Eccentric Rotating Mass
FI – Funneling Illusion
FPGA – Field-programmable Gate Array
GPS – Global Positioning System
HMD – head-mounted display
HRTF – Head-related Transfer Function
HTMR – High-Threshold Mechanoreceptor
HUD – Head-up Display
LiDAR – Light Detection and Ranging
LRA – Linear Resonant Actuator
LTMR – Low-Threshold Mechanoreceptor
ITQ – Immersion Tendency Questionnaire
PNS – Peripheral nervous system
PWM – Pulse-width Modulation
PQ – Presence Questionnaire
RQ – Research Question
SMA – Shape Memory Alloy
TN – Tactile Notification
TP – Tactile Pattern
USB – Universal Serial Bus
VIP – Visually Impaired People
VPZ – Virtual Point Zero
VR – Virtual Reality
VV – Viking Village

List of Figures

1.1	HapticHead prototype, side and front view. The blue markers indicate ear openings, and the red markers indicate the locations of 10 actuators that form a “ring” around the face. The prototype contains a total of 24 actuators on three concentric ellipses.	3
2.1	Fundamental structure of human nerve cells. Note that the structure changes based on the specialization of the cell. Some specialized cell elements may be altered or not present compared to this general structure (e.g., not all axons spawn axon collaterals). Figure source: own work, with background from [68, 205, 230]	10
2.2	Schematic curve of the action potential of a neuron in case of a successful firing (red curve) and a curve where the excitation threshold was not reached (blue curve). The successful firing can be split into four phases: (a) initiation phase, (b) depolarization, (c) repolarization, and (d) hyperpolarization. Figure adapted from [205, Fig. 2.3]	11
2.3	Mechanoreceptive afferents (tactile receptors) in the human skin. Adapted from [201]; enriched by additional information on hair afferents by [1]. Note that there are substantial differences between glabrous and hairy skin. This image shows afferents present in both glabrous and hairy skin. Not all of the shown hair afferents are present on all types of hair (see text).	14
2.4	Detail view of tactile receptors in the skin. From left to right: Merkel disk ($A\beta$ SA1-LTMR), Ruffini corpuscle ($A\beta$ SA2-LTMR), Meissner corpuscle ($A\beta$ RA1-LTMR), Pacinian corpuscle ($A\beta$ RA2-LTMR). Adapted from [201].	16
2.5	Overview of different actuator types. 1 - solenoid, 2a - generic voice coil actuator, 2b - bone conduction speaker, 3a linear resonant actuator, 3b - Lofelt actuator [125], 4 - 8 and 12 mm eccentric rotating mass actuators, 5 - two different sizes of piezoelectric actuators.	22
2.6	The tactile funneling illusion. In case two simultaneous vibrotactile stimuli are presented on the skin, the user may feel only a single stimulus somewhere between the stimuli locations if the two sources are close enough together.	32
3.1	Measuring the FI and centralizing bias on the head. A study participant showing two perceived actuator locations on the forehead.	36

3.2	Our prototype built after the template of Kerdegari et al. [105]. Seven actuators (A1-A7) are placed at scale positions 11 to 26 cm, with 2.5 cm distance between the centers. Tolerances are less than 1 mm.	40
3.5	The experiment UI was used by the experimenter to guide a participant through the study and record data. The two vertical lines can be moved by clicking and dragging them along the scale.	43
3.6	Heat map of mean absolute localization accuracies for all single actuator locations. From left to right: frontal, left, and back view of the model head. The data from symmetric head regions is merged to reduce noise. The color scale is viridis (a perceptually uniform color scale), ranging from 0.53 cm as the minimum error on the forehead to 1.78 cm on the headSides.	46
3.7	Localization error of each tactor on the forehead and chin, comparable to Fig. 4 right in [105]. Deviation towards the left is shown in blue, deviation towards the right in red. If the N for an actuator does not sum up to 24 (number of participants, median of 3 repetitions for each location), this means that the other trials were within ± 0.5 mm of the correct location. A1 is the left-most actuator, and A7 is the right-most actuator.	47
3.8	Occurrence frequency of the FI in two-actuator trials for different actuator distances at different head regions. Blue bars show trials in which a participant indicated a single location, red bars show trials in which a participant indicated two locations. Error bars represent the standard deviation between participants.	48
3.9	Actual and perceived distance between locations with multiple stimulation points. Error bars represent standard deviation between participants. In case a participant indicated only a single location when multiple stimulations were given, a distance of zero is assumed.	49
3.10	Average subjective head sensitivity rating. Participants were able to sort the head regions by sensitivity from score 6 (highest) to score 1 (lowest).	50
3.11	Diverging stacked bar chart of subjective results of our experiment, measured through the final questionnaire.	50
4.1	First HapticHead concept. Notice that all 20 actuators are located on one of three rings around the head and the actuator above the ear is > 4 cm away from the ear opening.	58
4.2	First HapticHead prototype	62
4.3	Revised HapticHead concept. Notice that all 24 actuators are located on one of three rings around the head, and the actuator above the ear is > 4 cm away from the ear opening. The forehead and chin regions contain a higher density of actuators intended to increase precision for guidance applications and enable more users to experience the funneling illusion, increasing comfort due to smooth transitions between actuators.	63

4.4	Second, refined HapticHead prototype, side and front view. Notice actuators located on the outside, the flexible chinstrap, and five instead of three actuators on each, the forehead and chinstrap. Positions of the 10 forehead and chin actuators forming a “ring” around the face are marked in red.	64
4.5	System overview of the HapticHead tactile around-the-head feedback system.	64
4.6	Third HapticHead prototype, hidden under a beanie for improved aesthetics.	65
4.7	HapticHead motor driver board v2. Developed in collaboration with Tim Diente.	65
4.8	Specialized, discontinued prototype with stronger actuators. Actuator type: PM-307-103 [162]	67
4.9	Specialized, discontinued prototype with low-frequency Lofelt actuators [125]. This prototype was developed during a Master’s thesis by Jonas Bock. Left: HapticHead with 24 Lofelt L5 actuators in custom casings; top-right: custom driver board with 12 dual MAX98306 stereo amplifiers and a custom FPGA from the MultiWave project [186] to generate sinus wave inputs for the driver board; bottom-right: Single Lofelt actuator and custom 3D printed casing designed by Jonas Bock. Image shown with friendly permission by Lofelt GmbH.	68
5.1	Descriptions of all 30 patterns. Left/right/reversed refer to different variants. These different variants may have different directions (e.g., left/right) or normal/reversed starting/ending points and play direction (e.g., forward/backward). Note that all static pattern with a dynamic stimulus location have a 100 ms overlap between active actuators to make the patterns feel smoother. The actuators we use need an active time of greater than 40 ms due to their lag time as they would otherwise not produce any stimulus.	73
5.2	Visualizations of representative static pattern with a dynamic stimulus location. From left to right: BackHeadInsideOut, CircleLeftForwardSlow, NavLeft, SpiralLeftForward, TopHeadXSimultaneous. Starting point(s) are marked in blue, ending point(s) are marked in darker green. The order of actuators is marked in bright green numbers and through the arrows on the paths. Note that there are simultaneous activations for some patterns.	74
5.3	Overall confusion matrix of all 30 patterns, N=1800 trials (30 patterns x 3 repetitions x 20 participants).	75
5.4	Post-questionnaire results of the Distinguishability of Patterns Experiment about properties of the 30 implemented TPs, $N = 20$	77
5.5	Internal and front view of the modified Samsung Galaxy S6 dummy. . . .	78

5.6	Comparison of the mean success rates (message type and priority separately) of Brown et al. [24] and our experiment. Significant differences between our conditions are marked (Wilcoxon signed-rank test, sign. level $p=0.05$).	81
5.7	Results of the post-questionnaire of the Tactile Pattern Notification Experiment about recognizing TPs on the head vs. a smartphone (N = 20). Significant differences between statements found using Wilcoxon signed-rank test. Significance level $p=0.005$	83
6.1	The final version of <i>VRTactileDraw</i> in action. Users wear an HTC Vive Pro VR headset and a tactile user interface. A user draws symmetric strokes on the 3D model head in VR (left) and then replays the resulting pattern (right). During drawing, the user can feel the resulting tactile actuation.	88
6.2	Curve interface (left) and drawing interface (right) from prior work [94].	89
6.3	Design space of TP editors after [152], modified to include <i>VRTactileDraw</i> and various other recent works.	91
6.4	First version of the head-up display with annotated elements.	98
6.5	System usability scale (SUS) questions and answers.	100
6.6	Questionnaire on intuitiveness and interface design.	101
6.7	Questionnaire on comfort.	102
6.8	AttrakDiff: Portfolio of results (left), Diagram of word pairs (right).	103
6.10	Mapping of the controller menu in the final prototype. Some functions were not present in the preliminary and main user studies. Left: selection controller, right: drawing controller. Menu items can be selected by moving the thumb in the appropriate direction on the touchpad and then pressing down. Pressing down multiple times without fully lifting the thumb triggers the action repeatedly. Some menu items lead to an adjustable slider, which is confirmed by pressing down (bottom center).	105
6.11	Final version of the HUD. Recording mode is currently active and a stroke is being drawn.	105
6.12	<i>VRTactileDraw</i> driving a 52-actuator tactile vest mockup (left) and the users' view inside an HTC Vive (right).	106
6.13	<i>VRTactileDraw</i> driving the MultiWave prototype vest [187]: A wearable vibrotactile vest containing 76 actuators (left) and the user's view inside an HTC Vive (right). The user is replaying a TP and simultaneously feels vibrations as indicated by the heatmap.	107
7.3	Questionnaire on appropriateness of vibrotactile feedback	114
8.2	Attention funnels with a tiny red crosshair in the view's center. Visual feedback from the user's perspective.	120
8.3	Boxplots of completion times for all conditions with merged data from all participants.	122

8.6	Selection time by angular distance between the start box and the target center.	124
8.7	Selection time by yaw (horizontal heading) distance between the start box and the target center.	124
8.8	Selection time by elevation angle between the ground plane and the target center.	125
8.9	Average learning effect – time. Merged data, all participants. Curves: Gaussian weighted moving average (width=3, blue=visual, green=vibrotactile, red=auditory).	125
8.10	Average learning effect - success rate. Merged data, all participants. Curves: Gaussian weighted moving average (width=3, blue=visual, green=vibrotactile, red=auditory).	126
8.11	Subjective results of the Guidance Comparison Experiment. Diverging stacked bar chart: scales in percent, and absolute values on bars.	127
8.12	Side and front view of modeled actuator positions. Does not fit perfectly due to arbitrary size and asymmetries of the Styrofoam head. The refined guidance algorithm uses triangles between actuators, including a virtual point zero (in green) between the eyes.	128
8.13	Intensity calculation visualization.	128
8.14	Subjective results of the Guidance Precision and Real-World 3D Guidance Experiments. Diverging stacked bar chart: scales in percent, and absolute values on bars.	132
8.15	10 items for the Real-World 3D Guidance Experiment. From left to right (height): Three books (0.9 m), a pen (1.8 m), a Lego piece (1.1 m), a ball (1.7 m, 6 cm diameter), a screwdriver (0.8 m), a remote control (1.4 m) and two balls (3 cm and 12 cm diameter, on the ground).	133
9.1	Visualization of patterns. Blue: start of pattern; green: end of pattern. The UP pattern uses two actuators as end points, because the top of the head is less sensitive to tactile stimuli (see chapter 3 and [144]) and needs a stronger tactile impulse to emphasize the direction at the end of the pattern.	145
9.2	Results of the post-questionnaire of the Pattern Recognition Performance Study about the suitability of the chosen four tactile patterns (TPs) for conveying the four instructions (N = 10).	146
9.3	Photo and model of obstacle course. Color coding: green – start/end zones; red – obstacle; blue – stairs.	149
9.4	Model of the staircase with measurements. Top view in center of Fig. 9.3b. The stairs are 80 cm wide.	149
9.5	Participant climbing the stairs while being accompanied by one of the experimenters for safety reasons.	150


9.6	Number of times participants needed assistance when on the stairs and number of obstacles hit (including participants just brushing the obstacles with their clothes). Participants 12 and 13 are VIPs, the others are sighted, but all of them were blindfolded. There were 24 runs per participant, and there could be multiple hits or needs for assistance in a single run. . . .	151
9.7	Mean time needed per run (across all participants). The means are smoothed by a weighted moving average with a window size of 3. The weights are 0.25, 0.50, and 0.25.	151
9.8	Boxplots showing the median absolute deviations of each participant from the optimal path for all trials. Participants 12 and 13 are VIPs.	152
9.9	The median (left=red, right=blue) deviation per trial for all participants (N=15). The numbers above the boxplots show the percentage of time a participant spent on the left of the optimal path and the average deviation to the left. Ideally, these numbers should be 50% and 0 cm; deviations from them indicate a systematic error. Participants 12 and 13 are VIPs. .	153
9.10	Subjective results of the post-study questionnaire.	154
9.11	Confusion matrices of the most interesting results from the Static Patterns Refinement Study. Cell values are absolute with $N = 100$ per row.	159
9.12	Top view of the obstacle course for the Improved System Validation Experiment, including dimensions. Cardboard obstacle (C1-C5) dimensions are 53.5 cm×33 cm (height 61.5 cm). Starting and ending zones are marked in blue. The dimensions of the stairs are given in Fig. 9.4.	161
9.13	The six possible ways (including reverses) for the Improved System Validation Experiment. Start and end zones are shown in green and blue. The black line shows the smoothed optimal way pre-calculated by our navigation algorithm.	162
9.14	Improved System Validation Experiment – mean left and right median deviations from the path. The red boxplot shows mean left deviations, the blue boxplot mean right deviations. Note that P2's results have to be viewed with caution due to broken cables of actuators (see text).	166
9.15	Improved System Validation Experiment – wobble of different participants on way 5 (left) and way 1 (right). The red line shows the shoulder tracker average of the respective participant.	167
9.16	Questionnaire results of the Improved System Validation Experiment. . .	168
11.1	Vision of a future 2D micro-navigation system for people with visual impairments, based on the original HapticHead concept.	180

List of Tables


2.1	Low-Threshold Mechanoreceptors (LTMRs) in glabrous skin and their properties. Excerpt from [155, p. 541].	15
2.2	Low-Threshold Mechanoreceptors (LTMRs) longitudinal lanceolate endings (LLEs) surrounding hair follicles in hairy skin. Table generated from data available in [1, 155].	19
2.3	Overview of head-based tactile systems, their use cases, and their limitations compared to the HapticHead concept presented in this thesis.	30
3.1	Mean absolute localization accuracy and misclassification of single actuators for different head regions (Fig. 3.3a, blue labels).	45
3.2	Midline bias for different head regions. A positive value for the left actuators (negative for right) shows the average bias towards the midline (red background). A blue background shows a bias away from the midline. Head regions with significant differences between the groups of left and right actuators are marked with a star. For both headSide regions, A1 is on the front and A7 on the back of the head.	45
5.1	Means of recognition rates of the top 15 TPs in the Distinguishability of Patterns Experiment and correlating Likert ratings (0: strongly disagree – 4: strongly agree), sorted by rec. rate.	76
5.2	Tactile Pattern Notification Experiment: Assignment of patterns to message type (spatially distributed rhythm) and source (location) for the HapticHead condition.	80
8.1	Task completion times and success rates for different feedback conditions.	121
9.1	Improved System Validation Experiment – participants overview. Note that P2 was excluded from the average results due to broken cables of actuators (see text).	165

A.1 Exemplary User Study Questionnaire

This section contains an exemplary user study questionnaire from the Improved System Validation Experiment (see section 9.7). All other experiments used questionnaires very similar to this one.



Human-Computer
Interaction Group



Leibniz
Universität
Hannover

HapticHead Blindennavigation - Einverständniserkl...

1 Einverständniserklärung zur Studie: HapticHead Blindennavigation

Diese Studie wird am Fachgebiet Mensch-Computer-Interaktion der Leibniz Universität Hannover durchgeführt. Die Experimentatoren sind Oliver Beren Kaul und Benjamin Simon.

Falls Sie nicht über ausreichend Restsehfähigkeit verfügen, so werden Ihnen alle im Rahmen der Studie benötigten Informationen, Einverständniserklärungen und Fragebögen von einem der Experimentatoren vorgelesen. In diesem Fall ist auch eine Audioaufzeichnung dieses Gesprächs zur Dokumentation Ihres Einverständnisses zu den Einverständniserklärungen nötig. Falls Sie damit nicht einverstanden sind, so können Sie nicht an der Studie teilnehmen und es wird an dieser Stelle abgebrochen.

[Beginn Audioaufzeichnung im Falle der Zustimmung]

Sind Sie mit der Aufzeichnung dieses Gesprächs zur Einholung Ihres Einverständnisses zur Studie HapticHead Blindennavigation einverstanden?

Da in dieser Studie ein Hindernisparcours inklusive einer zweistufigen Treppe durchlaufen wird, sind folgende Risiken zu beachten:

- Sturzrisiko. Mindestens einer der Experimentatoren wird stets in Ihrer direkten Nähe sein und Sie bestmöglich absichern bzw. auffangen. Dennoch ist ein Sturz nicht vollkommen auszuschließen.
- Überanstrengungsgefahr. Die Studie wird insgesamt ca. 60 Minuten dauern und wird im Gehen bzw. Stehen durchgeführt. Falls Sie sich unwohl oder überanstrengt fühlen, so können Sie jederzeit eine Pause einlegen oder die Studie abbrechen.

1.1 Wichtig! Unter den folgenden Umständen darf ich nicht an der Studie teilnehmen:

- Bei Schädigung sensorischer Nerven oder überempfindlicher Haut
- Bei Kopfschmerzen am Tag der Studie
- Bei sonstigen körperlichen Schmerzen / Unwohlsein, welche die Aufmerksamkeit beeinträchtigen
- Bei klastrophobischen Neigungen oder Ängsten

[] Keiner dieser Punkte trifft auf mich zu.
[] Mindestens einer der Punkte trifft auf mich zu.

1

1.2 Datenaufzeichnung

Ich wurde darüber informiert, dass während der Studie Daten aufgezeichnet, elektronisch gespeichert und zur Auswertung der Studie herangezogen werden. Die aufgezeichneten Daten werden ausschließlich für die wissenschaftliche Nutzung anonymisiert ausgewertet. Zu den aufgezeichneten Daten gehören neben den in den Fragebögen abgefragten Daten Ihre Bewegungs-Trajektorien durch den Hindernisparcours in Form von Positionsdaten und das eventuelle Berühren von Hindernissen.

Damit bin ich

einverstanden.

nicht einverstanden.

1.3 Unterschrift

Ich habe die oben genannten Risiken zur Kenntnis genommen und fühle mich gesundheitlich in der Lage, an der Studie teilzunehmen.

Ich habe den Überblick über die Studie verstanden. Ich nehme freiwillig und ohne Vergütung an dieser Studie teil. Ich habe das Recht die Teilnahme jederzeit und ohne Angabe von Gründen abzuberechen.

Vorname: _____ Nachname: _____

Ort, Datum: _____ Unterschrift: _____

2 Einverständniserklärung zu Foto-, Video- und Audioaufnahmen (optional)

Durch Unterzeichnung dieser Einverständniserklärung zu Foto-, Video- und Audioaufnahmen gebe ich dem Fachgebiet Mensch-Computer-Interaktion der Leibniz Universität Hannover die Erlaubnis, entsprechende Aufnahmen von mir, die während meiner Studienteilnahme angefertigt worden sind, zu nutzen. Diese Erlaubnis gebe ich freiwillig.

Für dieses Einverständnis erhalte ich weder jetzt noch in der Zukunft eine Vergütung. Die Aufnahmen dürfen ausschließlich zu wissenschaftlichen Zwecken verwendet werden.

2.1 Einverständniserklärung zur internen Nutzung (optional)

Die Aufnahmen dürfen intern im Rahmen der Analyse des Experiments ausgewertet werden. Die Aufnahmen dürfen weiterhin als Ganzes oder in Teilen zusammen mit anderen Aufnahmen im internen, wissenschaftlichen Kontext (z.B. in der universitären Lehre) gezeigt werden.

Damit bin ich

- einverstanden.
 nicht einverstanden.

2.2 Einverständniserklärung zur Veröffentlichung (optional)

Die Aufnahmen dürfen als Ganzes oder in Teilen zusammen mit anderen Aufnahmen im öffentlichen wissenschaftlichen Kontext (z.B. auf Konferenzen) oder Online in Form von z.B. Videos, welche wissenschaftliche Veröffentlichungen begleiten, gezeigt werden.

Ein Beispiel einer Online-Veröffentlichung sind Teilausschnitte Ihrer Navigationsversuche im Hindernisparcours als Teil eines YouTube Videos im Kanal des Fachgebiets Mensch-Computer Interaktion oder im Kanal einer wissenschaftlichen Konferenz. Dieses Video würde eine wissenschaftliche Veröffentlichung begleiten und zusätzlich Ergebnisse aus dem Experiment präsentieren. Dabei wären Sie als Person in Bewegung zu sehen, inklusive Ihrer Kleidung. Ihr Gesicht wäre jedoch fast vollständig von einer Schlafmaske und dem Prototypen verdeckt.

Diesem Punkt kann nur zugestimmt werden, wenn Sie sich aufgrund von Restsehfähigkeit oder der Einschätzung einer Vertrauensperson in der Lage sehen, Ihren visuellen Eindruck im Video nachträglich zu kontrollieren und ggfs. Ihre Zustimmung zurückzuziehen, sollten Sie dann nicht mehr mit einer Veröffentlichung einverstanden sein.

Damit bin ich

- einverstanden.
 nicht einverstanden.

2.3 Unterschrift

Mit meiner Unterschrift stimme ich den oben genannten Ausführungen zu.
Die Einverständniserklärungen zu den Aufnahmen können jederzeit durch Streichung folgenlos zurückgenommen werden. In diesem Fall werden alle bis dahin entstandenen, personenbezogenen Daten unverzüglich gelöscht.

Vorname: _____ Nachname: _____

Ort, Datum: _____ Unterschrift: _____

Ende der rechtlichen Absicherung

1 Überblick über die Studie

Dieser Überblick dient dazu, Ihnen das Experiment vorzustellen und Sie auf Ihre Rechte als freiwilliger Versuchsteilnehmer hinzuweisen. Bitte fragen Sie nach, wenn etwas unklar sein sollte. Es geht in dieser Studie um ein taktiles Interface am Kopf, welches auf seine Eignung zur Navigationshilfe durch einen Hindernisparcours (Mikronavigation) für Menschen mit visuellen Einschränkungen untersucht wird.

Vielen Dank, dass Sie an dieser Studie teilnehmen. Wir erheben Daten von mehreren Teilnehmern. Diese Daten helfen uns, interaktive Systeme zu evaluieren und zu verbessern. Es geht um die Evaluierung des Systems und nicht um Ihr individuelles Abschneiden im Experiment. Die Daten werden in der Regel gemittelt und der Trend aller Daten ist für unsere Forschung von Interesse.

Zu Beginn der Studie werden Sie gebeten einen Fragebogen auszufüllen, in dem u.a. Alter, Geschlecht und bisherige Erfahrungen mit taktilem Feedback beziehungsweise der taktilen Wahrnehmung am Kopf erfragt werden.

Im Verlauf der Studie untersuchen wir ein taktiles Interface am Kopf auf seine Eignung zur Navigationshilfe für Menschen mit visuellen Einschränkungen. Dazu tragen Sie einen Prototypen auf dem Kopf, welcher aus einer Badekappe mit einem Kinngurt und 24 kleinen vibrotaktilen Aktuatoren (ähnlich dem Vibrationsmotor in einem Smartphone) besteht. Dieser Prototyp wird Ihnen durch intuitive taktile Muster den Weg durch einen Hindernisparcours zeigen, welcher aus Pappkartons und einer zweistufigen Treppe besteht. Es gibt mehrere Routen, welche vom System zufällig ausgewählt werden. Nicht alle der Routen beinhalten die Treppe.

Während des Tests tragen Sie weiterhin eine Schlafmaske, um eventuelle Rest-Seh-Fähigkeiten auszuschließen. Mindestens einer der Experimentatoren ist während der gesamten Navigationsaufgaben, ohne Sie zu berühren, in Ihrer unmittelbaren Umgebung, um mögliche Stürze zu verhindern.

Nachdem Sie die Navigationsaufgaben bewältigt haben, werden Sie gebeten, einen kurzen Fragebogen über Ihre Erfahrungen während der Studie auszufüllen.

Die Studie wird insgesamt ca. 60 Minuten dauern. Sie haben das Recht die Teilnahme an dieser Studie jederzeit und ohne Angabe von Gründen abubrechen (Schokolade gibt es trotzdem!). Machen Sie insbesondere bei Unwohlsein von diesem Recht Gebrauch. Weiterhin können Sie zwischen den einzelnen Versuchen jederzeit eine Pause einlegen.

2 Allgemeine Angaben

2.1 Teilnehmernummer

2.2 Geschlecht

- weiblich
 männlich
 anderes: _____

2.3 Alter

2.4 Dominante Hand

- links
 rechts
 beidhändig

2.5 Art der Sehbehinderung - angeboren oder seit welchem Alter?

2.6 Restsehfähigkeit

2.7 Benutzen Sie einen Blindenstock?

- Ja
- Manchmal, aber nicht immer
- Nein

3 Abschlussfragebogen

3.1 Bitte bewerten Sie die folgenden Aussagen:

	Trifft gar nicht zu			Trifft völlig zu	
Die Art der Navigation hat mir gefallen	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich habe mich während der Navigation sicher gefühlt	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich habe die statischen Muster durcheinander gebracht (STOP, HOCH, RUNTER, ACHTUNG)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Die Art der Navigation war für mich intuitiv	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich konnte mir die statischen Muster gut merken (STOP, HOCH, RUNTER, ACHTUNG)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Das Erkennen der Treppe durch die taktilen Muster fiel mir leicht	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Während der Navigation wusste ich genau, wo ich im Parcours war	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Das Treppensteigen wird durch das System vereinfacht	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Diese Art der Navigation war mir zu unsicher	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

3.2 Bitte bewerten Sie die folgenden Aussagen:

	Trifft gar nicht zu			Trifft völlig zu	
Ich könnte mir vorstellen, diese Art der Navigation in der realen Welt zu verwenden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich habe das Vibrations-Feedback zu schwach gespürt	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich habe das Navigations-Feedback als störend empfunden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich habe mich an das Navigations-Feedback nach einigen Versuchen gewöhnt	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich fand die Feedback-Position für das Vibrations-Feedback am Kopf passend	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich würde mir Gedanken über die Ästhetik des Systems machen, auch wenn es unter einer Mütze versteckt ist	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich kann mir die regelmäßige Benutzung des Systems im Alltag vorstellen bezüglich Wohlbefinden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich kann mir die regelmäßige Benutzung des Systems im Alltag vorstellen bezüglich Hilfe bei der Navigation	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich kann mir vorstellen, dass ein solches System bei der Mikronavigation den Blindenstock ersetzen kann	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Ich kann mir vorstellen, dass ein solches System bei der Makronavigation den Blindenstock ersetzen kann	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bei diesem System fehlt mir das taktile Feedback des Blindenstocks bezüglich Untergrund-Beschaffenheit (z.B. steinig oder glatt)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

4 Abschlussfragebogen (Fortsetzung)

4.1 In welchen sozialen Kontexten können Sie sich die Verwendung des Systems vorstellen?

4.2 In welchen sozialen Kontexten können Sie sich die Verwendung des Systems -nicht- vorstellen?

4.3 Haben Sie weitere Anmerkungen oder Anregungen?

Curriculum Vitae

Full Name	Oliver Beren Kaul
Date of Birth	October 8 th , 1987
Place of Birth	Gehrden, Germany

Education

Aug 2015 – May 2021	Human-Computer Interaction (Doctor of Philosophy) Leibniz University Hannover
Oct 2011 – May 2015	Computer Science (Master of Science) Leibniz University Hannover
Oct 2007 – Sep 2011	Computer Science (Bachelor of Science) Leibniz University Hannover
Jun 2007	Abitur Gymnasium unter den Eichen, Uetze, Germany

Professional Experience

Dec 2013 – Feb 2020	Mobile Developer miovent.de
Oct 2014 – Mar 2015	Research Assistant TNT Group, Leibniz University Hannover
Jan 2012 – Dec 2014	Research Assistant SRA Group, Leibniz University Hannover

Complete List of Publications

Link to my *Google Scholar* profile:

<https://scholar.google.com/citations?user=UZCq61sAAAAJ>

Journal Papers

Oliver Beren Kaul et al. “Around-the-Head Spatial Tactile System for Supporting Micro Navigation of People with Visual Impairments”. In: *ACM Transactions on Computer-Human Interaction* 28.4 (Aug. 2021), p. 34

Full Papers

Oliver Beren Kaul et al. “VRTactileDraw: A Virtual Reality Tactile Pattern Designer for Complex Spatial Arrangements of Actuators”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (2021)

Oliver Beren Kaul et al. “Design and Evaluation of On-the-Head Spatial Tactile Patterns”. In: *19th International Conference on Mobile and Ubiquitous Multimedia*. New York, NY, USA: ACM, Nov. 2020, pp. 229–239. ISBN: 9781450388702. DOI: 10.1145/3428361.3428407

Oliver Beren Kaul et al. “Vibrotactile Funneling Illusion and Localization Performance on the Head”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, Apr. 2020, pp. 1–13. ISBN: 9781450367080. DOI: 10.1145/3313831.3376335

Oliver Beren Kaul and Michael Rohs. “HapticHead: A Spherical Vibrotactile Grid around the Head for 3D Guidance in Virtual and Augmented Reality”. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. New York, New York, USA: ACM Press, 2017, pp. 3729–3740. ISBN: 9781450346559. DOI: 10.1145/3025453.3025684

Short Papers

Oliver Beren Kaul et al. “Increasing presence in virtual reality with a vibrotactile grid around the head”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 10516 LNCS (2017). Ed. by Regina Bernhaupt et al., pp. 289–298. ISSN: 16113349. DOI: 10.1007/978-3-319-68059-0_19

Late Breaking Work

Oliver Beren Kaul et al. “Mobile Recognition and Tracking of Objects in the Environment through Augmented Reality and 3D Audio Cues for People with Visual Impairments”. In: *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2021. DOI: 10.1145/3411763.3451611

Oliver Beren Kaul et al. “3DTactileDraw: A Tactile Pattern Design Interface for Complex Arrangements of Actuators”. In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2019, pp. 1–6. ISBN: 9781450359719. DOI: 10.1145/3290607.3313030

Oliver Beren Kaul and Michael Rohs. “HapticHead: 3D Guidance and Target Acquisition through a Vibrotactile Grid”. In: *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2016, pp. 2533–2539. ISBN: 9781450340823. DOI: 10.1145/2851581.2892355

Oliver Beren Kaul et al. “Follow the Force: Steering the Index Finger towards Targets using EMS”. in: *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. New York, NY, USA: ACM, May 2016, pp. 2526–2532. ISBN: 9781450340823. DOI: 10.1145/2851581.2892352

Workshop Papers

Oliver Beren Kaul and Michael Rohs. “Concept for Navigating the Visually Impaired using a Tactile Interface around the Head”. In: *Proceedings of the CHI 2019 Workshop on Hacking Blind Navigation*. 2019. URL: <https://blindnavigationchi19.wordpress.com/>

Oliver Beren Kaul and Michael Rohs. “Requirements of Navigation Support Systems for People with Visual Impairments”. In: *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. New York, NY, USA: ACM, Oct. 2018, pp. 680–683. ISBN: 9781450359665. DOI: 10.1145/3267305.3267685

O.B. Kaul and M. Rohs. “Wearable head-mounted 3D tactile display application scenarios”. In: *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct, MobileHCI 2016*. 2016. ISBN: 9781450344135. DOI: 10.1145/2957265.2965022

Declaration

I declare that this thesis has been composed solely by myself and that it has not been submitted, in whole or in part, in any previous application for a degree. Except where states otherwise by reference or acknowledgment, the work presented is entirely my own.

Hannover, 28. Mai 2021

Oliver Beren Kaul