
SAARLAND UNIVERSITY

Faculty of Mathematics and Computer Science
Department of Computer Science
Dissertation



Hand-based Illusions for Haptics in Virtual Reality

Dissertation zur Erlangung des Grades des
Doktors der Ingenieurwissenschaften (Dr.-Ing.)
der Fakultät für Mathematik und Informatik
der Universität des Saarlandes

vorgelegt von
Martin Feick (M.Sc.)
Saarbrücken
2024

Date of the Colloquium: 10th of March, 2025

Dean: Professor Dr. Roland Speicher

Examination Board

Chair: Professor Dr. Jürgen Steimle

Reporter: Professor Dr. Antonio Krüger
Professor Dr. Anthony Tang
Dr. Mar Gonzalez-Franco
Professor Dr. Albrecht Schmidt

Scientific Assistant: Dr. Martin Schmitz

Notes on style:

The majority of the work presented in this thesis was conducted in collaboration with other researchers and students. For this reason, the scientific plural “we” is used throughout this thesis. References to web resources (e.g., websites or web articles) are provided in footnotes as URLs (long URLs have been shortened). Additionally, links to videos of our work and to contributed open-source repositories are provided as QR codes in the respective chapters. References to scientific publications (e.g., articles in journals or conference proceedings) are provided in the Bibliography at the end of this thesis.

"Sometimes it is the people no one can imagine anything of who do the things no one can imagine."

Alan Turing—The Imitation Game

ACKNOWLEDGMENTS

As I reflect on my journey, I am grateful for the support and encouragement I received from my colleagues, friends, and family. Successfully completing my doctoral dissertation would not have been possible without **YOU!** I could probably write another manuscript, similar to the length of this dissertation, to cover everyone who played a role in my academic and personal development over the past years. So, even if you do not find your name mentioned here, I thank you from the bottom of my heart!

First and foremost, I would like to express my deepest gratitude to my primary supervisor, **Antonio Krüger**, for the opportunity to pursue my PhD in his research group. I felt very welcome, and you have created an environment I enjoyed being part of. You allowed me to follow my passion and explore research directions that excited me, for which I am very grateful. Your support in paper publishing, sending me abroad, career advice, and fighting for me when things turned against me shows your incredible character—THANK YOU. To my secondary supervisor, **Anthony Tang**, who has tremendously impacted my life, academically and personally. Eight years ago, I reached out to you because I was interested in doing my Bachelor's thesis in your lab in Calgary, Canada. You took me under your wing without knowing me, and several years later, we are still working together. I genuinely enjoyed our regular meetings, where we (most of the time) chatted and discussed things that had very little to do with my PhD topic. You are among the most outstanding people I know, and you shaped me into an HCI researcher.

Further, I would like to thank **Mar Gonzalez-Franco** and **Albrecht Schmidt** for agreeing to review my thesis. Your work has always inspired me, and I am honored to have you on my committee. I am also deeply grateful for my past supervisors and academic mentors, **Nic Marquardt**, **Scott Bateman**, **Lora Oehlberg**, **Ehud Sharlin** and **Jürgen Steimle**, for teaching me a lot about how to do HCI research and giving me a skillset that was vital in my academic development. A special thank you to **André Miede**, who encouraged me to jump into the unknown. I would like to extend my appreciation to **Stefanie Müller**, who accepted me as a visiting PhD at M.I.T.

To all my lab mates at the Ubiquitous Media Technology Lab, I am very grateful for the camaraderie and support we shared throughout this journey, even during the challenging times of COVID-19. Thanks to **Dimitar Valkov** for sharing an office with me and being hypercritical about everyone's work. To **Denise Kahl** and **Martin Schmitz** for all the discussions, ideas and advice! To **Felix Kosmalla** and **Julian Wolter** for being amazing engineers and keeping our equipment alive. I would also like to thank **Anke Hirsch** and **Nina Knieriemen** for taking the time to advise me on experimental designs and **Rudolf Siegel** for teaching me about Bayesian inference. Another thank you goes to all the people at the German Research Center of Artificial Intelligence, especially the COS department. To **Michael Feld**, **Tim Schwartz**, **Florian Daiber** and **Pascal Lessel** for being such amazing project leaders, allowing me to focus on my PhD studies and being very accommodating. A massive THANK YOU goes to **Iris Lambrecht**, **Gundula Kleiner**, **Lisa-Marie Merziger**, and **Christine Pyttlik** for making the impossible possible =) I would also like to thank **Margaret De Lap** for proofreading many of my publications.

A special thanks to **André Zenner**, whose incredible support was essential to many projects in this thesis. I appreciate your time and greatly enjoyed our collaborations, joking, advising students, organizing seminars, and going for dinner after an exhausting day of pilots, studies, or photo shootings. To **Donald Degraen** for being such a strong researcher, always knowing what to do next. I learned quite a lot from working with you. Thank you to **Niko Kleer**, whom I have known since the first semester of my computer science studies for countless hours discussing research, live and playing games.

A big thank you goes to all my collaborators; especially, to **Kora Regitz**, the *GOAT* research assistant. Your support during many studies was beyond what I could ever ask for! I would also like to thank all the students who selected me as their thesis advisor. A special thanks goes to **Fabian Hupperich** and **Simon Seibert**, as the results of our joint studies made it into parts of this dissertation.

Thank you to my friends for their support and encouragement. Your belief in me provided the motivation I needed during challenging times. Finally, to my family, especially my parents **Annette** and **Mario Feick**, for their unconditional support in success and setbacks. For sponsoring my research visits to Calgary and London, always encouraging me to take the next step, not putting limits on myself.

Finally, to my lovely wife, **Denise Feick**, your love and support have been my anchor. Thank you for your input, patience, understanding, and sacrifices that made this journey possible.

ABSTRACT

Virtual Reality (VR) enables us to dive into artificially generated worlds, creating the illusion of being elsewhere. However, this illusion falls apart when using our hands to interact with objects inside VR because nothing can be physically touched. Haptic feedback is crucial to our everyday interactions in reality, so its sudden absence can disrupt the immersive nature of VR. To address this, a single physical proxy object can approximate virtual objects' properties to form a combined visuo-haptic illusion. We introduce four novel proxy-based approaches that render tactile and kinesthetic haptic feedback for object interactions in VR. To alter proxies' perceived properties, we explore perceptual illusions that visually offset virtual hand interactions from their real-world counterpart. Here, we contribute three novel hand-based illusions that can simulate different haptic effects. However, introducing offsets between what users see and what they feel risks disrupting the VR experience. Therefore, we set out to quantify the undetectable offset for various types of interactions, properties of proxies, and users' virtual representations to understand the techniques' application limits. Finally, we present a method for continuously monitoring and tailoring hand-based illusions to individuals' sensitivity to offsets. Together, this thesis advances the field of haptics for hand-based object interactions in VR.

ZUSAMMENFASSUNG

Die Virtuelle Realität (VR) ermöglicht es uns, in künstlich erzeugte Welten einzutauchen und die Illusion zu erleben, an einem anderen Ort zu sein. Sobald wir jedoch unsere Hände benutzen, um mit virtuellen Objekten zu interagieren, zerbricht die Illusion, da nichts berührt werden kann. Haptisches Feedback ist für unsere alltäglichen Interaktionen von entscheidender Bedeutung, sodass sein plötzliches Fehlen das VR-Erlebnis stört. Um dieses Problem zu lösen, kann ein einzelnes physisches Proxy die Eigenschaften virtueller Objekte annähern und so eine Illusion der simulierten Objekte erzeugen. Wir stellen vier neuartige Ansätze zur Verbesserung des taktilen und kinästhetischen haptischen Feedbacks mittels Proxys vor. Um die haptischen Fähigkeiten dieser zu erweitern, präsentieren wir drei Wahrnehmungssimulationen, die Simulationen haptischer Eigenschaften durch Entkopplung des visuellen und haptischen Sinnes ermöglichen. Die Einführung von Abweichungen zwischen dem, was die Nutzer sehen, und dem, was sie fühlen, kann jedoch das VR-Erlebnis beeinflussen. Daher haben wir die unbemerkbare Abweichung für verschiedene Arten von Interaktionen, Proxys und Repräsentationen der Nutzer quantifiziert, um die Grenzen der Techniken zu verstehen. Abschließend präsentieren wir eine Methode zur kontinuierlichen Anpassung der Abweichung an die individuellen Wahrnehmungsgrenzen von Nutzern. Die Arbeit leistet einen wertvollen Beitrag zum Feld der Haptik für handbasierte Interaktionen in VR.

PUBLICATIONS

Preamble. Parts of this dissertation, including text passages, figures, tables, ideas, concepts, implementations, presented applications and uses cases, studies and experiments, results, discussions, and conclusions, have previously appeared in scientific publications and their supplementary materials. In the following, we list the publications that cover this dissertation’s central contributions. I (**author**) have supervised and guided the development of all prototypes and studies presented in this dissertation. **Collaborators** include publication co-authors and people in acknowledgments.

7 x Full Conference Papers

[98] **Martin Feick**, Kora Regitz, Lukas Gehrke, André Zenner, Anthony Tang, Tobias Jungbluth, Maurice Rekrut, and Antonio Krüger. Predicting the Limits: Tailoring Unnoticeable Hand Redirection Offsets in Virtual Reality to Individuals’ Perceptual Boundaries.

*In Proceedings of
ACM UIST 2024*

⇒ *Parts of this paper appear in Chapter 2 and Chapter 6*

Author only: initial idea, writing—original draft, open-sourcing, project administration; **Author + collaborators:** conceptualization, study execution, data analysis, writing—review & edit, visualizations; **Collaborators under author’s supervision:** prototype development

[103] **Martin Feick**, André Zenner, Simon Seibert, Anthony Tang, and Antonio Krüger. The Impact of Avatar Completeness on Embodiment and the Detectability of Hand Redirection in Virtual Reality.

*In Proceedings of
ACM CHI 2024*

⇒ *Parts of this paper appear in Chapter 2 and Chapter 5*

Author only: initial idea, data analysis, writing—original draft, project administration; **Author + collaborators:** conceptualization, writing—review & edit, visualizations; **Collaborators under author’s supervision:** prototype development, study execution

[104] **Martin Feick**, André Zenner, Oscar Ariza, Anthony Tang, Cihan Biyikli, and Antonio Krüger. Turn-It-Up: Rendering Resistance for Knobs in Virtual Reality through Undetectable Pseudo-Haptics.

*In Proceedings of
ACM UIST 2023*

⇒ *Parts of this paper appear in Chapter 2 and Chapter 4*

Author only: study execution, developing applications, writing—original draft, project administration; **Author + collaborators:** initial idea, conceptualization, writing—review & edit, visualizations, data analysis; **Collaborators under author’s supervision:** prototype development

*In Proceedings of
UIST 2023*

[92] **Martin Feick**, Cihan Biyikli, Kiran Gani, Anton Wittig, Anthony Tang, and Antonio Krüger. VoxelHap: A Toolkit for Constructing Proxies Providing Tactile and Kinesthetic Haptic Feedback in Virtual Reality. \Rightarrow *Parts of this paper appear in Chapter 3 and Chapter 2*

Author only: initial idea, data analysis, writing—original draft, project administration; **Author + collaborators:** conceptualization, prototype development, writing—review & edit, open-sourcing, visualizations; **Collaborators under author’s supervision:** study execution, developing applications

*In Proceedings of
IEEE VR 2023*

[97] **Martin Feick**, Kora Regitz, Anthony Tang, Tobias Jungbluth, Maurice Rekrut, and Antonio Krüger. Investigating Noticeable Hand Redirection in Virtual Reality using Physiological and Interaction Data. \Rightarrow *Parts of this paper appear in Chapter 2 and Chapter 6*

Author only: initial idea, writing—original draft, open-sourcing, project administration; **Author + collaborators:** conceptualization, study execution, writing—review & edit, visualizations, data analysis; **Collaborators under author’s supervision:** prototype development

*In Proceedings of
ACM CHI 2022*

[99] **Martin Feick**, Kora Regitz, Anthony Tang, and Antonio Krüger. Designing Visuo-Haptic Illusions with Proxies in Virtual Reality: Exploration of Grasp, Movement Trajectory and Object Mass. \Rightarrow *Parts of this paper appear in Chapter 2 and Chapter 5*

Author only: initial idea, data analysis, writing—original draft, project administration; **Author + collaborators:** conceptualization, study execution, writing—review & edit, visualizations; **Collaborators under author’s supervision:** developing applications, prototype development

*In Proceedings of
ACM CHI 2021*

[95] **Martin Feick**, Niko Kleer, André Zenner, Anthony Tang, and Antonio Krüger. Visuo-haptic Illusions for Linear Translation and Stretching using Physical Proxies in Virtual Reality. \Rightarrow *Parts of this paper appear in Chapter 2 and Chapter 4*

Author only: prototype development, study execution, data analysis, developing applications, writing—original draft, visualizations, project administration; **Author + collaborators:** initial idea, conceptualization, writing—review & edit

1 x Journal Article

*In Frontiers in
Virtual Reality 2023*

[93] **Martin Feick**, Donald Degraen, Fabian Hupperich and Antonio Krüger. MetaReality: Enhancing Tactile Experiences using Actuated 3D-Printed Metamaterials in Virtual Reality. \Rightarrow *Parts of this paper appear in Chapter 2 and Chapter 3*

Author only: initial idea, project administration; **Author + collaborators:** conceptualization, data analysis, writing—original draft, writing—review & edit, visualizations; **Collaborators under author’s supervision:** prototype development, study execution

1 x Late-Breaking-Work

[94] **Martin Feick**, Niko Kleer, Anthony Tang, and Antonio Krüger. The Virtual Reality Questionnaire Toolkit.

⇒ *Parts of this Late-Breaking-Work appear in Chapter 7*

Author only: initial idea, prototype development, writing—original draft, visualizations, open-sourcing, project administration; **Author + collaborators:** conceptualization, writing—review & edit

*In Proceedings of
ACM UIST 2020
Adjunct*

1 x Demo Paper

[102] **Martin Feick**, André Zenner, Simon Seibert, Oscar Javier Ariza Nunez, David Wagmann, Juliana Helena Keller, Anton Wittig and Antonio Krüger. MoVRI: The Museum of Virtual Reality Illusions.

⇒ *Parts of this Demo Paper appear in Chapter 7*

Author only: writing—original draft, project administration; **Author + collaborators:** initial idea, conceptualization, writing—review & edit, visualizations; **Collaborators under author’s supervision:** prototype development

*Demo at Saarland
Informatics Campus
2023*

Other Publications

During my time as a doctoral student, I have collaborated successfully with other researchers, leading to several publications across VR, human-robot interaction, multimodal interaction, and artificial intelligence. While these projects do not directly appear in my thesis, they helped me sharpen my research skills and gain extensive knowledge in various interdisciplinary fields. This section also covers my first-author publications, which are not part of my dissertation.

[101] **Martin Feick**, Xuxin Tang, Raul Garcia-Martin, Alexandru Luchianov, Roderick Huang, Chang Xiao, Alexa Siu and Mustafa Doga Dogan. Imprinto: Enhancing Infrared Inkjet Watermarking for Human and Machine Perception.

*In Proceedings of
ACM CHI 2025*

[65] Donald Degraen, **Martin Feick**, Serdar Durdyev, and Antonio Krüger. Prototyping Surface Slipperiness using Sole-Attached Textures during Haptic Walking in Virtual Reality.

*In Proceedings of
ACM MUM 2024*

[381] André Zenner, Chiara Karr, **Martin Feick**, Oscar Ariza, and Antonio Krüger. Beyond the Blink: Investigating Combined Saccadic & Blink-Suppressed Hand Redirection in Virtual Reality.

*In Proceedings of
ACM CHI 2024*

[189] Niko Kleer, **Martin Feick**, Amr Gomaa, Michael Feld, and Antonio Krüger. Bridging the Gap to Natural Language-based Grasp Predictions through Semantic Information Extraction.

*In Proceedings of
IEEE/RSJ IROS
2024*

- In Proceedings of
IEEE ROMAN
2024* [190] Niko Kleer, Ole Keil, **Martin Feick**, Amr Gomaa, Tim Schwartz, and Michael Feld. Incorporation of the Intended Task into a Vision-based Grasp Type Predictor for Multi-fingered Robotic Grasping.
- In Companion of
ACM/IEEE HRI
2024* [187] Niko Kleer, **Martin Feick**, Florian Daiber, and Michael Feld. Towards Remote Expert Supported Autonomous Assistant Robots in Shopping Environments.
- In Proceedings of
ACM VRST 2023* [378] André Zenner, Chiara Karr, **Martin Feick**, Oscar Ariza, and Antonio Krüger. The Detectability of Saccadic Hand Offset in Virtual Reality.
- In Proceedings of
ACM UIST 2022
Adjunct* [100] **Martin Feick**, Anthony Tang, and Antonio Krüger. HapticPuppet: A Kinesthetic Mid-air Multidirectional Force-Feedback Drone-based Interface.
- In Proceedings of
IEEE/RSJ IROS
2022* [188] Niko Kleer, **Martin Feick**, and Michael Feld. Leveraging Publicly Available Textual Object Descriptions for Anthropomorphic Robotic Grasp Predictions.
- In Proceedings of
IEEE-RAS
Humanoids 2022* [186] Niko Kleer and **Martin Feick**. A Study on the Influence of Task Dependent Anthropomorphic Grasping Poses for Everyday Objects.
- In Workshop of
ACM CHI 2021* [120] Tobias Geib and **Martin Feick**. Everyday Objects for Volumetric 3D Sketching in Virtual Reality.
- In Workshop of
ACM CHI 2021* [76] Tim Düwel, **Martin Feick**, and Antonio Krüger. Considering Interaction Types for Geometric Primitive Matching.

Supervised Theses

In total, I have supervised two completed Bachelor's theses (with one of them contributing to parts of this dissertation) and six completed Master's theses (with one of them contributing to parts of this dissertation).

- Master Thesis
completed in 2025* 1. Investigating the Impact of Long-term Exposure to Hand Redirection in VR. Kora Regitz.
- Bachelor Thesis
completed in 2023* 2. The Impact of Avatar Completeness on Embodiment and the Detectability of Hand Redirection in VR. Simon Seibert.
⇒ contributed to Chapter 5 and publication [103].
- Master Thesis
completed in 2023* 3. Advancing Real-time Class-Incremental Continual Learning through Experience Replay of VQPs. Eiram Mahera Sheikh.

- | | |
|--|--|
| 4. MetaReality: Dynamic Metamaterials for Visuo-Haptic Texture Perception in VR. Fabian Hupperich.
⇒ contributed to Chapter 3 and publication [93]. | <i>Master Thesis
completed in 2022</i> |
| 5. Hybrid Intelligence for Product Layout Maintenance in Retail Environments. Hanna Schier. | <i>Master Thesis
completed in 2022</i> |
| 6. Task Intent Aware Grasp Type Prediction for Multi-Fingered Robotic Grasping. Ole Keil. | <i>Bachelor Thesis
completed in 2022</i> |
| 7. Simulating Slipperiness of Ground Surfaces in Virtual Reality using Everyday Materials. Serdar Durdyev. | <i>Master Thesis
completed in 2022</i> |
| 8. Designing Pen-based Controller for Volumetric 3D Sketching in VR. Tobias Geib. | <i>Master Thesis
completed in 2022</i> |

CONTENTS

1	INTRODUCTION	1
1.1	Motivation	1
1.2	Problem Statement	2
1.3	Research Questions	4
1.4	Method & Approach	6
1.5	Contributions	8
1.6	Overview	13
2	BACKGROUND & RELATED WORK	17
2.1	The <i>Vision</i> of Virtual Reality	17
2.1.1	A Brief History of Virtual Reality	17
2.1.2	Conceptual Realm of Virtual Reality	20
2.1.3	Current and Envisioned Applications	28
2.2	Human Perception	30
2.2.1	Perceiving our Environemnt	30
2.2.2	Visual Perception	31
2.2.3	Haptic Perception	34
2.2.4	Multisensory Perception	37
2.3	The Grand Challenge of Haptics in Virtual Reality	38
2.3.1	Active Haptic Feedback	40
2.3.2	Passive Haptic Feedback	46
2.3.3	Conclusion	52
2.4	Dynamic Proxies for Virtual Reality	53
2.4.1	Construction of Proxies	53
2.4.2	Fabrication of Proxies	56
2.4.3	Conclusion	58
2.5	Hand-based Illusions for Virtual Reality	59
2.5.1	Designing Illusions with Proxies	60
2.5.2	Detectability of Illusions	68
2.5.3	Conclusion	77
3	CHANGING PROPERTIES OF HAPTIC PROXIES	79
3.1	Constructing Multimodal Haptic Proxies	79
3.1.1	Introduction	79
3.1.2	<i>VoxelHap</i> Toolkit	80
3.1.3	Technical Evaluation	90
3.1.4	Experiment 1—The Impact of Functionality	90
3.1.5	Results	94
3.1.6	Experiment 2—The Impact of Shape Fidelity	95
3.1.7	Results	96
3.1.8	Discussion of <i>VoxelHap</i>	98
3.1.9	Conclusion & Contributions	99

3.2	Enhancing Tactile Perception through Visuo-Haptics . . .	101
3.2.1	Introduction	101
3.2.2	Metamaterials for Tactile Texture Perception . . .	102
3.2.3	Experiment 1: Tactile Metamaterial Perception . . .	108
3.2.4	Results	110
3.2.5	Experiment 2: Visuo-Haptic Perception	112
3.2.6	Results	116
3.2.7	Discussion of Metamaterials	122
3.2.8	Conclusion & Contributions	126
3.3	Summary	127
4	DESIGNING HAND-BASED HAPTIC ILLUSIONS	129
4.1	Towards Pseudo-Haptic Resistance	129
4.1.1	Introduction	129
4.1.2	Undetectable <i>Pseudo-Haptic Resistance</i>	130
4.1.3	Experiment	132
4.1.4	Results	137
4.1.5	A Model of <i>Pseudo-Haptic Resistance</i>	141
4.1.6	Applications & Use Cases	142
4.1.7	Discussion	143
4.1.8	Conclusion & Contributions	145
4.2	Towards Visuo-Haptic Translation	146
4.2.1	Introduction	146
4.2.2	Undetectable <i>Visuo-Haptic Translation</i>	146
4.2.3	Experiment	148
4.2.4	Results	153
4.2.5	Applications & Use Cases	158
4.2.6	Discussion	159
4.2.7	Conclusion & Contributions	161
4.3	Summary	162
5	DETECTABILITY OF HAND-BASED ILLUSIONS	165
5.1	Factors Influencing Detectability	165
5.1.1	Introduction	165
5.1.2	Impact of Grasp, Object Mass and Motion	166
5.1.3	Experiment 1: Grasp Type and Trajectory	170
5.1.4	Results	173
5.1.5	Experiment 2: Grasp Type and Object Mass . . .	175
5.1.6	Results	177
5.1.7	Movement and Individual Differences	178
5.1.8	Effects on Applications & Use Cases	181
5.1.9	Discussion	182
5.1.10	Conclusion & Contributions	183
5.2	Effects of Avatar Representation on Detectability . . .	185
5.2.1	Introduction	185
5.2.2	Experiment	186
5.2.3	Results	190

5.2.4	Discussion	193
5.2.5	Conclusion & Contributions	195
5.3	Summary	196
6	INDIVIDUALIZING HAND-BASED ILLUSIONS	199
6.1	Physiological and Interaction Data during Illusions . .	199
6.1.1	Introduction	199
6.1.2	Physiological & Interaction Data	200
6.1.3	Experiment	201
6.1.4	Results	207
6.1.5	Discussion	218
6.1.6	Conclusion & Contributions	219
6.2	Tailoring Illusions to Users' Sensitivity	222
6.2.1	Introduction	222
6.2.2	Experiment	222
6.2.3	Results	228
6.2.4	Predicting Perceptual Boundaries	234
6.2.5	Discussion	236
6.2.6	Conclusion & Contributions	239
6.3	Summary	240
7	CONCLUSION	243
7.1	Summary & Dissertation Contributions	243
7.1.1	Major Contributions to RQ1	243
7.1.2	Major Contributions to RQ2	245
7.1.3	Major Contributions to RQ3	248
7.1.4	Major Contributions to RQ4	251
7.1.5	Further System Contributions	254
7.2	Future Work & Conclusion	257
7.2.1	Future Directions for Proxy Design	257
7.2.2	Future Directions for Hand-based Illusions . . .	259
7.2.3	Individualizing Hand-based Illusions	261
7.2.4	Long-Term Exposure to Hand-based Illusions .	262
7.2.5	Conclusion	263
	LIST OF FIGURES	265
	LIST OF TABLES	268
	LIST OF ABBREVIATIONS	269
	BIBLIOGRAPHY	271

INTRODUCTION

1.1 MOTIVATION

Imagine Lisa, a medical student who wants to practice her surgical skills in a safe environment under realistic circumstances without the risk of severely hurting a patient. She can read, watch, and learn about the procedures, but without practical experience, replicating the stress and sense of responsibility, it is challenging for her to move from theory to practice confidently. This situation happens not only to our medical student, Lisa, but to many people who touch the unknown, whether it is in acquiring new skills, executing tasks in extreme environments, or using expensive and complex tools. There exists a high demand for affordable technology that enables us to immerse people in real-world or theoretical scenarios in a safe and responsible manner.

Challenge that affects everyone

VR may be this technology, allowing humans to experience Immersive Virtual Environment (IVE) through different sensory modalities, creating the feeling of ‘being there’ [315]. In this state, Lisa would truly believe and feel as though she is present in the computer-generated environment despite being physically located in a different space. Achieving a sense of presence is a fundamental goal in VR, as it enhances the immersive nature of the experience. Nowadays, users immerse themselves by using Head-Mounted Display (HMD), offering high-resolution stereoscopic visual and spatial auditory rendering. Driven by technological advancements and more affordable hardware systems, VR has experienced a constant rise over the last decade, with many companies and researchers contributing to its development. By 2024, there exist numerous commercial and even DIY¹ VR headsets. We have come a long way since VR’s historical roots in the 1960s, but yet, we have not achieved Evan Sutherland’s vision of the *Ultimate Display* [331]—a technology that can stimulate all the human senses. Despite the improvements of VR devices’ visual and auditory rendering capabilities, other senses, such as smell, taste, and touch, remain poorly supported. In particular, the absence of the sense of touch prevents VR from reaching its full potential because many of the envisioned use cases and applications, such as in entertainment, training, and therapy, often involve interacting with the environment [201]. Imagine if Lisa reaches out to touch, e.g., a virtual tool with her hands; however, with current VR systems, she would only find ‘thin air’, i.e., she

Virtual Reality may help address this challenge

The sense of touch is missing

¹ Hackaday webpage: <https://tinyurl.com/3yymp3p4>. Last accessed: Nov 1, 2024

*Absence of haptic
feedback*

would reach through the virtual tool. This implausible interaction can lead to semantic violation [265] and consequently disrupts the feeling of presence [310] and can even lead to Lisa experiencing VR sickness—a condition being extremely uncomfortable and usually leading to users immediately taking off the HMD [85]. Despite these consequences, not much has changed since the early days of physical controllers, providing basic controls, e.g., buttons and vibrotactile feedback, as a means for interacting with virtual content.

1.2 PROBLEM STATEMENT

*Why is this
important?*

In current consumer-grade VR systems, users use controllers with buttons and thumbsticks in their hands. All interactions with virtual objects are done through these controllers. As a consequence, object interactions in VR feel disembodied from their real-world counterpart because, e.g., objects' physical properties such as their shape, size, mass, or textures are missing, i.e., there is nothing to grab and hold. However, humans use haptic sensory information to make sense of their environment and their interactions within it [283]. Researchers found that adequately representing haptic sensations leads to greater realism and presence [153, 158], more natural human responses [310] and thus, enhances the overall VR experience. These advantages combined promote better skill transfer into the real world [166]. *So, how can we develop a VR system that provides appropriate haptic feedback for Lisa as she trains with virtual objects?*

*Active and Passive
haptic feedback*

Here, researchers have proposed the concept of Active Haptic Feedback (AHF), which uses computer-controlled actuation to produce the haptic stimulus corresponding to Lisa's desired interaction. These systems use sophisticated algorithms and physics simulations to compute the required force feedback [319]. A well-known example is the *PHANTOM haptic device* [234], a pen-shaped device that applies directed forces on the fingertip, allowing users to experience a diverse range of virtual objects and their characteristics. Despite these advantages, AHF approaches suffer from high mechanical complexity and are often bulky because of the actuators needed to produce human-scale forces—thus making them less suitable for everyday use. To overcome these limitations, Hinckley et al. [151] introduced Passive Haptic Feedback (PHF), which utilizes passive real-world objects, so-called proxies, as '*stand-ins*' for ideally multiple virtual objects [158, 256]. Generally, proxies aim to replicate the relevant physical properties of virtual objects given an interaction. They are simple and inexpensive and do not require actuators or complex algorithms to deliver haptic feedback. PHF has been shown to be an effective technique in various domains, but having a dedicated proxy for each virtual object in the IVE that resembles its physical properties quickly becomes impractical [304]. Because users would need a large number of proxies

that are highly similar to the virtual objects present in the IVE [256]. Moreover, this would result in them having to physically reconstruct the IVE with exact proxy replicas, which ultimately defeats the purpose of VR [158].

As a consequence, the endless range of possible haptic sensations, their complexity, the adaptability of devices to different situations, their ergonomics, cost and form factors have prevented the previously mentioned approaches from becoming viable options for everyday use. Hence, there is a clear need for inexpensive techniques that allow us to trigger or manipulate haptic sensations on demand.

This is where hand-based illusions come in. They enable VR designers to simulate haptic stimuli without ‘*real haptics*’—purely based on visual manipulation of the IVE and the interactions within it. Hand-based illusions in VR exploit humans’ perceptual inaccuracies [128], ‘*tricking*’ Lisa’s perception so that she believes that she experiences haptic feedback, even though physically she does not. The advantage is that these techniques are software-based and require only programmatic interventions but not complex hardware. This way, they can be easily applied to any new and existing VR application. Lécuyer et al. [210] showed that *Pseudo-Haptics*, one type of VR illusion, can create haptic effects similar to those of an equivalent physical device. These promising early empirical studies paved the way for developing many hand-based illusion techniques (see Figure 1.1). For example, in VR, researchers simulate haptic feedback for mid-air interactions in the complete absence of real haptics [317]. Other approaches combine illusions with proxies, e.g., to change their perceived physical properties such as shape [20], size [30], mass [288, 377] or texture [69], as well as combining them with active haptic devices to increase their haptic resolution [2, 323]. However, illusions in VR are by no means limited to only haptic effects. They can facilitate interactions [3] and manipulate the perception of users’ body [180] and space [235, 279, 329]. In summary, several prototypes and empirical studies demonstrated the effectiveness of illusions in VR.

However, they work by introducing an offset between what Lisa sees and what she feels. As a result, the techniques come with an inherent risk of her noticing them and, thus, may disrupt the immersive nature of the VR experience [208]. This is unlikely to happen with smaller offsets but becomes more likely with larger offsets, marking the bottleneck of using illusions in VR. On the one hand, larger offsets are required to achieve stronger perceptual effects, which in turn increase the risk of disrupting the immersive experience. Yet, very little is known about the extent to which illusions can be used while remaining unnoticeable, leading us to the following overarching objective driving this work:

*Perceptual illusions
in VR*

*Potential of illusion
techniques*

*Drawbacks of
illusion techniques*

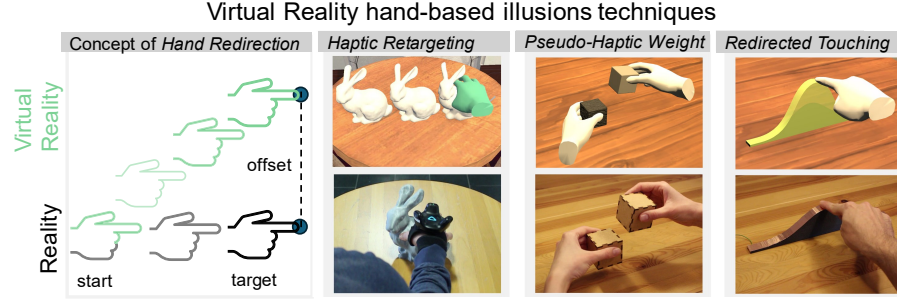


Figure 1.1: During hand movements under the influences of *Hand Redirection* [384], the position of a user’s virtual hand is gradually offset from the position of the real hand. This technique can be used to create a variety of effects. For example, the three famous hand-based illusion techniques displayed in our *Museum of VR Illusions* [102]. *Haptic Retargeting* [17] can create the illusion of three physical bunnies (left) by redirecting a user’s real hand to the same proxy bunny. *Pseudo-Haptic Weight* [288] can simulate weight difference between two equal cubes (middle) by vertically slowing down users’ real-world movements with the ‘heavier cube’. *Redirected Touching* [193] can simulate a larger physical shape by vertically offsetting the virtual hand during interaction with the proxy (right).

Exploring the boundaries of unnoticeable hand-based illusions to enhance haptics in VR.

1.3 RESEARCH QUESTIONS

To explore this general objective, we split our work into four parts, corresponding to the four overarching research questions outlined and discussed below.

Multimodal haptic
feedback through
proxies

NOVEL PROXY DESIGN Haptic proxies have been a central part of VR research in the past decade (further discussed in Chapter 2). Many approaches focus on adding hardware to existing VR controllers, allowing them to effectively render certain kinds of haptic feedback (e.g., the center of mass [382]). However, (1) approaches often focus only on one type of haptic feedback, and (2) they severely limit the type of interactions and rely on button controls to select and manipulate virtual objects inside IVEs. As a result, users can not leverage the dexterity of their hands and interact with virtual objects in direct and embodied ways [184, 349]. Therefore, we ask the question:

RQ1: How can we design proxies that can change their perceived kinesthetic and tactile properties?

NOVEL ILLUSION TECHNIQUES Proxy objects underlie the physical constraints of fabrication, materials, and energy management that cannot be easily overcome. Many haptic VR controllers and proxies,

including our approaches developed in (RQ1), suffer from this general limitation. However, illusion techniques can help to address these limitations as they can trick users' perception into experiencing, for instance, haptic feedback without the need of physically rendering it [227]. Yet, while many illusion techniques can change the haptic perception of a proxy, they come at a cost—users can often easily notice them. Thus, they can severely disrupt the immersive nature of the experience inside an IVE. Therefore, we looked into:

RQ2: How can we design unnoticeable hand-based illusions that expand the scope of proxy-based interactions?

INVESTIGATE CONTRIBUTING FACTORS While designing and evaluating our unnoticeable hand-based illusion techniques in RQ2, we identified factors that influence the unnoticeable offset. Moreover, a string of ongoing research in this field confirmed our observations and showed that the reported detection thresholds for the unnoticeable offset are subjected to change, depending on, e.g., task difficulty [83]. This is of great interest because the ever-growing body of illusion techniques that have been presented in the last decade is directly affected by this. A commonly used type of hand illusion that also builds the foundation for the techniques developed in RQ2 is called Hand Redirection (HR) [384]. The technique offsets the position of the real hand from the position of the virtual hand, and because users compensate for this offset, the system can redirect the real hand to a different position (see Figure 1.1). To study the robustness of our novel, unnoticeable hand-based illusion techniques and to understand how perceptual thresholds can generalize to other types of interactions or user groups, we investigate:

RQ3: Which factors related to the user, interaction, and proxy extend or limit the unnoticeable offset during hand-based illusions?

NOVEL METHOD TO CALIBRATE MAGNITUDE OF ILLUSIONS Our findings (RQ3) suggest a complex relationship between many factors and the amount of unnoticeable offset. Thus, seamless integration of illusions into VR experiences, which do not result in noticeably disrupting presence, remains challenging. Especially, the difference in individuals' sensitivity to offsets [144] greatly affects detectability of HR. As a result, it is necessary to investigate alternative ways to adapt the offsets to users and the VR experience to exploit their potential beyond controlled lab environments. This is relevant for the constantly growing field of illusion techniques, as it would allow them to become an effective tool in VR interaction and haptic design. Therefore, our last research question explores:

RQ4: How can we tailor unnoticeable hand redirection to individual users' perceptual boundaries?

Creating haptic sensations through unnoticeable visual offsets

Do illusions remain of a theoretical nature, or can they be used in practical VR settings?

Apply illusions beyond lab environments

1.4 METHOD & APPROACH

*General research
approach*

We explored our research questions through a set of established scientific methods from the fields of Human-Computer Interaction (HCI), VR, psychophysics, neuroscience, and psychology. The projects presented in this dissertation share the same underlying research approach, consisting of four phases: Conception, Implementation, Evaluation, and Analysis. Generally, we followed a waterfall model, subsequently running through the phases, despite phases one and two, which were done in an agile-like fashion, i.e., early prototypes and pilot studies helped us to decide whether the outcome from the conception phase needed to be refined in another iteration or it helped us to narrow down which concepts and ideas to follow. The following outlines the general methodology for each of the four phases. More details can be found in the corresponding sections.

*Iterative research
cycle*

Conception Phase. This phase consisted mainly of developing ideas based on the related work and/or previous projects. Once we identified a research gap, we developed early prototypes, proof-of-concept implementations, and sketches to illustrate and further refine ideas. Some concepts were also explored in small pilot studies to test their validity early on. In this phase, we also formulated hypotheses for the evaluation phase.

Implementation Phase. After establishing a promising concept, we developed prototypes, often consisting of hardware and software. For building the hardware prototypes in Chapter 3 (**RQ1**) and Chapter 4 (**RQ2**), we used fused deposition modeling (FDM) 3D printing, laser cutting, electrical circuits, and mechanical components. On the software side, we mainly used computer-aided design (CAD) software to design the 3D models, 3D engines to implement the IVE, and low-level microcontroller programming to control our prototypes (**Q1–4**). The commercial hardware stack included various VR systems, 3D tracking systems, eye trackers, and physiological sensors that can be found in the corresponding sections. Once we obtained a functional prototype, we tested our hypothesis.

*Lab studies are the
core of our
evaluation strategy*

Evaluation Phase. To address our research questions, we performed rigorous user studies and experiments to collect scientific data, helping us evaluate our concepts and prototypes. The core element of our evaluation strategy was controlled lab studies.

Lab studies. In our controlled lab studies with human subjects, we applied qualitative and quantitative methods. For example, we used quantitative measures such as task performance, Likert scales, and standardized and custom questionnaires (**Q1–4**). We also applied

qualitative strategies by conducting semi-structured interviews and using observations and field notes to better understand the experiences of participants with our prototypes (Q1–2). We also applied methods from the psychophysics (Q2–4) and neuroscience (Q4) domains in our experiments, which we will discuss further in Chapter 2.

Analysis Phase. After collecting qualitative and quantitative data, we analyze it using established scientific methods (e.g., thematic analysis [38] and/or statistical testing [74, 311]). We collected scientific evidence through rigorous hypothesis testing using traditionally frequentist but also Bayesian statistics (Q1–4). Furthermore, we applied linear dimensionality reduction techniques such as Principal Component Analysis (PCA) and Machine Learning (ML) to our data to gather more insights (RQ4). For our analysis, we used a combination of R², JASP³, Python⁴ and Matlab⁵ with a variety of existing toolboxes. The exact procedures are described in the corresponding sections.

APPROACH TO RESEARCH QUESTIONS In RQ1, we explore two novel approaches to proxy design that allow reconfiguration and dynamic shape-change to adapt to the constantly changing requirements, given the dynamic nature of IVEs. There exists a growing body of work on haptic devices (discussed in Chapter 2), presenting concepts and implementations that effectively provide haptic feedback for certain kinds of interactions. In Section 3.1.1, we conceptually combine many of the approaches presented in the literature by designing, fabricating, and implementing a block-based construction toolkit called *VoxelHap*. It allows users to construct handheld-sized proxy objects capable of providing tactile and kinesthetic haptic feedback. In Section 3.2, we specifically focus on changing users' tactile perception by combining dynamically changing proxies with visual texture overlays. The central idea surrounds the design of 3D printable metamaterials that can change their tactile properties, hardness and roughness upon compression. We investigate their potential as multipurpose proxies by purposely introducing mismatches between the haptic and visual sense and study multisensory integration in case of conflicting sensory information.

*VoxelHap toolkit
allows proxy
reconfiguration*

*Metamaterials for
shifting tactile
perception*

In RQ2, we investigate three novel hand-based illusions with proxies to increase their haptic rendering capabilities. In contrast to RQ1 and the majority of the related literature, we elicit haptic sensations without users noticing the visual manipulation, and thus not disrupting presence [314]. We focus on two simple but common user interface (UI) elements, rotation of a knob (see Section 4.1) and

² R webpage: <https://www.r-project.org/>. Last accessed: Nov 1, 2024

³ JASP webpage <https://jasp-stats.org/>. Last accessed: Nov 1, 2024

⁴ Python webpage: <https://www.python.org/>. Last accessed: Nov 1, 2024

⁵ Matlab webpage: <https://tinyurl.com/mtmwym8t>. Last accessed: Nov 1, 2024

*Three novel illusion
techniques*

translation of a slider (see Section 4.2). In addition, we present a technique that allows a single container proxy to provide haptics for virtual containers of varying diameters (see Section 5.1). As part of this research question, we consider not only passive proxies but also AHF devices that produce counterforces. The central idea is to obtain a more holistic understanding of the effects that unnoticeable illusions can create, e.g., how does a pseudo-haptic force compare to a physical force? In the second part of this chapter, we also begin exploring how aspects of the interaction affect the unnoticeable offset between the real and the virtual world, paving the way to **RQ3**.

*Factors that
influence
unnoticeable
illusions in VR*

In **RQ3**, we start our systematic analysis of potential factors that limit or extend the unnoticeable offsets. Although researchers in the field have already explored effects such as distraction levels [384] or task difficulty [83], there is still a lot of untouched ground. Therefore, we look at various factors related to the user, interaction, and proxy to contribute to a better understanding of how and to what extent designers can apply illusions to new and existing VR experiences. This includes (1) users' general sensitivity to visuo-haptic offsets, (2) how users perform interactions, (3) the physical properties of the proxy they use, and (4) users' body representations inside the IVE. Yet, there exists no framework for designers on how to apply hand-based illusions in VR that consider these factors.

*Novel method to
dynamically tailor
HR to individuals'
sensitivity*

In **RQ4**, we built on top of the results of **RQ3**, attempting to establish a novel methodology in the realm of illusion techniques that constantly monitors users' current sensitivity to visuo-haptic offsets and dynamically adjusts them. Our goal is to predict the amount of unnoticeable HR offset regardless of the interaction, illusion technique, or VR application. Inspired from previous work that detects errors and mismatches [118, 265] in VR, in Section 6.1 we investigate how humans physiologically respond to HR offsets above their Detection Threshold (DT). This includes Electrocardiogram (ECG), Eletrothema Activity (EDA), Respiratory Rate (RSP), Electroencephalogram (EEG), eye tracking, and interaction data. In (Section 6.2), we investigate whether our approach can be used to distinguish between interactions under the influence of HR Below, At and Above an individual's perceptual boundary.

1.5 CONTRIBUTIONS

The goal of this dissertation is to improve proxy-based haptic feedback for VR (**RQ1**) by applying unnoticeable hand-based illusions techniques (**RQ2**), understanding their application limits (**RQ3**) and tailoring them to individuals' perceptual boundaries (**RQ4**). To achieve this, we designed and implemented four research prototypes

and eight testbeds and conducted eleven lab studies with a total of 206 participants to evaluate our systems and techniques and collect empirical evidence for the theoretical contributions of this work.

First, we focus on developing reconfigurable and dynamic proxies to promote more versatile and adaptable proxy devices. Next, we look at inexpensive software-based techniques that allow designers to trigger haptic feedback purely based on visual manipulations while ensuring that users cannot notice them. This paves the way for our thorough investigation of potential factors that influence the detectability of such illusions. Finally, we investigate an alternative method to dynamically tailor hand-based illusions to individual perceptual boundaries. In this dissertation, we make six major contributions:

*Contributions
connect to each other*

C1: We design, fabricate, and study four novel haptic proxies that can render a wide range of kinesthetic and tactile haptic effects.

C2: We establish three novel undetectable hand-based illusion techniques for VR, *Pseudo-Haptic Resistance*, *Visuo-Haptic Translation* and *Visuo-Haptic Rotation*, through a series of perceptual experiments.

C3: We provide a set of design factors concerning user representation, interactions, and the proxy that should be considered when designing unnoticeable illusions for VR experiences.

*Core contributions
of this dissertation*

C4: We present a novel method that monitors users' interactions and their physiological body responses to predict how much offset can be used while remaining unnoticeable to an individual.

C5: We develop a range of use cases and application scenarios that showcase how our systems and techniques can be applied in IVEs.

C6: We open-source a large part of our work, including frameworks, experimental data, hardware, software, and two data sets.

These key contributions and their subcontributions can be divided into theoretical, design, and technical contributions, which we discuss below.

Theoretical Contributions. In this dissertation, we present several theoretical contributions related to the fields of HCI, VR, haptic proxy design, and perceptual illusions. In Chapter 3, we contribute to the domain of proxy design by presenting two novel concepts concerning the combination and reconfiguration of haptic proxies. In Section 3.1.1, we conceptually combine several haptic feedback

*Multipurpose
proxies for texture
perception*

modalities, pushing beyond single controllers and devices that support one type of haptic feedback (C1), and provide evidence in favor of the uncanny valley of haptics theory[29] . In Section 3.2, we are the first to combine the design of 3D-printed metamaterials with haptic proxies, and by overlaying visual textures, we demonstrate the potential of the approaches to influence tactile perception (C1). Our findings support the theory of multisensory integration [80] while demonstrating the limits of the visual-dominance phenomenon (C3).

*Hand-based illusions
for interface
elements*

In Chapter 4, we present two novel illusions with proxies based on HR, creating the illusions of resistance (*Pseudo-Haptic Resistance*) and allowing a single proxy slider to act as a stand-in for multiple virtual sliders of different lengths (*Visuo-Haptic Translation*) (C2). Our effects exploit the visual-dominance phenomenon, and by determining DTs through psychophysical methods, we are the first to ensure that our techniques (1) trigger haptic sensations and (2) remain unnoticeable to users (C3). Furthermore, we also provide a first multisensory integration model for *Pseudo-Haptic Resistance*, explaining the relationship between physical resistance and perceived ‘resistance’ created by visual manipulation (C2). In this way, we demonstrate that a pseudo-force is perceptually equal to a physical force [210].

*Understanding the
limits of hand-based
illusions*

The aspect of unnoticeability was further investigated in Chapter 5, where we contribute to theoretical knowledge about the robustness of hand-based illusions for several factors, e.g., the absences and presences of a full-body avatar or a proxy, a user’s movement trajectory, proxy type and mass, but also aspects related to users such as their proprioceptive sensitivity (C3). We also introduce our third unnoticeable illusion technique, *Pseudo-Haptic Rotation*, that allows a single physical container to provide haptics for multiple virtual containers of different diameters (C2). Our findings provide valuable insights about the factors that limit or extend the unnoticeable offset. Our experiments contribute towards a better understanding of the dominating factors leading to a semantic violation [265].

*Tailor hand-based
illusion to perceptual
boundaries*

In Chapter 6, we present a novel method to tailor a HR illusion to individual perceptual boundaries, i.e., DTs of different magnitudes. The method provides an alternative to psychophysical DT experiments and may be used in an on-the-fly fashion, regardless of the underlying VR experience (C4). To the best of our knowledge, this approach marks a significant step forward in the field of VR illusions that are designed to remain unnoticeable. Moreover, the method could potentially be applied beyond our hand-based illusion techniques developed and studied in this dissertation.

Design Contributions. We also make design-related contributions with respect to the VR systems, haptic devices, illusions, use cases, and applications.

Section 3.1.1 outlines the design of a block-based construction kit, *VoxelHap*, that enables users to build and reconfigure haptic proxies on demand (C1). To achieve this, we designed the Haptic Proxy Description Format (.hpdf), a semi-generalizable description format for proxy design, which includes a proxy’s kinematics, a functionality log, and a construction plan. A design tool that takes .hpdf as input and guides the user through the assembly process (C6). In addition, we designed example scenarios that illustrate the capabilities of the toolkit (C5). In Section 3.2, we contribute to the design of metamaterials (i.e., 3D models) that can dynamically adjust their physical properties with respect to hardness and roughness upon compression (C1). Furthermore, we studied how visual overlays in VR affect the perception of these designs by purposely introducing mismatches between the haptic and the visual sensory information. The results help VR designers to make informed decisions on how to effectively shift user’s tactile perception in the absence of, e.g., a proxy texture that matches the virtual texture.

*Support proxy
design and
construction process*

*Metamaterial
designs that change
roughness and
hardness*

Following up on the idea of purposely introducing mismatches between the haptic and the visual sensory modality, we designed three novel illusion techniques that designers can apply to new and existing VR experiences to improve their perceived haptic resolution (C2). Through our experiments, we determined the limits of how much offset can be used while remaining undetectable to users. These findings can inform VR designers on how to incorporate our results into their design process (C3). To facilitate this transfer process, we also present a range of use cases and applications that demonstrate the application of these techniques in IVES (C5).

Chapter 5 investigates the design variables that should be considered when including HR-based illusions into VR experiences that aim to remain unnoticeable. It also provides boundaries in the form of DTs for various interactions with and without proxies (C3).

*Understanding the
design of VR
illusions*

In Chapter 6, we present a novel method that allows researchers and designers to expose users to perceptually different magnitudes of HR. The robustness of this procedure is demonstrated in Section 6.2. Provided further validation, it would ultimately support VR designers in creating immersive experiences that utilize unnoticeable VR illusions that can be automatically tailored to individuals’ perceptual sensitivities (C4).

Technical Contributions. Many of our theoretical and design contributions ultimately led to technical contributions.

*Research prototypes
consist of hardware
and software*

For example, the *VoxelHap* hardware was realized by using FDM 3D printing, custom-designed mechanics, microelectronics, off-the-shelf actuators, and sensors. It was implemented using low- and high-level programming (C1). In addition, we designed a PCB to optimize the form factor; the schematics and Gerber files are open-source to ensure reproducibility and promote transparency (C6). Our technical evaluation of *VoxelHap*, together with the used cases and applications presented, provides a holistic assessment of its capabilities and limitations, especially relevant for VR designers and practitioners (C5). Although the fabrication of metamaterials using FDM 3D printing in Section 3.2 is not novel in itself, we demonstrate that our cell designs and form factor can still be fabricated using conventional FDM 3D printing (C1). We also open-source the 3D models of our metamaterial designs and the fabrication pipeline (C6).

*Common interface
elements knob and
slider*

In addition to the prototypes necessary to study the three novel hand-based illusion techniques in Chapter 4 and Section 5.1, we built a set of use cases and applications, demonstrating the technical feasibility, the broad application space, and the adaptability of the techniques (C5). The proxies for the two common interface elements, knob and slider, presented in Chapter 4 have sensing and actuation mechanisms, and by combining them with existing HR techniques, can render a variety of virtual object interactions and haptic sensations (C2). We implemented demo scenarios to educate the HCI/VR community and the general public on this topic, allowing them to experience these illusions first-hand that would otherwise remain inaccessible and abstract for the reader (C5). As part of this mission, we also developed the *Museum of Virtual Reality Illusions* (MoVRI) [102] in Section 7.1.5.2, which exhibits several illusion techniques for VR discussed in this dissertation (C5).

*Experiencing
illusions first hand*

Open-source toolkit

Chapter 5 and Chapter 6 are merely of theoretical and design nature and, therefore, did not yield a prototype or artifact besides the developed VR testbeds. However, to investigate our hypothesis and collect data within IVEs, we developed the open-source⁶ *VRQuestionnaireToolkit* [94] (discussed in Section 7.1.5.1), which allows us to collect experimental data without having to remove participants from the IVE (C6). Researchers highlighted the importance of collecting this data inside the IVE to not disrupt presence [299]. To this end, we still support and maintain the toolkit, and it has received considerable attention in the HCI/VR community.

⁶ Author's GitHub: <https://tinyurl.com/dzakk2bh>. Last accessed: Nov 1, 2024

Finally, in Chapter 6, we contribute two data sets with physiological features recorded while exposing participants to HR offsets of different perceptual magnitudes (**C6**). In addition, Section 6.2 contributes a multimodal classifier that can distinguish noticeable from unnoticeable HR and, thus, can predict whether an individual’s perceptual boundary for HR was exceeded (**C4**). The long-term vision for the classifier is that it substitutes a DT experiment, allowing constant monitoring and dynamic adaptation of the HR offset based on participants’ sensitivity to illusions.

Data sets ensure reproducibility

1.6 OVERVIEW

This dissertation is organized as follows.

Chapter 2 briefly introduces the field of VR and its history before diving into the basics of human perception. Next, we provide an overview of haptics in VR, followed by introducing relevant theories about perception and how VR illusions exploit them. In addition, we will discuss work related directly to our research questions.

Related literature

In Chapter 3, we investigate reconfigurable proxies with the goal of rendering haptic feedback for as many virtual objects as possible (**RQ1**). We focus on handheld-sized objects, presenting the *VoxelHap* toolkit with an end-to-end pipeline that allows users to construct and reconfigure haptic proxies on demand (Section 3.1.1). *VoxelHap* proxies offer kinesthetic and tactile haptic feedback, which we evaluated in two user studies. In addition, we performed a technical evaluation and developed a set of use cases and applications that demonstrate its capabilities. To improve its haptic resolution, we combine *VoxelHap* with existing hand-based illusion techniques to render a rich set of haptic sensations. Next, we focus on tactile perception and proxies that can change their tactile properties dynamically upon actuation to add another string to (**RQ1**) (Section 3.2). Specifically, we designed and 3D-printed metamaterial patterns that change their hardness and roughness properties when compressed. In a small lab experiment, we evaluated five metamaterial designs, each in four compression states, to investigate how humans perceive them through one-finger touch. We investigated the two most promising metamaterial patterns in a second study, where we added visual texture overlays in VR. We purposely introduced mismatches between the haptic and the visual texture to investigate how we can shift users’ perception and their tolerance to visuo-haptic conflicts.

Changing haptic properties of proxies

In Chapter 4, we follow up on the promising results of Section 3.2 by presenting two novel hand-based illusions with proxies (**RQ2**). For both techniques, we showcase potential use cases and applications,

*Illusions that
expand proxy-based
interactions*

demonstrating how they can expand the scope of proxy-based interactions for two common interface elements, knob and slider in VR. In Section 4.1, we present *Pseudo-Haptic Resistance*, which was created by applying visual offsets to the rotation of a knob. We established the technique through a three-stage psychophysical experiment. First, we investigate how much offset remains unnoticeable for users given this type of interaction, followed by how much offset is needed to create the sensation of resistance change. Finally, we study how our pseudo-force compares to a physical force and present a first model explaining their relationship. Section 4.2 demonstrates how a similar technique can be used to manipulate the perceived length of virtual sliders embodied by a single physical proxy slider. We determined to what extent this visual manipulation remains unnoticeable to users for varying slider lengths and in the presence of additional proprioception cues. As part of this experiment, we also looked at the potential side effect of exposing users to this type of illusion. Specifically, we consider factors such as task performance and investigate potential individual differences.

*Factors that affect
detectability of
illusions in VR*

This is further explored in Chapter 5, where we perform a systematic analysis of the factors that may affect the potential unnoticeable offset (**RQ3**). This chapter is divided into two parts: (1) Section 5.1 investigates factors related to the interaction with and without proxies, physical properties of those, aspects related to the interaction such as movement trajectory and participants' backgrounds related to their proprioceptive sensitivity, whereas (2) Section 5.2 explores potential effects related to the visual representation (i.e., the avatar) of the user within an IVE. In Section 5.1, we also present our third unnoticeable illusion technique called *Pseudo-Haptic Rotation*. From here, we discuss a set of guidelines that should help designers incorporate unnoticeable illusions into new and existing VR experiences.

*Tailor HR illusion to
perceptual
boundaries*

In Chapter 6, we focus on HR to investigate whether users' physiological responses (EEG, EDA, ECG, RSP) and interaction data can reveal if they are exposed to HR above their personal DT (**RQ4**). In Section 6.1, we study the potential of using this type of data to distinguish noticeable HR from no HR. First, we calibrate the noticeable HR offsets for each participant, followed by collecting data at their individual HR thresholds. Then, we perform an exploratory data analysis to identify the most promising modalities. In Section 6.2, we model users' ability to discriminate HR DTs at 25%, 50%, and 75% detectability. We investigate if we can distinguish the perceptually different magnitudes of HR using EEG, eye tracking, and movement data. Then, we train a multimodal classifier to predict whether users approach, reach, or exceed their personal DT.

Finally, Chapter 7 concludes this dissertation with a summary of our work, reflects and discusses our contributions, and provides recommendations for future research directions.

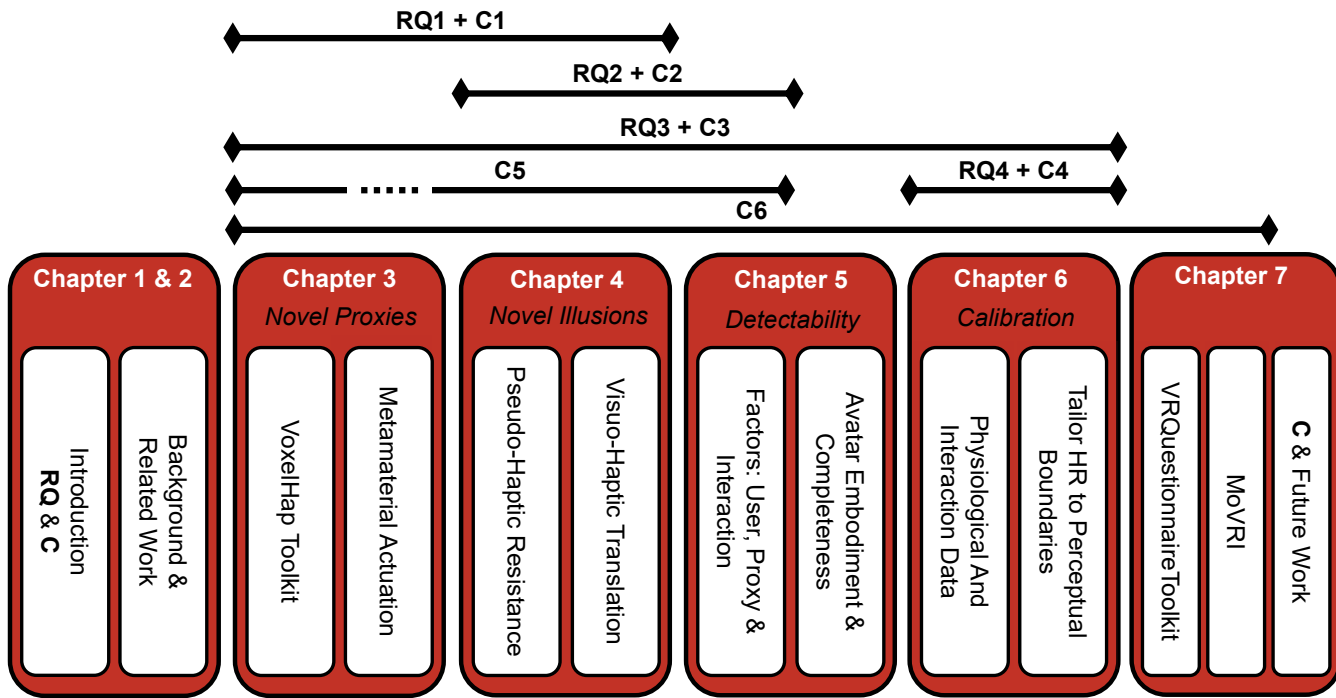


Figure 1.2: Overview of this dissertation, including chapters, research questions (RQ) and contributions (C).

BACKGROUND & RELATED WORK

2.1 THE *vision* OF VIRTUAL REALITY

The vision of entering and experiencing a different ‘*reality*’ attracts many people, especially in the science fiction domain. The list of movies and television shows that apply such storylines is long, reaching from living in a simulated world like in the “The Matrix”¹ movies to uploading human consciousness to an external world. These ideas resonate with people because they open up a wide range of exciting possibilities, such as visiting places of the past, seeing and interacting with family members who are long gone, overcoming disabilities, or even sustaining life beyond death. With VR, we are closer to this science fiction vision than ever before because VR allows Lisa to experience computer-generated digital environments through her sensory organs—seeing, touching, smelling, tasting, and hearing [249]. In the long-term vision, Lisa could physically interact with the artificial environment as it would be real [165]. Ivan Sutherland, a VR pioneer and creator of what would most regard as the world’s first VR system, describes this vision as “*The ultimate display would, of course, be a room within which the computer can control the existence of matter. A chair displayed in such a room would be good enough to sit in. Handcuffs displayed in such a room would be confining, and a bullet displayed in such a room would be fatal*” [331]. Putting the severe consequences of this vision aside, today’s VR is a powerful tool for many use cases across various fields, further described in Section 2.1.3. In the following section, we briefly discuss VR’s historical roots, establish a conceptual framework for VR, and outline current and future practical applications.

*VR can shape the
future of society*

2.1.1 A Brief History of Virtual Reality

“When anything new comes along, everyone, like a child discovering the world, thinks that they’ve invented it, but you scratch a little and you find a caveman scratching on a wall is creating virtual reality in a sense. What is new here is that more sophisticated instruments give you the power to do it more easily. Virtual reality is dreams.”—Morton Heilig [142] p.9.

Many think of VR as a novel technology; however, its roots go back to the 1850s, with stereoscopic viewers for still imagery being one of the first devices that allow users to experience the illusion of a 3D image. Stereoscopic devices use two images from slightly different

¹ IMBD webpage: <https://tinyurl.com/4zevcduv>. Last accessed: Nov 1, 2024




Figure 2.1: Left: Audience watches 3D film at the Telekinema on the South Bank in London during the Festival of Britain in 1951. (Source: The National Archives UK under OGL v1.0 license²). Middle: Morton Heilig's Sensorama Simulator [146]. (Source: History of Information³). Right: The Sword of Damocles by Sutherland. (extracted from [331]; ©1968 ACM).

Long history goes
back to 19th century

angles and present one to the left and one to the right eye. The brain combines the two pictures, resulting in the illusion of depth. In 1861, Oliver Wendell Holmes created the first handheld-size stereoscopic viewer, sharing many similarities with today's HMD (see Figure 2.2). However, instead of viewing still images, the technology evolved throughout the 20th century to film, triggering commercial interest and leading to the emergence of 3D Cinemas in the 1950s/60s (Figure 2.1: left). At the same time, multisensory viewing experiences made their first appearance. One prominent example is the *Sensorama Simulator* by Heilig [146] in 1962, where users could sit inside and view a stereoscopic 3D film with synchronized sound, smell, and haptic feedback, consisting of vibrations, tilting of the user's seat and facial wind (Figure 2.1: middle). While these systems allowed users to experience the illusion of 3D films, they were limited to one 2D viewing plane, like any other 2D screen-based device.

The first VR wave

Sutherland [331] changed this in 1968 by presenting a ceiling-mounted system, *Sword of Damocles*, which laid the foundation for modern-day HMDs (Figure 2.1: right). The system could track the user's head and display stereoscopic images rendered by a computer according to the user's head gaze. Thus, users could now be 'inside' a virtual environment and 'feel' surrounded by it. This fundamental change led to a development period with improvements to HMDs' Field of View (FoV), resolution, optics, and the introduction of hand-tracking gloves, allowing users to pick and place virtual objects through natural hand interactions. In 1984, the first commercial VR company VPL⁴ ("Virtual Programming Languages") appeared, selling HMDs, VR software, and tracking gloves. A few years later, a

² National Archives UK: <https://tinyurl.com/3s9eb4zx>. Last accessed: Nov 1, 2024

³ History of Information: <https://tinyurl.com/2jd92pmf>. Last accessed: Nov 1, 2024

⁴ Wikipedia: <https://tinyurl.com/y7kapnee>. Last accessed: Nov 1, 2024



Figure 2.2: Left: Brewster-type stereoscopic viewer from 1870. (Source: Alessandro Nassiri⁵ under [CC BY-SA 4.0](#)) Right: Modern Head-mounted Display. (Source: Wikimedia Commons⁶)

novel VR system Cave Automatic Virtual Environment (CAVE) was introduced by Cruz-Neira et al. [61], consisting of 3–6 wall stereo rear-projection screens with motion tracking (e.g., using IR markers). Users wore stereoscopic LCD shutter glasses that are synchronized with the computer, which generates an image for each of the user's eyes based on motion tracking, i.e., a user's viewpoint. This enabled a much larger FoV while users could move around freely inside the CAVE. In the following years, VR technology developed slowly and almost stopped because of hardware costs and low computational power. It was simply unaffordable and far from an everyday technology, so the public lost interest in it. As a result, funding was lacking, and many companies, such as VPL, faced bankruptcy, leading to a “VR winter” according to Jerald [165], who adopted the terminology previously used in the context of AI.

*Introduction of
CAVES*

Winter is coming

While military and scientific community research in VR continued over this period, advancements in computing would ultimately lead to a VR renaissance in the early 2010s. Oculus Rift's HMD Kickstarter⁷ project raised \$2.4 million, and a couple of years later, they released the first affordable consumer-grade VR systems, with many companies to follow. Mobile VR systems such as Google Cardboard⁸, introduced by Google I/O in 2014 or Samsung Gear VR⁹ in 2015, entered the market. They harnessed the increasing computational power of consumer smartphones, redirecting the public's attention to VR. Ever since the market has grown substantially with today's landscape of affordable VR systems at its all-time high. Many big tech companies such as Meta¹⁰, Apple¹¹, and Sony¹² offering HMDs at various price points. While CAVE has been updated to CAVE2

The second VR wave

5 Wikipedia: <https://tinyurl.com/36cubczz>. Last accessed: Nov 1, 2024

6 Wikimedia Commons: <https://tinyurl.com/mb53dtmb>. Last accessed: Nov 1, 2024

7 Kickstarter webpage: <https://tinyurl.com/ye29ya3z>. Last accessed: Nov 1, 2024

8 Google webpage: <https://arvr.google.com/cardboard/>. Last accessed: Nov 1, 2024

9 Samsung webpage: <https://tinyurl.com/2d283zhu>. Last accessed: Nov 1, 2024

10 Meta webpage: <https://www.meta.com/de/en/quest/>. Last accessed: Nov 1, 2024

11 Apple webpage: <https://tinyurl.com/mr3z66dn>. Last accessed: Nov 1, 2024

12 Playstation webpage: <https://tinyurl.com/42a6u73t>. Last accessed: Nov 1, 2024

13 Statista webpage: <https://tinyurl.com/2x52yu7x>. Last accessed: Nov 1, 2024



Figure 2.3: Left: A user standing inside a CAVE VR system [61]. (Source: Wikimedia Commons¹⁴; public domain). Right: A user wearing a modern HMD VR system. (Source: Image by Achin¹⁵ bm from Pixabay¹⁶)

*Predictions for
development of VR
technology*

[89], moving away from projection to LCD screens, HMDs are the dominant technology today, especially in the consumer market due to their affordability. Statista¹³ expects the Augmented Reality (AR) and VR market to grow at an annual rate of 8.97% from 2024 to 2029, suggesting that we have not even reached its peak.

VR technology looks back at a long history. However, research efforts have only recently accelerated with the introduction of consumer-ready VR devices, still making it a young and rapidly growing field of research. This has become apparent with the rising number of publications since the early 2010s, presenting novel hardware, software, design, and interaction concepts for VR.

2.1.2 *Conceptual Realm of Virtual Reality*

2.1.2.1 *Form of Realities*

*Classifying VR
systems*

Today, we look at a wide technology landscape that aims at seamlessly blending virtual and real-world content. As a result, the terminology surrounding VR systems can sometimes be confusing. Therefore, Milgram and Kishino [242] introduced the *Reality-Virtuality Continuum* depicted in Figure 2.4, covering a spectrum from the real environment to a fully virtual environment. Without getting too philosophical, the real environment is the world we live in, whereas the virtual environment is artificially generated and does not contain real-world content. Systems in between that use both real and virtual content are classified as Mixed Reality (MR) systems. MR can further be split into AR and Augmented Virtuality (AV). The latter describes virtual environments that are augmented with real-world content. This can be achieved by capturing a real-world model of, for example, a pet moving around and bringing it into the virtual envi-

¹⁴ Image of CAVE: <https://tinyurl.com/kj4rahkk>. Last accessed: Nov 1, 2024

¹⁵ Creator webpage: <https://tinyurl.com/y5wu42wy>. Last accessed: Nov 1, 2024

¹⁶ Pixabay webpage: <https://tinyurl.com/bdct9vb8>. Last accessed: Nov 1, 2024




Figure 2.4: *Reality-Virtuality Continuum* by and adapted from Milgram and Kishino [242].

ronment. On the other hand, AR uses the real world as a reference and augments it by, e.g., projecting virtual content onto physical objects in the environment. Ideally, this content seamlessly blends into the real world. Initially, the continuum focused on display technology, i.e., visual sense, but it can be generalized to other sensory modalities such as haptics shown by Jeon and Choi [163] (see Figure 2.10). This dissertation uses the conceptual space of the presented continuum to classify existing and novel approaches.

*Difference between
AR and AV*

2.1.2.2 Concepts of Immersion and Presence

“VR is about psychologically being in a place different than where one is physically located, where that place may be a replica of the real world or may be an imaginary world that does not exist and never could exist.”—Jason Jerald [165], p. 45.

IMMERSION Immersion refers to the technicality of the system and is, therefore, objective. It can be seen as the interface through which users experience the virtual environment with their sensory receptors. Presence is the subjective result of how users respond or experience different levels of immersion. Researchers describe six dimensions of immersion, which we briefly outline below [165, 310].

*Being immersed in
VR*

- **Extensiveness** describes the range of sensory modalities presented to the user, such as visual, auditory, and haptic.
- **Matching** is the congruence between the sensory modalities, e.g., visual rendering corresponds to head movements.
- **Surroundness** is the extent of panoramic cues, e.g., wide FoV and 3D audio.
- **Vividness** is the rendering quality, e.g., resolution of the HMD or audio bitrate.
- **Interactivity** refers to users being able to make changes to and in the environment.
- **Plot** is the narrative or the story that users experience in the artificial world.

Terminology around
immersion can be
misleading

It is important to note that VR literature often uses the terminology immersion to describe technical aspects of a VR system [310]. However, it is still an ongoing discussion whether it is appropriate to exclusively reserve the terminology immersion for this purpose because of conflicting definitions. Nilsson et al. [255] proposed a taxonomy outlining three critical concepts of immersion: (1) System immersion, a property of the system used to present the virtual world; (2) Narrative immersion as a perceptual response to that system; and (3) Challenge-based immersion as the subjective response to challenges presented. They take a holistic and interdisciplinary standpoint beyond the use of VR. To avoid ambiguity, we adopted their proposed terminology system immersion for the rest of this dissertation.

The feeling of being
there

PRESENCE To achieve that Lisa truly feels that she is in and a part of a virtual environment, we want her to experience the sense of presence—“*the feeling of being there*” according to Slater [310]. It is very challenging to describe this subjective feeling of being present in a virtual environment without having experienced it. Note to all the readers of this dissertation who have not had the chance to feel present in an IVE. Given the affordable VR technology today, we strongly recommend experiencing this once because it is truly fascinating. Broadly speaking, presence refers to users’ psychological and physiological state, where they do not perceive the technology [165]. The sense of presence can be very strong—leading to people even becoming anxious in IVEs. This makes it interesting for psychological exposure therapy, e.g., for treating fear of heights [112].

Place illusion
happens when SCs
align with users’
expectations

According to Slater [310], feeling present in a virtual environment can be achieved through *Place Illusion* (PI). PI relies on the Sensorimotor Contingencies (SCs) supported by the VR system. If the SCs in the virtual environment approximate those from the real world, it can create the illusion of being inside the virtual environment. Traditionally, this “*strong illusion of being in a place in spite of the sure knowledge that you are not there*” [310] has been referred to as telepresence, originating from the domain of teleoperating remote robots. Jerald [165] describes this as *The Illusions of Being in a Stable Spatial Place* (p.47), which is regarded as the most important part of presence. An example of PI is navigation through natural locomotion inside IVEs, i.e., walking or walking in place, which enhances presence compared to teleporting using a pointing and button device [350]. It is important to note that VR systems will most likely remain limited regarding the SCs that they can support, and thus, “(System) immersion provides the boundaries within which PI can occur” [310].

On the other hand, *Plausibility Illusion (Psi)* describes the perception of Lisa that what is happening in the virtual environment is really happening to her, “even though she knows for sure that it is not” [310]. It encompasses events that refer to the users over whom they have no direct control. For example, Pan and Slater [266] studied how shy males respond to a virtual woman smiling at the user and asking a question. The males smile back and respond to her question, even though they know no real person is there. The virtual environment appears to address and react to you, advancing the perceived reality of the virtual world. This is an example of *Psi*.

Events are really happening

The Illusion of Self-Embodiment is another dimension of presence, describing the user’s perception of having and identifying with a body inside the virtual environment [165]. Here, the appearance of the virtual body does not necessarily need to match a user’s physical body to result in a strong feeling of embodiment. Ultimately, this illusion allows Lisa to experience the same sensations in a virtual body as in her biological body [180]. The sense of embodiment has three main subcomponents: (1) the sense of *Self-Location*, describing the location of the virtual body in relation to her own body, (2) the sense of *Agency*, having direct control over her virtual body, i.e., movements of Lisa’s own body correspond to her virtual body movements, and (3) the sense of *Body Ownership*, which describes the perception of owning her virtual body, which can be achieved through multisensory integration [80]. For example, in the famous *Rubber Hand Illusion* by Botvinick and Cohen [35], visuo-tactile stimuli are provided to an artificial and the user’s own hand synchronously. After about one minute, users perceive the artificial hand to be their own, and they respond to stimuli provided to the artificial hand [172]. Experiencing ownership over the virtual body is a prerequisite for an embodiment illusion to occur [130]. Researchers have determined many variables that positively affect the sense of embodiment towards a virtual body (i.e., an avatar). For example, the completeness [84] or the visual fidelity [136] of a virtual avatar can increase the sense of embodiment. Despite the increasing sense of embodiment, full-body avatars have also been found to promote more realistic responses to stimuli [124], demonstrating the strength of the illusion.

The feeling of our bodies inside an IVE

Ownership is crucial for embodiment

The Illusion of Social Presence describes the perception that Lisa can communicate with others [165]. For example, people with a fear of public speaking showed a significant increase in signs of anxiety when giving a presentation in front of a virtual audience compared to an empty room [312]. While social presence focuses on the perception and the quality of the virtual avatars, the illusion of *Co-Presence* focuses more on the psychological interaction within

Meeting other humans in virtual environments

the group and thereby has two dimensions [46]. The subjective feeling of perceiving others as well as being perceived by them [313]. Often, the two concepts of social presence and co-presence are used interchangeably in the literature, but they provide slightly different angles on the user's perception of being with others (computer- or user-controlled) [165].

*Can I touch this?
The role of haptics.*

The Illusion of Physical Interaction is arguably the most important in the scope of this dissertation. Imagine Lisa is reaching out to grasp the virtual tool; however, she only finds 'thin air', leading to breaks in presence [314], because she expected a touch sensation upon making contact with the object. Ideally, the haptic sensation matches the user's expectation; however, it does not necessarily need to be realistic [165]. In Chapter 3, we look at ways how we can develop proxy objects that can provide realistic haptic sensations for a variety of objects and interactions. Chapter 4 and Section 5.1 present three techniques that rely on visuo-haptic offsets to overcome some of proxies' general limitations. These parts focus on not disrupting the illusions of physical interaction to maintain high levels of presence.

*How to measure
breaks in presence?*

The feeling of being present in a virtual environment is highly subjective, but it is often used to assess VR systems quality, allowing researchers and designers to compare them to each other. However, measuring subjective feelings is challenging, and researchers argue that it might only be possible to measure breaks in presence rather than presence itself. Thus, the goal for any VR system is to minimize the risk of a BIP as much as possible. Standard methods to measure presence have been questionnaires such as presence questionnaire (PQ) [365], slater-usoh-steed (SUS) questionnaire [351] or igroup presence questionnaire (IPQ) [296], counting BIP [314], physiological responses such as Skin Conductance Response (SCR) or heart rate [241], multi-modal matching method [31], or sentiment analysis [22].

2.1.2.3 *Stimulating The Senses*

VR systems stimulate Lisa's senses through multisensory, i.e., visual, auditory, haptic, olfactory, and gustatory feedback. Below, we provide an overview of the current technology landscape.

Seeing is believing

VISUAL When people think of VR systems, they associate it with stimulating the visual sense. One cost-effective way to do this is through a *Desktop VR* system, which uses a regular 2D monitor to display a 3D computer-generated environment. Users can move around in these environments using input devices such as mouse and keyboard, but they will not perceive stereoscopic 3D vision. However, by using shutter glasses, one can create the illusion of depth. *Fish Tank VR* systems enhance this by adding 3-Degrees of Freedom (DoF)

head gaze tracking, rendering the environment according to the user's viewpoint [356]. This provides a much more natural way to engage with the 3D environment than with traditional input devices; however, they suffer from limited FoV. Humans have about 200° horizontal and 135° vertical FoV when looking straight ahead, which these systems do not support [165]. The CAVES mentioned earlier solve this problem by providing a wider FoV and allowing greater freedom of movement in an approximately 3 m³ physical environment [61]. To this day, cost and space requirements severely limit the application domains of CAVES, which are used for specific exhibitions or research purposes (see Figure 2.3: left). The best compromise between form factor, cost, and space is (HMDs) such as HTC Vive¹⁷ or Meta Quest¹⁰. Today's HMDs offer high-resolution stereoscopic vision with 6-DoF tracking, but still cannot cover humans' full FoV. VR technology has significantly advanced in the last decade, producing HMDs with different capabilities designed for various use cases and applications. For example, the HTC ecosystem relies on outside-in tracking stations (i.e., base stations) that emit infrared to determine the position and orientation of the HMD, the controllers, and the trackers that can be attached to objects or the user's body. On the other hand, Meta Quest uses inside-out tracking and, thus, does not require any external tracking stations. Hence, it can only track objects and the environment in front of the headset through the built-in cameras, but it is a very portable device. High refresh rates of display and tracking ensure strong SCs, which ultimately leads to high levels of system immersion. Additionally, recent advancements in computer vision enabled high-precision hand and finger, face, and even body tracking that can realistically render users in IVEs, facilitating a presence illusion.

Moderns HMDs are powerful, enabling realistic visual renderings

Inside-out vs. Outside-in tracking

AUDITORY Sound is an essential aspect of immersive VR experiences, improving the sense of depth, space, and realism. There are several audio techniques, including mono, stereo, Dolby surround, and 3D or spatial audio. They can be experienced using speakers or headphones, which is the current standard in VR systems. 3D audio is often superior as it provides information about the sound source's location on the horizontal and vertical plane, as well as information about the distance to the sound source. Thus, it can enhance the VR experience by improving users' spatial awareness. This can be achieved using approaches such as binaural cues, head-related transfer function, head movement, and reverberation [41]. Today's VR systems feature high refresh rates, allowing for synchronized audio feedback during user interactions within an IVE. For example, when Lisa touches an object, the contact sound can provide cues about its properties, such as its material. This feedback contributes to a more intuitive

Can you hear it?

¹⁷ Vive webpage: <https://www.vive.com/de/>. Last accessed: Nov 1, 2024

and responsive experience, enabling her to understand the impact of her actions. Finally, auditory feedback is also essential for improving accessibility for visually impaired people, making VR more inclusive.

*Haptics is at the
very core of this
dissertation*

HAPTIC By adding tactile and kinesthetic haptic feedback to IVEs, they become more engaging, realistic, and interactive. However, to this day, VR systems mainly rely on controller-based devices with buttons and joysticks. These devices use vibration or rumble motors to render tactile feedback for interactions within the IVE. However, the simple nature of vibrations, which can only be changed in their duration, amplitude, and frequency, often results in ambiguous mappings between interaction and haptic feedback. The next generation of actuators, which can be found in the PlayStation's DualSense®¹⁸ controller uses linear resonant actuators or voice-coil actuators that can produce linear directional forces. This step forward allowed for much more convincing haptic feedback for certain kinds of interactions with controllers. Nevertheless, considering the dexterity of our hands and the richness of interactions, it becomes clear that a controller with embedded actuators cannot render the wide range of haptic sensations known from physical environments. Another commercially available type of device is haptic gloves. For example, MANUS Gloves¹⁹ use vibration motors on each finger, providing tactile stimuli upon hand interactions within the virtual environment. SenseGloves²⁰ take this further by exerting forces on the user's fingers corresponding to the interaction inside IVEs, for example, blocking the fingers when grasping an apple. Although not in the scope of this dissertation, haptic suits such as TESLASUIT²¹ or TactSuit²² enable full-body haptics. As rendering haptic feedback for hand interactions in VR is one of the main objectives of this dissertation, we use Section 2.3 to elaborate further on the landscape of haptic technology and concepts concerning this area.

*Very few commercial
companies exist*

OLFACTORY According to Merriam Webster²³, olfaction is defined as the sense of smell, which is perceived by the stimulation of the receptors in the nose. Scents can be delivered to the nose through olfactory displays, usually consisting of stocked components of odors that are vaporized using airflow, heat, or atomization. The individual odor components are then blended and presented to the user through a ubiquitous, wearable, or handheld display. A recent survey article by Tewell and Ranasinghe [341] provides a comprehensive overview of the current state of the art in olfactory display research for VR. Hariri

What about smell?

¹⁸ DualSense® webpage: <https://tinyurl.com/4653nadf>. Last accessed: Nov 1, 2024

¹⁹ MANUS webpage: <https://www.manus-meta.com/>. Last accessed: Nov 1, 2024

²⁰ SenseGloves webpage: <https://www.senseglove.com/>. Last accessed: Nov 1, 2024

²¹ Teslasuit webpage: <https://tinyurl.com/msa9ftzp>. Last accessed: Nov 1, 2024

²² bHaptics webpage: <https://www.bhaptics.com/>. Last accessed: Nov 1, 2024

²³ Merriam-Webster: <https://tinyurl.com/4jndcfzn>. Last accessed: Nov 1, 2024



Figure 2.5: Left: Thermal device that actuates the taste receptors of the tongue, creating the sensation of sweetness. (Source: extracted from Karunanayaka et al. [176] ©2018). Middle: Chemical modulators can alter users' taste perception. (Source: extracted from Brooks et al. [43] ©2023). Right: Electrical stimulation device that produces stereo-smell. (Source: extracted and cropped from Brooks et al. [44] ©2021).

et al. [143] proposed digitizing the sense of smell by directly stimulating the nasal conchae through weak electrical stimulation with varying frequencies and currents. Yet, besides early explorative studies describing a range of difficulties and a wide range of perceived sensations, no working prototype has been presented. While replicating odors might be challenging to achieve, Brooks et al. [44] demonstrate that electrical stimulation of the trigeminal nerve can be used for stereo-smell, i.e., communicating the direction or location of an odor (see Figure 2.5: right). However, commercially available devices are still hard to find, with Olorama²⁴ and OVR Technology²⁵ being the most known companies that target professional use cases and applications but not everyday consumers.

Challenging to digitalize smell

GUSTATORY Simulating the perception of taste is similar to olfactory, still in its early days. Karunanayaka et al. [176] developed a thermal device that actuates the taste receptors of the tongue, which creates the sensation of sweetness (see Figure 2.5: left). Ranasinghe [277], a pioneer in gustatory HCI research and co-authors, presented several systems relying on eletrotactile simulation of the tongue, creating sour, bitter, salty, and sweet sensations. They also introduced the digital flavor synthesizing device, which applies electrical and thermal stimulation as well as smell sensations to simulate minty, spicy, and lemony flavors [278]. Brooks et al. [43] propose using chemical modulators in the mouth before eating to alter users' taste perception (see Figure 2.5: middle). The modulators temporarily change the response of taste receptors, selectively suppressing salty, umami, sweet, or bitter and even transforming sour into sweet. These exciting and promising strings of research for simulating or redirecting taste sensations in VR demonstrate the potential of gustatory feedback for virtual and augmented environments. Much has happened in recent years; however, to this day, no commercially available device exists.

Maybe even taste?

Research still in its early days

²⁴ Olorama webpage: <https://tinyurl.com/5n8s267w>. Last accessed: Nov 1, 2024

²⁵ OVR webpage: <https://ovrtechnology.com/>. Last accessed: Nov 1, 2024

2.1.3 *Current and Envisioned Applications*

*VR has an impact on
various fields*

VR can provide unique experiences for users where they feel a strong sense of presence. This enables VR to create experiences that go far beyond what any other technology is capable of. Since there is not really a limit to what can be visualized inside an IVE, users can experience hypothetical scenarios or situations of the past, which is impossible in the real world, safely practice tasks in dangerous environments, or travel and experience remote places that require too many resources to visit. There exists a large body of diverse use cases and applications, ranging from simulation, training, communication, arts, and design to healthcare and entertainment [201].

*Training the
unexpected*

Simulation and training is one of VR's most promising use cases by trying to replicate the real world as much as possible. Here, Lisa can learn new skills in a safe and responsible way under realistic conditions—made possible by the presence illusion. Xie et al. [369] published a review article on skill training in VR, covering domains such as medical, military, transportation, workforce, interpersonal skills, or training first responders.

Overcoming distance

*Connecting with
others*

Communication and collaboration is another branch of VR application that gained significant attention during the COVID-19 pandemic²⁶ and the announcement of the Metaverse²⁷. VR has always been a promising technology for meeting, playing, and interacting in a shared environment. With the absence of real-world connections, its potential, especially to create the illusion of social and co-presence, became even more obvious to the general public. People started connecting in virtual environments such as VRChat²⁸ to cope with the situation. The author of this dissertation attended many conferences in remote settings, e.g., the 2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)²⁹ happened inside a virtual environment VirBELA³⁰. Fashion week³¹, musical festivals³², and other events moved to the Metaverse.

Professional applications in **arts and design** include designing architecture and products. For example, the computer-aided design (CAD) software Rhino3D³³, which we used to model many of the

²⁶ WHO webpage: <https://tinyurl.com/24jaym69>. Last accessed: Nov 1, 2024

²⁷ Meta webpage: <https://about.meta.com/en/metaverse>. Last accessed: Nov 1, 2024

²⁸ VRChat webpage: <https://hello.vrchat.com/>. Last accessed: Nov 1, 2024

²⁹ ISMAR conference: <https://ismar2020.ismar.net/>. Last accessed: Nov 1, 2024

³⁰ VirBELA webpage: <https://www.virbela.com/>. Last accessed: Nov 1, 2024

³¹ Vogue webpage: <https://tinyurl.com/bdepwckh>. Last accessed: Nov 1, 2024

³² Tomorrowland festival webpage: <https://tinyurl.com/4vuzkmhj>. Last accessed: Nov 1, 2024

³³ Rhino3D webpage: <https://www.rhino3d.com/>. Last accessed: Nov 1, 2024

³⁴ Rhino3D GitHub page: <https://tinyurl.com/47fkdyndv>. Last accessed: Nov 1, 2024

3D objects in this dissertation, has a VR plugin³⁴ to facilitate working with 3D objects. Moreover, VR is a powerful tool for information and visualization [79] and immersive analytics [111]. Companies offer immersive VR applications to help customers plan their new kitchen or living room³⁵. Car manufacturers provide virtual showrooms that facilitate the configuration of a car³⁶. The list of VR supporting applications covers many industry sectors.

Being immersed in content

Healthcare can greatly benefit from VR, treating anxiety, PTSD, and phobias³⁷ [112]. Other examples include supporting the physical rehabilitation of stroke patients who can practice motor skills and coordination³⁸ or distracting patients from pain by engaging in virtual environments. Due to VR's customizability, treatments can be tailored to individual needs to improve the therapy outcome. Finally, VR has also been proposed as a tool for remote diagnosis of neurodegenerative diseases such as Parkinson [264].

Supporting medical recovery

In recent years, we have seen several industry sectors developing VR technology for their use cases; however, **entertainment** is, to this day, arguably still the most significant industry investing in the development of VR technology. Here, VR gaming marks the forefront with games such as the infamous Beat Saber³⁹ attracting avid VR gamers. We have seen VR-enhanced rollercoasters⁴⁰, exhibitions⁴¹, and VR arcades⁴². VR movies are gaining popularity and enabling new ways of experiencing content. For example, in "The Soloist VR"⁴³ on Oculus TV, users can experience Alex Honnold's free solo climbs of challenging mountain faces from a first-person view. In the travel industry⁴⁴, VR can be used to capture the beauty and scenery of destinations around the world. While this industry wants to promote travel, the author of this dissertation sees this as a chance to (1) reduce traveling by visiting virtual replicas without any crowds and environmental impact and (2) allow people with disabilities or people who cannot afford traveling to experience the uniqueness and diversity of the world in their living room. This vision is inspired by VR museum tours⁴⁵ that already exist today, allowing people to virtually enter, explore, and experience the most famous museums of our modern world.

Explore, engage, and play!

Novel ways of consuming content

35 VR kitchen planner: <https://tinyurl.com/2u77sv5d>. Last accessed: Nov 1, 2024

36 VR car planner: <https://tinyurl.com/mt9mc9ck>. Last accessed: Nov 1, 2024

37 PsyTech webpage: <https://psytechvr.com/>. Last accessed: Nov 1, 2024

38 Cureosity webpage: <https://www.cureosity.com/>. Last accessed: Nov 1, 2024

39 Beatsaber game webpage: <https://www.beatsaber.com/>. Last accessed: Nov 1, 2024

40 Europa-Park webpage: <https://tinyurl.com/e9pe6ew3>. Last accessed: Nov 1, 2024

41 Wonderland webpage: <https://tinyurl.com/2ethn3d>. Last accessed: Nov 1, 2024

42 Virtual area games: <https://virtual-area.de/>. Last accessed: Nov 1, 2024

43 VR movie "The Soloist" : <https://thesoloist-VR.com/>. Last accessed: Nov 1, 2024

44 Article VR tourism: <https://tinyurl.com/59zrc87h>. Last accessed: Nov 1, 2024

45 Article VR museums: <https://tinyurl.com/mr46hycc>. Last accessed: Nov 1, 2024

2.2 HUMAN PERCEPTION

To design effective VR systems, we need to understand how humans perceive their environment through their senses. However, this is still not fully understood, and different models exist that explain how perception works. In the following section, we introduce the most relevant work that helps us to explain the sensory processes concerning the topics in this dissertation.

2.2.1 *Perceiving our Environment*

*Perception combines
sensory inputs with
previous experiences*

*Our senses help us
to understand our
environment*

Human perception is the process through which we interpret and make sense of sensory information from the (virtual) environment. It involves multiple steps, from detecting stimuli to organizing and interpreting them. To establish a common ground for the remaining section, we need to distinguish between Sensation and Perception [165]. Sensation is the process where sensory organs (e.g., eyes and ears) detect external stimuli (light and sound) and convert them into neural signals. Through transduction, these stimuli are converted into neural signals and transmitted to the brain. The brain then processes this information, integrating it with past experiences and knowledge, and finally interprets it. Thus, perception is the brain's interpretation of these signals, which add meaning and context, forming a subjective experience of the environment. Traditionally, we consider the five human senses: seeing, hearing, touching, tasting, and smelling [249] (p.694 and onwards). They are complemented by the proprioceptive sense, responsible for knowing where our own body parts are without looking at them, and the vestibular sense, often referred to as the sense of balance. In the context of this dissertation, we focus on stimulating the visual, touch, and proprioceptive senses to create compelling VR experiences.

*Actually seeing vs.
expecting to see a cat*

Literature states two types of perceptual processing, bottom-up and top-down [165]. In bottom-up processing, perception starts with sensation, which the brain then combines to form a complete picture. For example, Lisa might see shapes and colors first, then recognize them as a cat. In contrast, top-down processing involves using prior knowledge, expectations, and experience to interpret sensory information. For example, if Lisa expects to see a cat in the driveway, she is more likely to perceive an ambiguous shape as a cat. VR uses continuous bottom-up stimuli that are convincing enough to create a presence illusion [310], overwriting top-down data, i.e., *"I am just wearing an HMD"*. However, once presence is achieved, the relevance of top-down processing increases since stimuli become more complex [165]. When Lisa starts exploring the environment, e.g., by turning her head, she issues a motor command, sending an efferent signal from her central

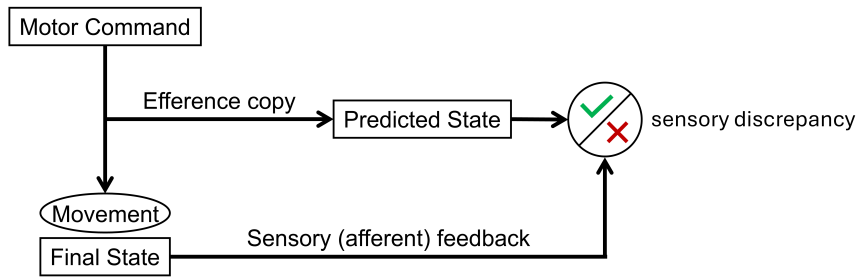


Figure 2.6: Sensorimotor Contingencies Model by and adapted from Gonzalez-Franco and Lanier [128] under CC BY 4.0.

nervous system to the muscles. At the same time, an efference copy containing the expected multisensory result of the predicted state is generated. This is then compared to the multisensory afferent input sensations (e.g., vision, touch, proprioception). If the expected state in Lisa's brain matches the information that arrives through the sensory afferent, then Lisa experiences a strong VR illusion according to Gonzalez-Franco and Lanier [128]. Thus, maintaining a high state of presence in a virtual environment requires a continuous SCs loop, such as depicted in Figure 2.6. Otherwise, a large enough sensory discrepancy between afferent multisensory inputs and the efferent predicted state can disrupt the experience. An important aspect of SCs is that they can be learned/updated, depending on, e.g., system immersion. For example, Lisa experiences a virtual environment through a head-tracked HMD with limited FoV. Usually, Lisa could quickly learn the effect of head movements on her visual perception, i.e., due to the limited FoV, she needs to turn her head further [310]. In the following, we will outline the primary perceptual mechanism for the visual and haptics (tactile and proprioceptive) as they are most relevant to this dissertation.

Comparing sensory inputs and predictions

How does VR achieve the feeling of being present?

2.2.2 Visual Perception

Visual perception begins when light falls (i.e., photons) on the photoreceptor in the eye [60]. It passes through the cornea, pupil, and lens, where it is focused onto the retina. The retina contains photoreceptor cells, including rods, primarily responsible for seeing during low illumination, and cones, responsible for colored and detailed vision in high illumination conditions. As described above, these photoreceptors convert light into electrical signals via transduction, which are then transmitted through the optic nerve to the brain. The process of why we perceive depth and three-dimensional vision is complex and contains many puzzle pieces. Generally, we use a combination of binocular and monocular depth cues that allow the brain to construct a sense of spatial depth [121]. Stereopsis arises from binocular vision, where each eye captures a slightly different

Different parts of the eye are responsible for seeing color and details

Why do we see in
3D?

image due to its horizontal separation. The brain processes these disparities, calculating depth and distance. This is also why most modern HMDs use two separate displays, rendering one image from a slightly different angle for each eye. Key monocular cues, such as relative size and interposition of objects, texture, shadows, and motion parallax, contribute to depth perception, offering important spatial information even when viewed with only one eye.

Visual perception differs in the center from the peripheral [165]. Most notably, central vision has high acuity, is color sensitive, and is optimized for daytime conditions, whereas peripheral vision is overall less sensitive to color and slow motions but highly sensitive to light in darker conditions and fast motions. This is important because some of the visual illusions outlined in the next section exploit these effects to trick human perception.

2.2.2.1 Visual Illusions

“What is reality anyway? Just a collective hunch!”—Lily Tomlin.

Tricking our
perception into
making errors

Our brains constantly try to make sense of our environment, inferring assumptions and expectations from past experiences. As a result, it is impossible to perceive a purely objective reality. Despite the feeling of “knowing the truth” [165], it is pretty easy to trick our perception into making errors. Below, we display three different visual illusions to illustrate how we can foul our perception. There exist many other types of illusions concerning the perception of depth, size, and motion. In VR, we can ‘play’ with the assumptions and expectations of our perception to create novel experiences and effects.

All lines have the
same length

2D ILLUSIONS A classical example: *Which of the two lines is longer?* The surprising answer is both have the same length (see Figure 2.7: middle). This phenomenon is known as the *Müller-Lyer illusion* [251], showcasing how the brain can misinterpret depth cues and size consistency. To this day, there exists no consensus about why this illusion works so well, and different explanations have been provided. One early explanation is that the brain interprets the arrowheads “< >” as objects that are closer, e.g., a convex corner of a room. In contrast, the “> <” suggests objects that are far away, such as the concave corner of a room, which distort the brain’s interpretation of the line length. More recent evidence suggests that the visual system picks up on depth cues, attempting to extract 3D information and, by doing so, misinterprets them. Nevertheless, most likely, a combination of different perceptual errors leads to an overestimation of the line length with “> <” arrows, even though both lines are of equal length.




Figure 2.7: Examples of optical illusions. Left: Expanding Hole illusion [202]. (Source: extracted from Laeng et al. [202] under CC BY 4.0). Middle: Müller-Lyer illusion [251]. (Source: extracted and cropped from Wikipedia⁴⁶ under CC BY-SA 3.0). Right: Kanizsa Triangle illusion [175]. (Source: extracted from Wikipedia⁴⁷ under CC BY-SA 3.0).

BOUNDARY COMPLETION ILLUSIONS Looking at this image, most people see a triangle that does not exist (see Figure 2.7: right). This phenomenon is rooted in the brain's tendency to "*fill in the gaps*" [165], according to the Gestalt law of closure and good continuation [324]. The brain constructs the invisible edges or contours that are not physically present in the image. They are inferred based on the arrangement of surrounding elements, allowing us to perceive objects as whole or continuous. For example, in the famous *Kanizsa Triangle illusion* [175], the three Pac-Man-like shapes positioned at the corners lead the brain to perceive a complete triangle, even though the triangle's edges do not exist. The brain also uses prior knowledge and expectations to interpret incomplete images according to top-down processing.

The triangle's edges do not exist

MOTION ILLUSIONS This type of illusion, which Laeng et al. [202] coined "*expanding hole*", combines dynamic sensations of motion and a gradually expanding central region, despite being a static image. This works because the brain processes visual information with slight delays, compensating for this lag by predicting how and where objects will move. Motion illusions such as expanding hole exploit this prediction system by presenting patterns or contrasts that suggest motion, even when none is present. It is complemented by high contrast (black and white) and an alternating color pattern that is known to confuse the brain's motion detection system. This triggers involuntary eye movements (microsaccades) interacting with these patterns or repetitive textures in an image, causing the brain to interpret them as moving. Note that the motion illusion becomes stronger when the image is viewed full-screen on a larger monitor because peripheral vision is more sensitive to these motion cues.

The image does not move

⁴⁶ Wikipedia webpage: <https://tinyurl.com/2s3fhun8>. Last accessed: Nov 1, 2024

⁴⁷ Wikipedia webpage: <https://tinyurl.com/5n73uxtU>. Last accessed: Nov 1, 2024

2.2.3 Haptic Perception

*Haptics is essential
in our everyday life*

Can you 'grasp' it?

Haptic perception refers to how humans interpret and understand their environment through touch. It involves integrating multisensory information from the skin, muscles, tendons, and joints to perceive properties such as texture, shape, temperature, pressure, and weight [211]. Haptic perception develops in our early childhood and is essential for fine motor skills [39]. Thus, it plays a crucial role in our everyday lives, e.g., in grasping and manipulating objects. Grasping is a fundamental and necessary type of interaction to assess, use, and leverage the full potential of objects in our environment. For instance, there are various ways to grasp a simple mug [106], and grasping becomes even more complex when looking at interactions such as opening a bottle, where humans seamlessly transition between several grasping types [62]—demonstrating the dynamics and unpredictability of these interactions. Humans choose the correct grasping type based on the underlying task requirements [45, 62, 105] and objects' characteristics [106] (particularly, the shape of the object [62]). These variables remain entangled and can, therefore, only be considered holistically. Cutkosky [62] and Feix et al. [107] proposed grasping taxonomies that broadly distinguish between power (intermediate) and precision grasps (see Figure 2.8). As the name suggests, power grasps are used to manipulate heavier/larger objects or when dexterity is secondary. On the other hand, precision grasps are primarily for fine-grained manipulations. Hence, different muscle groups are involved when changing or adjusting the grasping type [325]. These kinematic differences build the foundation for our investigations about the robustness of hand-based illusions in Chapter 5. Generally, haptic perception can be divided into two broader categories, tactile and kinesthetic, which we discuss below.

2.2.3.1 Tactile Perception

*Skin receptors allow
us to sense tactile
properties*

Tactile perception involves detecting and interpreting sensory information through the skin, allowing the recognition of touch, pressure, texture, temperature, and pain. Mechanoreceptors sense physical pressure or vibration when the skin contacts an object or surface. These include receptors like Meissner's corpuscles (for light touch and low-frequency vibrations), Pacinian corpuscles (for high-frequency vibrations), Merkel cells (for pressure and coarse texture), and Ruffini endings (for skin stretch) [211, 249]. Thermoreceptors respond to changes in temperature in the range of 5°C to 45°C, and nociceptors detect pain or harmful stimuli such as heat, pressure, or chemicals. The three receptors convert mechanical, thermal, or noxious stimuli into electrical signals through the spinal cord to the primary somatosensory cortex in the brain via the somatosensory pathways [249]. This area organizes sensory input according to the

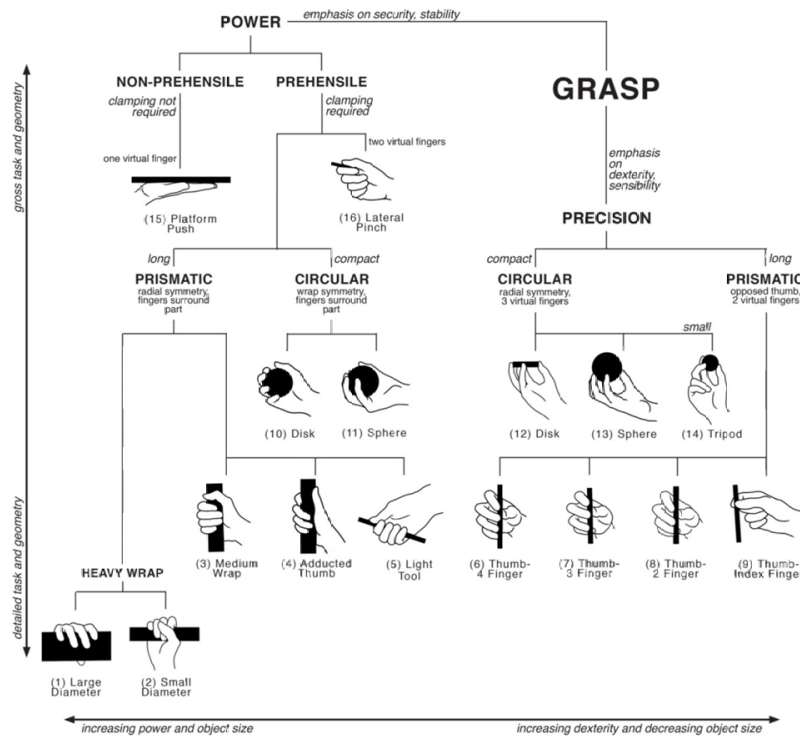


Figure 2.8: Grasp type taxonomy by Cutkosky [62] (Source: extracted from Zheng et al. [389]; ©2011 IEEE).

body regions where the stimuli were detected, creating a sensory ‘map’ that allows us to localize and differentiate between different types of touch. Once processed, the brain integrates this information with prior knowledge and context, allowing us to perceive the texture, shape, or temperature of objects.

During hand interactions with physical objects, tactile material properties play an essential role. According to Okamoto et al. [263], they consist of five dimensions: hardness (hard–soft), warmth (warm–cold), macro roughness (uneven–flat), fine roughness (rough–smooth), and friction (moist–dry, sticky–slippery). To design effective VR systems in Chapter 3, we need to understand how humans interact with (unknown) physical objects. Klatzky et al. [184] state that humans intuitively use *exploratory procedures* to gain accurate information about them. They provide examples of direct exploratory procedures that include:

- **Lateral motion** of the fingers across an object’s surface to explore its texture.
- **Applying pressure** to an object with a finger to explore its hardness.
- **Following the contour** of an object with the fingers to explore its shape.

Five dimensions of tactile material properties

Exploratory procedures to gather information about object properties

- **Enclosing** an object with the hand to explore its global size and volume.
- **Unsupported holding** an object with the hand to explore its weight.
- **Static contact** between an object and the hand to explore its temperature.

2.2.3.2 Kinesthetic Perception

Where is my body?

Kinesthetic is the perception of our own body movements, which is made possible through the sense of proprioception. It is essential for tasks like walking and lifting objects, as it provides constant feedback about the body's position and movements, enabling motor control and coordination even without the need for visual monitoring [274]. For example, without proprioception, we would need to look at our feet while walking to avoid stumbling. The proprioceptive sense uses specialized sensory receptors located in muscles, tendons, and joints, known as proprioceptors [249]. They continuously monitor the position and movement of body parts, detecting changes in muscle length, joint angle, and the force being applied. These include muscle spindles, which detect changes in muscle length and tension, and Golgi tendon organs, which sense the force and tension in tendons during muscle contractions [274].

*Processing of
proprioceptive
sensory inputs*

*Dynamic touch to
assess objects'
kinesthetic
properties*

The mechanical changes are converted into electrical signals, then transmitted via peripheral nerves to the spinal cord and brain. Here, they are processed in the cerebellum and the primary somatosensory cortex. The cerebellum plays a crucial role in coordinating movement, balance, and motor learning, while the somatosensory cortex helps to map the body's position in space (see later Figure 2.31). The brain integrates this kinesthetic information with other sensory inputs, such as visual and vestibular, ultimately allowing us to perceive body position, movement, and effort [249]. This also translates to the perception of object properties, helping us to, e.g., perceive their weight or inertia when holding and wielding objects with our hands. Such manipulations are often referred to as *dynamic touch* [349]. Through this, we can assess not only the properties of the object itself but also the properties of objects it comes into contact with. Thus, dynamic touch occurs in our everyday manipulatory activity of objects. It occurs statically and in motion, because the manipulated object affects the muscles and tendons of the hand and arm [349]. In contrast to tactile perception, dynamic touch focuses on muscles and tendons rather than skin deformations.

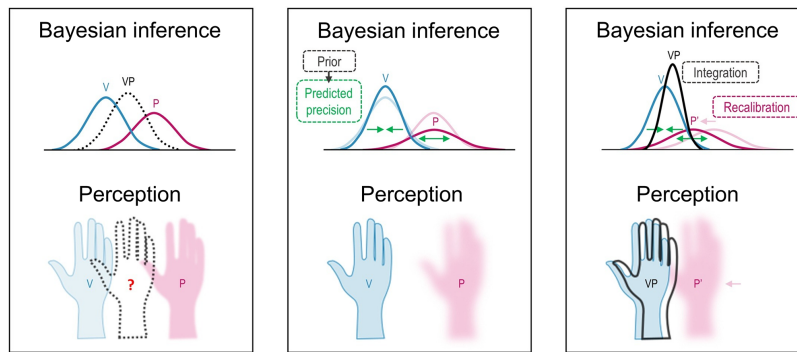


Figure 2.9: Bayesian inference during visuo-proprioceptive conflicts. Left: Shows Visual (V) and Proprioceptive (P) position of the hand. The sensory inputs are depicted as Gaussian probability distributions with equal variance (i.e., sensory noise). But where do we perceive our hand (VP) in case of conflicting sensory inputs? A naive multisensory estimate cannot explain the perceived hand position because prior experience and top-down processing affect the estimated precision of the sensory modality (Middle). Here, V hand position estimates have lower variance, i.e., less uncertainty than P. Right: V and P are integrated to form a combined multisensory estimate VP based on their weighted relative precision. This may also lead to recalibrating P'. Note, if the positional hand offset between V and P becomes too large, vision will typically be discarded in favor of proprioception, i.e., haptic dominance occurs [23]. (Source: Extracted and modified from Limanowski [219] with author's permission and under CC BY 4.0.)

2.2.4 Multisensory Perception

We experience our environment through multisensory perception, where sensations from sensory modalities are integrated to form a multisensory experience [81]. However, what happens if sensory modalities provide conflicting information? Ernst and Banks [80] suggest that sensory integration between vision and haptics can be described by maximum likelihood estimation (MLE). The proposed principle aims to minimize the variance or uncertainty in the overall perception. Sensations can be seen as probability distributions because of their variance, which is naturally related to the amount of noise in the sensory inputs. As an example, a probability distribution for visual and haptic sensory modality depicted in Figure 2.9, which follows a Gaussian probability density function [219]. The variance of each sensory probability distribution represents reciprocally the estimated precision of the sensory modality. MLE incorporates the degree to which each sensory modality contributes to the combined perception by a weighted sum that “adds the sensor estimates weighted by their normalized reciprocal variances” Ernst and Banks [80]. In neuroscience, the combination of sensory cues weighted by their relative precision is described as a process of Bayesian inference and may explain the phenomenon of sensory dominance and recalibration [219].

How does the brain integrate sensory inputs?

Sensory noise determines reliability

*Vision usually
dominates over
proprioception*

*There is a limit to
which the sensory
mismatch remains
plausible*

*Our beliefs and
expectations impact
our tolerance to
mismatches*

For example, visual estimates usually have lower variance than haptic (proprioception) estimates in settings that require estimating the spatial position of body parts [47]. Thus, in this case, vision likely dominates multisensory perception in conflicting sensory information. This is often referred to as the visual-dominance phenomenon, which has been observed in many studies that visually displace, e.g., users real from their virtual hand position [23, 35, 123]. The visual-dominance phenomenon is the core principle that this dissertation exploits, utilizing current VR technology's powerful capabilities in stimulating the visual sense to overcome the lack of haptics. However, vision does not always dominate over proprioception, as found in the study by Beers et al. [23]. The authors showed that the weighting of vision and proprioception varies with the direction of hand displacement. In the depth direction, participants relied more on proprioception than on vision. These findings demonstrate the potential realm of variables that contribute to the weighting of sensory modalities during multisensory integration.

As illustrated in Figure 2.9, the weights added to the sensory modalities depend not only on the sensory inputs' variance but are influenced by top-down processing [128, 219]. Top-down processing can compensate for the sensory variance of task-relevant stimuli while suppressing task-irrelevant stimuli. Gonzalez-Franco and Lanier [128] outline that it can even bias afferent sensory inputs in order to comply with a predicted (efferent) state, which may help to reduce uncertainty about limb positions [219]. Ultimately, the sensory correction of deficiencies is a powerful mechanism during perceptual inference, suggesting that VR systems do not need to be perfect and that our brains can compensate for some discrepancies or lacking resolution. Moreover, controlled visuo-haptic conflicts can elicit a design space for novel experiences in VR as further illustrated in Section 2.5.1.

2.3 THE GRAND CHALLENGE OF HAPTICS IN VIRTUAL REALITY

*Haptics is crucial to
create immersive
experiences*

Providing haptic feedback, i.e., creating physical forces between Lisa's body and the virtual objects, is one of the grand challenges in VR [165]. Many studies have shown the benefits of haptically supporting interactions in IVEs. For example, higher system immersion can lead to increased presence [158], enable intuitive and direct interaction with virtual objects [151], and improve task performance [166]. In the following section, we overview existing taxonomies in the field of haptics in VR to establish common ground for the remainder of this dissertation. Next, we review the two basic concepts of haptic feedback systems, namely passive and active, which help us classify systems and identify research gaps. Given the constant rise of VR




Figure 2.10: *Visual-Haptic Reality-Virtuality Continuum* by and adapted from Jeon and Choi [163].

technology and the increasing breadth of literature in the space, we use this section to provide an overview of the most influential haptic systems and concepts for hand-based interactions before focusing on the most relevant works in light of the contributions in this dissertation. A comprehensive overview of the literature on haptics in VR can be found in the doctoral dissertation by Zenner [380] and the survey on haptic solutions for VR interaction “*Can I Touch This?*” by Bouzbib et al. [36].

Given the rapid development in the field of haptic technology, there exist several ways how to classify haptic interfaces designed to support hand-based interactions in VR. Historically, research distinguishes approaches based on how the haptic stimuli are generated, resulting in two main classes: PHF and AHF devices. To reflect this, Jeon and Choi [163] proposed the *Visual-Haptic Reality-Virtuality Continuum* as an extension of the previously introduced *Reality-Virtuality Continuum* by Milgram and Kishino [242]. It adds a haptic dimension (x-axis) to the visual dimension (y-axis), resulting in a 2D continuum that describes the degree of reality and virtuality across the two dimensions. As shown in Figure 2.10, the continuum encompasses nine classes, with the real world being at the bottom left corner (visual and haptic reality) and the entirely virtual world + computer-generated haptic stimuli at the top right corner (visual and haptic virtuality). Our haptic devices developed in Chapter 3 fall into the visual virtuality haptic mixed reality category, according to the continuum.

Taxonomy that helps classify haptic systems

Given the diverse landscape of haptic technology, researchers also use the target sensory modality, tactile or kinesthetic, to classify the type of haptic device. In the following, we explain the two types of haptic feedback and provide examples of systems that haptically support hand-based interactions in VR. This provides a basis for our work presented in Chapter 3 and Chapter 4.

2.3.1 Active Haptic Feedback

*Active haptics uses
computer-controlled
actuation*

AHF uses computer-generated actuation, computing the appropriate haptic stimulus for a given interaction. This requires mathematical modeling of objects, their properties, and users' interactions to produce convincing haptic feedback through haptic rendering algorithms. Here, the advantages are that the forces can be dynamically updated to account for a variety of virtual object interactions [165]. AHF systems render kinesthetic forces targeting muscles, tendons, and joints, whereas tactile systems primarily focus on the human skin. Next, we provide an overview of active haptic technology from the perspective of groundedness: world-grounded, body-grounded, and ungrounded [165].

2.3.1.1 World-Grounded Devices

*Systems produce
forces relative to the
ground*

These haptic devices are physically attached to tables, walls, or the floor in a user's real environment. A prominent example of an AHF device is the *PHANToM* force feedback device [234] depicted in Figure 2.11 left. The system uses a robotic arm with a pen as an end-effector. Users interact with virtual objects through the pen. The system provides both 6-DoF position sensing and actuation of the end-effector to accurately render haptic feedback upon collisions with virtual objects. For example, *PHANToM* has been used in medical surgery training to simulate haptic feedback for drilling in bones or examining tissue [82].

*Producing
human-scale forces
often require large
actuators*

Rendering kinesthetic haptic feedback for human-scale forces is challenging and requires large actuators. An example is *HUG* (part of the *VR-OOS* system [286]) in Figure 2.11 right. It can sense human movements and simultaneously provide force feedback for arm-scale forces, e.g., in a virtual assembly simulation, where the systems can produce counterforce for lifting and moving objects. However, it becomes apparent that the physical limitations of actuators quickly lead to bulky and expensive systems. They also restrict users' freedom; for example, *HUG* requires users to remain in one place because the robotic arms are attached to a station; otherwise, they would be too heavy to wear as an exoskeleton.




Figure 2.11: Examples of world-grounded active haptics devices. Left: Image of PHANToM force feedback device extract from Massie, Salisbury, et al. [234]. Right: Image of a user using the HUG system. (Source: extracted and cropped from Sagardia et al. [286]; ©2015 IEEE).

2.3.1.2 Body-Grounded Devices

In contrast, body-grounded are directly attached to the user's body. This includes haptic gloves (or exoskeletons) that can render kinesthetic and tactile sensations. Hand-worn devices typically actuate or restrict users' movements of their limbs to create compelling kinesthetic sensations when grasping virtual objects. For example, *CyberGrasp*⁴⁸ pneumatically actuates users' fingertips, preventing them from penetrating virtual objects inside IVES. *Wolverine* [55] deploys a breaking mechanism to haptically support pinch-like grasping in VR. *Wireality* [88] shown in Figure 2.12 right, allows for the hand and individual finger joints to be programmatically blocked through the use of retractable strings. This way, the system can simulate kinesthetic forces when touching objects with different geometries. *Gravity* [54] extends rendering contact forces and weight by combining a breaking mechanism with vibrotactile feedback (see Figure 2.12: middle). While these systems focus on grasping rigid virtual objects, *CapstanCrunch* [305], a palm-grounded device uses capstan-based breaking to generate resistive forces that enable rendering compliance sensations when grasping virtual objects using a 2-finger pinch (see Figure 2.12: left). Strasnick et al. [327] introduced *Haptic Links*, an electro-mechanically actuated linking mechanism capable of rendering variable stiffness for bi-manual interactions with objects.

Systems that can be worn on the body

Handheld devices are effective, but constantly occupy the hands

2.3.1.3 Ungrounded Devices

Ungrounded or handheld devices mark the largest category of devices, including commercially used controllers such as the ones from HTC Vive or Meta Quest. Users hold these devices in their hand(s), which are spatially tracked, allowing them to perform unrestricted 6-DoF movements (see examples in Figure 2.13 and Figure 2.14). They typically use vibrotactile feedback through actuators such as eccen-

⁴⁸ CyberGrasp webpage: <https://tinyurl.com/5b49svsc>. Last accessed: Nov 1, 2024




Figure 2.12: Examples of body-grounded active haptics devices. Left: *CapstanCrunch*, a palm-grounded device that renders compliance when grasping virtual objects. (Source: extracted from Sinclair et al. [305]; ©2019). Middle: *Gravity* renders contact forces. (Source: extracted from Choi et al. [54]; ©2017). Right: *Wireality* simulates kinesthetic forces when touching objects. (Source: extracted and cropped from Fang et al. [88]; ©2020)

Variety of haptic
sensations can be
simulated

tric rotating mass (ERM) vibration motors, linear resonant actuators (LRAs), piezoelectric motors, or voice coil actuators. While current off-the-shelf systems mostly remain limited to basic vibrations to facilitate interactions with virtual objects, vibrotactile feedback has been successfully used in many research prototypes to convey a variety of tactile and kinesthetic haptic impressions. For example, it can simulate button presses [267], object properties such as compliance [179, 214] or softness of virtual objects [56], textures during mid-air interactions [328] and dynamic masses when moving objects [335]. In *TeslaTouch*, Bau et al. [21] showed that vibrotactile feedback at different frequencies and amplitudes can convey sensations of stickiness, smoothness, and friction during active finger exploration. Yet, designing vibrotactile feedback is challenging and can be done iteratively through an experienced haptic designer, directly recorded from physical interaction with a real surface (an approach called *Haptography* [198]) or through vocalization as shown by Degraen et al. [66] in *Weirding Haptics*.

Ungrounded active
kinesthetic haptic
devices

Handheld devices can create kinesthetic forces by using propulsion through directing air or by means of inertia (see Figure 2.13). Well-known examples include *Thor's Hammer* [147], which uses propellers to generate 3-DoF force feedback for mid-air interactions or *JetController* [354] which relies on air jets, reducing the form factor and improving the haptic rendering of the device. While such systems can provide compelling kinesthetic forces, they (by design) demand high power and produce wind and noise, which can impact the VR experience. On the other hand, ungrounded devices such as *MetamorphX* [145] change the inertia and viscosity of motion leveraging four control moment gyroscopes, i.e., flywheels, to provide 3-DoF moment feedback to the user's hand.




Figure 2.13: Examples of ungrounded active kinesthetic haptic devices. Left: *Thor's Hammer* uses propellers to generate force feedback for mid-air interactions. (Source: extracted and cropped from Heo et al. [147]; ©2015). Middle: *JetController* uses air jets to provide force feedback. (Source: extracted and cropped from Wang et al. [354]; ©2021). Right: *MetamorphX* simulates kinesthetic forces through inertia and viscosity changes. (Source: extracted and cropped from Hashimoto et al. [145]; ©2022).

Tactile feedback for hand interactions with virtual objects can be achieved by electrotactile feedback, skin deformation, and skin stretch. Devices are typically in direct contact with the skin to render tactile features. For example, *Tacttoo* by Withana et al. [364] is a feel-through temporary tattoo on the fingertip with embedded electrodes (taxels) that provide electrotactile feedback when touching virtual objects. This works by applying electric stimulation to the nerve stem of mechanoreceptors, which the brain interprets as mechanical vibrations. Groeger et al. [139] showed that taxels can be directly integrated into the fabrication process of proxy objects, e.g., by 3D printing. Shen et al. [301] presented *Fluid Reality*, a pneumatic shape-changing fingerpad array that can render high-resolution tactile feedback such as object geometries, textures, and compliance (see Figure 2.14: middle). Schorr and Okamura [295] proposed skin deformation on the fingertip to convey continuous feedback during virtual object exploration and manipulation, allowing users to perceive virtual objects of different mass, friction, and stiffness. Object shapes can be rendered with the help of extrudable and tiltable platforms such as the *NormalTouch* [28] haptic controller or actuated pin-based arrays like in *PoCoPo* [376] and the *TextureTouch* [28] controllers (see Figure 2.14: left). In *Gravity Grabber*, Minamizawa et al. [243] show that deformation on fingerpads created by a belt mechanism can create the illusion of weight, even in the absence of proprioceptive sensations (see Figure 2.14: right). Knierim et al. [191] proposed *TactileDrone*, quadcopters enclosed in a cage flying into the user to provide tactile feedback to various locations on the body. In one of the side projects during this dissertation, we developed *HapticPuppet* [100], expanding on the idea of using computer-controlled drones—puppeteering the users by actuating their limbs through body-attached strings to provide kinesthetic forces.

*Approaches for
simulating tactile
sensations*

*Creating a
kinesthetic illusion
through tactile
feedback*




Figure 2.14: Examples of ungrounded active tactile haptic devices. Left: *TextureTouch* uses pin-based arrays under the fingertip to render tactile feedback. (Source: extracted from Benko et al. [28]; ©2016). Middle: *Fluid Reality* a pneumatic shape-changing fingerpad array to render tactile feedback. (Source: extracted and cropped from Shen et al. [301]; ©2023). Right: *Gravity Grabber* applies deformation on fingerpads to create the illusion of weight. (Source: extracted and cropped from Minamizawa et al. [243]; ©20217).

2.3.1.4 Emerging Concepts for Haptic Feedback

In recent years, we have witnessed emerging technologies in the field of haptics in VR, such as *UltraHaptics* [49]. It uses computer-controlled ultrasonic transducers that emit ultrasound to a focal point (often the user's fingertip or palm), producing tactile haptic sensations upon making contact with the skin. This approach allows for contactless mid-air haptic feedback beyond the scope of VR. *AIREAL* [316] uses air vortices directed onto the human skin to provide tactile feedback mid-air that can be felt up to 1 m. Both are depicted in Figure 2.15.

Triggering
involuntary
movements

Kinesthetic forces can further be rendered by using systems that can electrically stimulate users' muscles or tendons through technologies such as EMS and Tendon Electrical Stimulation (TES). These approaches use electrodes that are attached to the users' skin and, by using electric impulses, elicit involuntary muscle or tendon contraction. The main difference between the two approaches lies in the placement of the electrodes so that EMS stimulates motor nerves and TES sensory nerves. Lopes et al. [224, 225] showed that EMS can




Figure 2.15: Invisible haptics. Left: *AIREAL* emits vortices that users can feel mid-air. (Source: extracted and cropped from Sodhi et al. [316]; ©2013 ACM). Right: *UltraHaptics* emits ultrasound to the user's palm, producing tactile haptics. (Source: extracted from Long et al. [223]; ©2014 ACM).




Figure 2.16: Examples of emerging haptic concepts. Left: *Haptic Source-Effector* stimulates the sensorimotor cortex to create haptic feedback. (Source: extracted and cropped from Tanaka et al. [336]; ©2024). Middle: *Mutual Turk* provides a shared physical prop through which pairs can exchange kinesthetic forces. (Source: extracted and cropped from Cheng et al. [51]; ©2017). Right: Electrical Muscle Stimulation (EMS) can enhance haptics for walls and heavy objects. (Source: extracted and cropped from Lopes et al. [225]; ©20217).

be used to simulate kinesthetic forces of virtual impacts, touching virtual walls, and carrying heavy virtual objects (see Figure 2.16: right). TES has been combined with other haptic feedback modalities, such as vibration and visual techniques, to simulate forces for virtual button presses [15]. However, TES and EMS induce tingling sensations and require skin attachment. To overcome these limitations, Tanaka et al. [337] investigated Magnetic Muscle Stimulation (MMS) using contactless electromagnetic coils that can actuate muscle up to 5 cm away, even through clothing. They further noted that participants reported significantly less discomfort and more realistic feedback of kinesthetic forces compared to traditional EMS.

Tanaka et al. [336] recently proposed a novel approach to provide haptic feedback by non-invasively stimulating the corresponding brain areas using Transcranial Magnetic Stimulation (TMS) (device see Figure 2.16: left). They presented the *Haptic Source-Effector*, a head-wearable device with a single magnetic coil that can be moved across the user's scalp to stimulate the sensorimotor cortex locally. Their studies showed that they can create about 15 different tactile and kinesthetic sensations on various parts of the body (e.g., hands, arms, legs, feet, and jaw). While participants reported different touch sensations, from tapping, vibrating, and tingling to pressing, the approach may change how we think about designing haptic experiences in the future.

Simulating haptic effect at the source

Finally, an interesting concept is to use pairs of users to provide haptic feedback to one another. Cheng et al. [51] call this mutual human actuation, and they proposed a system, *Mutual Turk* that provides a shared physical prop through which users can exchange kinesthetic forces (e.g., a fishing rod and kite). The two users experience different VR scenes, drilling a fish and letting a kite fly, but are un-

Human actuate other humans




Figure 2.17: Examples of passive haptic feedback. Left: Using real-world objects to facilitate manipulation of 3D neurosurgical visualizations. (Source: extracted from Hinckley et al. [151]; ©1994 ACM). Right: A passive haptic kitchen that provides touch feedback for a virtual kitchen. (Source: extracted from Insko [158]; ©2001).

aware of the fact that there is a human in the same room (see Figure 2.16: middle). By synchronizing the two experiences so that their way of manipulating the shared props is consistent across both virtual worlds, they add haptic feedback to one another's experience. While this does not allow for full control over the haptic stimulus, it is an exciting concept for VR designers that can playfully incorporate haptics into a multiuser VR experience.

2.3.2 Passive Haptic Feedback

PHF uses physical objects, often referred to as proxies, that act as 'stand-ins' for virtual objects without any computer-controlled or actuating components. The central idea is that proxies resemble physical properties of, ideally, multiple virtual objects such as their shape, size, weight, or material, giving physical embodied to what otherwise would remain purely virtual. Originating from the work by Hinckley et al. [151], who proposed real-world objects to facilitate manipulation of 3D neurosurgical visualizations, physical proxies made their way to IVEs to support interactions with 3DUIs [220] and to haptically enhance large-scale virtual environments [158]. In these settings, proxies were quickly prototyped with materials such as plywood, plexiglass, foam-, and cardboard and, as a result, only approximated virtual objects. However, researchers found that, even though a close match is desirable, proxies generally enhance the immersive experience and allow for more intuitive and direct interactions [158, 248]. This led to a growing interest in using proxies to haptically enhance interactions in IVEs. Hence, researchers proposed everyday objects as proxies [63, 150], fabricated low [248] and high fidelity proxies [326], designed proxy-based controller [302, 382] or developed proxy construction kits (i.e., toolkits) [16, 91] for VR. Nilsson et al. [256] formulated two central challenges for VR proxy design, *Colocation* and *Similarity*. *Colocation* refers to the proxy being in the same location as

*To the roots of
passive haptics*

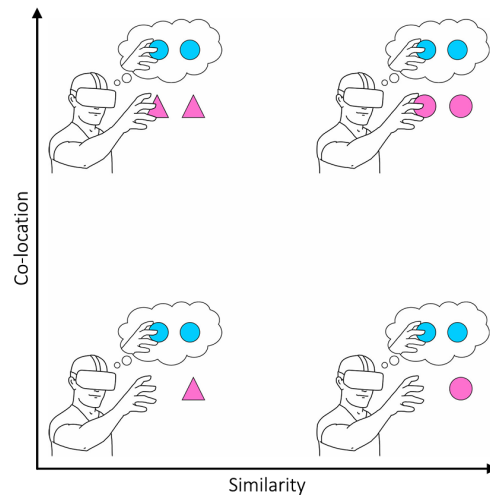


Figure 2.18: Challenges of *Colocation* and *Similarity*. (Source: extracted from Nilsson et al. [256]; ©2021 IEEE).

the virtual object, whereas *Similarity* refers to how closely the physical properties of the proxy match the virtual object's properties.

2.3.2.1 *Similarity of Proxies*

Users interact with proxies by touching, grasping, or holding the physical object; thus, using proxies undeniably leads to a multisensory experience. In contrast, AHF systems often aim to isolate the haptic feedback modality. As a result, proxies usually provide tactile and kinesthetic haptic feedback as a whole, conveying physical object properties such as weight and inertia, material and texture, as well as shape and size. While some discrepancy is acceptable, researchers looked into the effects of purposely introducing mismatches between proxy and virtual object to investigate potential effects. For example, Simeone et al. [304] varied the degree of mismatch between the virtual object's and proxy's material, temperature, size, shape, and weight. Based on their studies, Simeone et al. [304] recommend to “minimize mismatches with the manipulable parts of an object” and “to maintain correct proprioceptive feedback even if it conflicts with the real pose”. However, they also reported that participants found interacting with a similar proxy as engaging as interacting with an actual proxy replica. There exist several other studies that investigate how tactile [248, 342] and kinesthetic mismatches [362] affect the VR experience and performance, reporting similar results. Generally, it is desirable to approximate virtual object properties relevant to task and interaction. However, passive proxies have an inherent limitation because they cannot change their physical properties—making them inflexible, and physically reconstructing the virtual environment defeats the purpose of VR. To overcome this limitation, passive proxies can be equipped with actuators that can change their properties on de-

Does this feel like the object I am touching?

How close does it need to be?

mand, a concept Zenner and Krüger [382] coined Dynamic Passive Haptic Feedback (DPHF), a subcategory of mixed haptic feedback.

*Proxies that can
change their own
haptic properties
dynamically*

DYNAMIC PASSIVE HAPTIC FEEDBACK In DPHF, proxies do not host actuators to exert forces onto humans but to change their own physical properties. As a result, they only require small actuators because they do not need to produce human-scale forces.

*I can feel the weight
of the object*

Weight- and Inertia-Changing proxies, for example, *Shifty* [382] a rod-based controller with an internal weight that can be re-positioned, change their center of mass and, thus, the inertia of the controller upon interactions with virtual objects. Zenner and Krüger [382] showed that *Shifty* can enhance the perception of virtual objects varying in length, thickness, and weight. This is an effective way to increase the haptic rendering capabilities of passive proxies, allowing them to become truly multipurpose. Hence, it has informed the design of many VR proxies targeting kinesthetic and tactile haptic feedback. *Transcalibur* [302] applies weight-shifting in 2D to render the weight distribution of 2D objects, while *SWISH* [287] uses 3D weight-shift to simulate fluids. Rendering weight sensations of virtual objects, e.g., by transferring liquid mass into a controller, is a technical challenge. *MobileGravity* [173] is a handheld proxy with a liquid reservoir station that can perform weight changes up to 1 Kg within a few seconds. *ElastOscillation* [346] and *ElaStick* [285] are DPHF proxies that can create the kinesthetic forces of moving flexible objects such as casting with a finishing rod. *Drag:on* can simulate resistance by using a shape-changing mechanism to adjust its surface area [383].

*Continuous
exploration of
surfaces*

Texture- and Material-Changing proxies, for example, the *HairTouch* [213] handheld controller that uses two brushes close to the index fingertip, and by controlling the hairs' length and bending direction, it provides different stiffness and roughness sensations. *Haptic Palette* [68] is a handheld proxy with a rotatable disc that is augmented with different material samples. The system presents the material that is most suitable for the interaction to the user. Whitmire et al. [363] designed the *Haptic Revolver*, a handheld controller with an actuated wheel that raises and lowers underneath the finger to render continuous contact with a virtual surface. The exchangeable wheel can be equipped with different textures to provide a large range of sensations to the user. This controller combines features of DPHF and AHF because it can also render contact and shear forces to the index finger during virtual object exploration.

Shape- and Size-Changing proxies can change their physical form factors through actuation. For example, *PuPoP* [339] is a pneumatically inflatable airbag that user wear on their palm, approximating

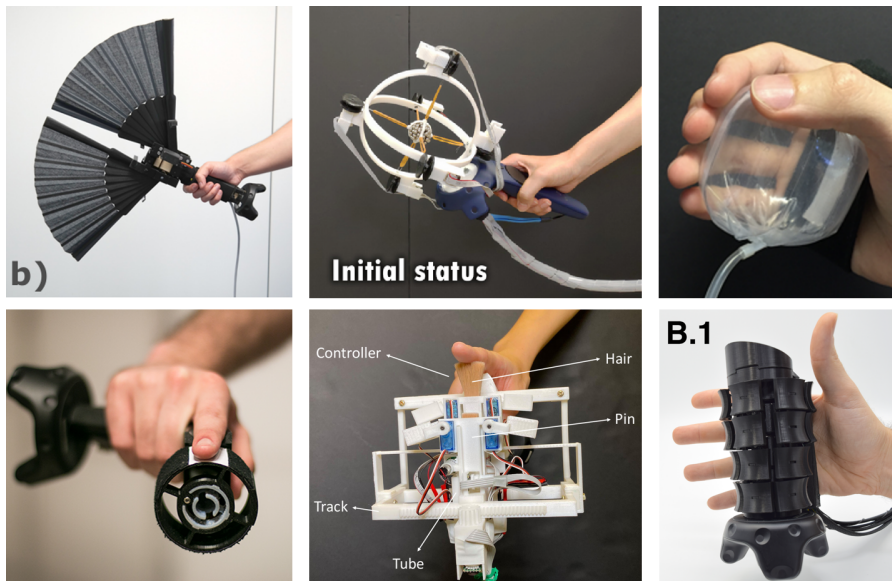


Figure 2.19: Examples of dynamic passive haptic feedback devices. Top left: *Drag:on* can simulate mid-air resistance by changing its surface area. (Source: extracted and cropped from Zenner and Krüger [383]; ©2019). Top middle: *ElastOscillation* provides 3D multilevel damped oscillation force. (Source: extracted from Tsai et al. [346]; ©2020). Top right: *PuPoP* is a pneumatically inflatable airbag that approximates the shape and size of virtual objects. (Source: extracted and cropped from Teng et al. [339]; ©2018). Bottom left: *Haptic Revolver* uses an actuated wheel that raises and lowers underneath the finger to render continuous contact with a virtual surface. (Source: extracted from Whitmire et al. [363]; ©2018). Bottom middle: *HairTouch*, a handheld controller that uses hairs to create stiffness and roughness sensations. (Source: extracted and cropped from Lee et al. [213]; ©2021). Bottom right: *X-Rings* uses expandable rings to render surfaces of virtual objects. (Source: extracted and cropped from Gonzalez et al. [135]; ©2021).

the shape and size of virtual objects. *PoCoPo* by Yoshida et al. [376] it a handheld pin-based shape display that can render various 2.5D shapes in the user's hand, and Gonzalez et al. [135] developed *X-Rings* a shape-changing proxy capable of rendering the surface of virtual objects through expandable rings.

Shape-changing
devices

In the following section, we overview the field of Encounter-Type Haptic Feedback (ETHF) that primarily focuses on the challenge of *Colocation* [256]. Arguably, there exist overlaps in the concepts of DPHF and ETHF, and some of the proxies and systems presented in the previous and the next sections fall into either category.

2.3.2.2 Colocation of Proxies

When exploring virtual environments, it is often undesired that our hands are occupied by a controller. Humans explore their environment through active touch, and while having a handheld controller-

Can I touch this
object over here?

type device is beneficial in many situations, it takes away users' freedom for active, unencumbered object exploration [211]. However, in the absence of a spatially tracked handheld controller, the question arises of how we can synchronize the position of proxies and their virtual counterparts in 3D space.

ENCOUNTER-TYPE HAPTIC FEEDBACK ETHF, sometimes referred to as *robotic graphics* [240], is a part of mixed haptic feedback in the continuum. It focuses on providing haptic feedback wherever and whenever a user interacts with objects inside the IVE. Systems often deploy robotic platforms or structures, presenting the appropriate physical proxy to the user upon interaction (see examples in Figure 2.20). ETHF systems come in various scales and, similar to AHF, can be categorized by their groundedness. For example, *Snake Charmer* by Araujo et al. [12] is a world-grounded system that uses a robotic limb with a replaceable proxy end-effector that can host a variety of shapes, textures, and materials. Through accurate hand tracking and prediction models, the system anticipates the targeted virtual object in the environment and quickly positions the appropriate proxy to haptically support the interaction. *Haptic-Go-Round* [155] uses a rotating platform that surrounds the user hosting several proxies that are suitable for the given VR experience. The system presents the proxy to the user that is most suitable for the interaction. In *HapticBots*, Suzuki et al. [333] use small-scale robots that can change their height and orientation through tiltable and height-adjustable platforms to haptically render continuous virtual surfaces and objects. *ShapeShift* [308] utilizes a spatially tracked shape display to render the shape and size of virtual objects. By using an omnidirectional robotic platform, the device can reposition itself and provide both vertical and lateral kinesthetic feedback. The above-mentioned systems require heavy actuation and engineering efforts that limit their practical use in everyday VR experiences. To address this problem, Wang et al. [355] proposed *MoveVR*, an encounter-type system using an off-the-shelf cleaning robot. The system can create kinesthetic effects of tension, resistance, and impact by changing the robot's moving speed, rotation, and position.

*Dynamically
re-position the proxy*

*Continuous
rendering of shape
and size*

*Overcoming
technical barriers*

*Body-grounded
encounter-type
haptic devices*

Examples of body-grounded encounter-type haptic devices include the previously mentioned *Haptic Revolver* and *X-Rings*. Another example is *PIVOT* developed by Kovacs et al. [196], a wrist-worn haptic device that renders grasping, catching, or throwing a virtual object in an IVE. It uses a proxy that can be moved in and out of the user's hand by a single actuated joint.

To this day, ungrounded encounter-type haptic devices rely on flying quadcopters, i.e., drones. First proposed by Yamaguchi et al.




Figure 2.20: Examples of encounter-type haptic feedback devices. Top left: *HoverHaptics* uses drones for mid-air positioning of proxies. (Source: extracted and cropped from Abtahi et al. [4]; ©2019). Top middle: *Snake Charmer* positions a proxy in space using a robotic end-effector. (Source: extracted and cropped from Araujo et al. [12]; ©2016). Top right: *HapticBots* is a swarm of robots that can provide continuous touch feedback. (Source: extracted and cropped from Suzuki et al. [333]; ©2021). Bottom left: *ShapeShift* renders virtual objects continuously through a spatially tracked shape display on wheels. (Source: extracted and cropped from Siu et al. [308]; ©2018). Bottom middle: *PIVOT*, a wrist-worn device that renders grasping, catching, or throwing a virtual object. (Source: extracted and cropped from Kovacs et al. [196]; ©2020). Bottom right: *Haptic-Go-Round* uses a rotating platform that surrounds the user hosting several proxies. (Source: extracted and cropped from Huang et al. [155]; ©2020).

[371], the authors showed that drones could be used to provide haptic feedback for interacting with surfaces. However, rendering force feedback through drones is most effective on the y-axis (1D) due to their mechanical design [1]. Abtahi et al. [4] presented *HoverHaptics* proposing drones for dynamic re-positioning of passive haptics and mid-air positioning of proxies. However, the noise and wind turbulence caused by drones remain a limiting factor to this day.

*Ungrounded
encounter-type
haptic devices*

Finally, Cheng et al. [51, 53] presented several systems that apply the principles of using humans for actuation rather than hardware. One example is the *TurkDeck* system, which instructs people who are not in VR to manually relocate proxies (wall and other artifacts) through floor projections and audio. They ‘build’ the physical world around the VR user to ensure that there is always something to touch. *ZoomWalls* [375] and *RoomShift* [332] are actuation-based systems that use robotic platforms to dynamically re-position room-scale proxies on demand. All three approaches are depicted in Figure 2.21.




Figure 2.21: Examples of encounter-type haptics at room-scale. Left: *ZoomWalls* that dynamically repositions walls at room-scale. (Source: extracted and cropped from Yixian et al. [375]; ©2020). Middle: *TurkDeck* uses people to reposition room-scale proxies. (Source: extracted and cropped from Cheng et al. [53]; ©2015 ACM). Right: *RoomShift* uses swarm robots that can reposition proxy furniture. (Source: extracted and cropped from Suzuki et al. [332]; ©2020).

2.3.3 Conclusion

Rendering haptic feedback for interactions in VR is crucial to enhance presence [158], enabling practical training and simulations that transfer [164, 257] with comparable task performances [222] from the virtual to the real world.

Opportunities and Drawbacks of AHF

AHF relies on mathematical simulations to compute haptic stimuli for a given user interaction inside the IVE. This has clear advantages as computer-controlled actuation is very flexible, allowing full control of the haptic feedback, and can be adjusted dynamically to cover a wide spectrum of interactions. The large body of presented systems and studies suggest that AHF is a very effective way to deliver compelling tactile and kinesthetic haptic sensations. However, AHF approaches are often costly and have a high mechanical and mathematical complexity. The devices consist of actuators and electronics that consume power and, therefore, require batteries or come with the downside of being tethered. Together with the mechanical components, they can quickly become bulky and restrictive. Moreover, larger actuators for human-scale forces pose safety risks as failures might harm users.

Opportunities and Drawbacks of PHF

In contrast PHF applies an inexpensive approach to haptics in VR that repurposing everyday objects [63, 150, 304], fabricated low [91, 248] and high fidelity proxies [69, 326]. As a result, they are inexpensive and typically have low complexity. PHF proxies do not host actuators and, thus, do not require additional hardware and mathematical simulations to compute haptic stimuli. While AHF approaches often focus on one type of haptic feedback, handheld proxies automatically deliver multimodal haptic feedback. This makes finding suitable proxy objects more challenging as designers drive to minimize the, for the interaction relevant, mismatches [150, 304]. However, “physically replicating the detail of every object that the

user might interact with would be costly and time-consuming; indeed it would nullify all the advantages that a VE could offer” Insko [158]. Yet, PHF proxies remain inflexible because they cannot adapt their physical properties to the constantly changing demands of IVEs.

The concept of DPHF combines the advantages of AHF and PHF by deploying low-cost actuators, allowing proxies to change their haptic properties such as center of mass dynamically [382], mid-air resistance [383], shape [302] and size [135]. This makes the approach versatile, especially for addressing the challenge of proxy *Similarity* [256], providing a solid base for our investigations surrounding (RQ1).

2.4 DYNAMIC PROXIES FOR VIRTUAL REALITY

This section discusses the most closely related literature relevant to our work in Chapter 3. Here, we focus on the challenge of proxy similarity and present two novel approaches for designing and developing proxies that can change their haptic properties. Parts of the text in this section have previously appeared in:

[92] **Martin Feick**, Cihan Biyikli, Kiran Gani, Anton Wittig, Anthony Tang, and Antonio Krüger. *VoxelHap: A Toolkit for Constructing Proxies Providing Tactile and Kinesthetic Haptic Feedback in Virtual Reality*.

*In Proceedings of
UIST 2023*

[93] **Martin Feick**, Donald Degraen, Fabian Hupperich and Antonio Krüger. *MetaReality: Enhancing Tactile Experiences using Actuated 3D-printed Metamaterials in Virtual Reality*.

*In Frontiers in
Virtual Reality 2023*

2.4.1 Construction of Proxies

In this part, we closely look at proxies for handheld (*dynamic touch* [349]) and *exploratory* [184] interactions. Considering VR experiences that target learning and simulation, like the one that Lisa finds herself in, require proxies that closely resemble their virtual counterpart. One way to address this is to use replicas for the most important virtual objects in the environment. However, this can quickly become costly and inflexible and make VR experiences only accessible to a few users who can afford them. To address this limitation, a set of proxies that can change their physical configuration in a computer-controlled (shape-changing) way [284, 388] or through manual configuration [91, 216, 391] can help.

2.4.1.1 Reconfiguration of Proxies

*Build it as you need
it*

*Modularity
increases flexibility*

McClelland et al. [239] presented *HaptoBend*, an example of a manual reconfigurable device, allowing users to create 2D plane-like shapes and multi-surface 3D shapes by bending them. *StrutModeling* [216] utilizes physical struts and hub primitives that enable users without a 3D modeling background to prototype 3D models. Struts can be adjusted in length and snap to magnetic hubs in any configuration and virtually represented through their capturing software. Zhu et al. [391] used Rubik's twists, a passive low-cost twistable artifact, allowing users to build interactive haptic proxies for various hand-held virtual objects. The shape approximations can be equipped with active components, e.g., buttons, to increase the range of supported interactions, making them more interactive. Li et al. [218] presented *HapLinkage*, allowing realistic simulations of, e.g., wrenches, pliers, scissors, and syringes by using a linkage mechanism that supports typical motion patterns. However, it may be challenging to construct, considering that users reportedly struggle with shape creation with such complex device mechanisms [391]. One way to address this was presented by Roudaut et al. [284] with *Cubimorph*, a modular interactive device that uses a hinge-mounted turntable mechanism to self-reconfigure in the user's hand to adapt to the requirements. Zhao et al. [388] used block-based swarm robots that self-assemble physical handheld-sized proxies for VR but are limited to block-like shapes. However, in contrast to *HapTwist* and *HapLinkage*, block-based approaches have been shown to allow lay users to easily construct rough approximations [91, 183].

*Block-like
construction has
been effectively used*

ActiveCube [183] is a device that allows users to construct and interact with 3D environments by using cubes with special functionality. They can be attached to each other, recognizing their 3D structure to ensure high shape similarity between the virtual and physical objects. The authors present a range of applications, including physical data visualization, spatial cognitive ability assessment, and creative storytelling and modeling of the environment. *RoBlocks* [298], is a computational construction kit consisting of sensor, logic, and actuator blocks. Users can arrange and combine individual blocks to explore and experiment with complex ideas in science, technology, engineering, and mathematics. Arora et al. [16] presented *VirtualBricks*, a modular block-based construction kit based on the LEGO® platform that enable passive haptics in VR. Besides regular bricks, they offer a set of feature bricks that allow mechanical rotation and translation of parts of the proxy, which can be tracked through embedded electronics. This allows for high visuo-haptic concurrency even during proxy exploration and manipulation. During his master studies, the author of this dissertation contributed to this area by developing *TanGi* [91], a toolkit for rapid construction of passive haptic proxies




Figure 2.22: Top left: *VirtualBricks* is a toolkit for constructing proxies. (Source: extracted and cropped from Arora et al. [16]; ©2019). Top middle: *StrutModeling* is a construction kit to model 3D objects. (Source: extracted from Leen et al. [216]; ©2017). Top right: *Cubimorph*, for designing modular interactive devices. (Source: extracted and cropped from Roudaut et al. [284]; ©2016 IEEE). Bottom left: Robotic assembly of haptic proxy objects. (Source: extracted and cropped from Zhao et al. [388]; ©2016). Bottom middle: *Hapto-Bend*, a passive shape-changing proxy. (Source: extracted and cropped from McClelland et al. [239]; ©2017). Bottom right: *HapLinkage* enables prototyping haptic proxies for hand tools. (Source: extracted and cropped from Li et al. [218]; ©2020).

using 3D-printed *Composable Shape Primitives* and *Manipulators* for rotation, translation, bending and stretching that can be reconfigured to approximate form and function of visual objects. In contrast to *VirtualBricks*, users can 3D print the shape primitives to allow for higher customization and high-fidelity shape and size approximation, demonstrated through two user studies. Together, Kitamura et al. [183], Arora et al. [16] and Feick et al. [91] highlight the benefits of block-based construction kits with respect to their great modularity, scalability, and ease of use, allowing to construct a large range of proxies for passive haptics. Block-based proxies can be enhanced by using components from active haptics, especially vibrotactile actuation. *Kooboh* [179] is a rigid block that can render various mechanical properties such as compliance and deformation when being squeezed or pressed by a user. This is made possible without any visual feedback, only by sensing input forces (e.g., from a pinch hand gesture) and rendering a vibrotactile stimulus according to their haptic algorithm. *VibeRo* [6] combines the force-driven synthesis of vibration with an additional visual deformation corresponding to the input force. Their results suggest that they can simulate various levels of stiffness with a single physical proxy. This is possible due to perceptual processes discussed in Section 2.2. We will present an overview of several other techniques that deploy such principles in Section 2.5.

Giving functionality to proxies

Adding AHF components

2.4.2 Fabrication of Proxies

*Digital fabrication
for VR proxies*

*3D printing allows
users to fabricate
objects with varying
physical properties*

The field of proxies in VR greatly benefits from the development of additive and subtractive manufacturing technology, such as 3D printing and laser cutting. Nowadays, several companies offer affordable systems for research labs but also for personal fabrication in the consumer market. As recent advancements in fabrication technologies support the manufacturing of highly detailed physical artifacts, they can be used to construct artifacts with varying haptic properties. As a result, researchers started to experiment with different fabrication methods to create more scalable and flexible proxy objects. In order to simulate tactile experiences, previous work has combined visual textures with influencing different tactile dimensions of texture perception [263]. For example, Degraen et al. [69] used 3D-printed hairs to enhance tactile one-finger explorations of surface materials in VR. The authors were able to influence the feeling of roughness and hardness by changing the length of the printed hairs, and by overlaying visual textures, they effectively increased the resolution of the proxy. Hair-like designs were further studied by Takahashi and Kim [334] to create haptic displays using perforated plates to alter the length and behavior of the available set of hairs. Other examples include methods for designing objects with desired mechanical behavior, such as elasticity [297] or deformation [34] through varying internal microstructures, or for fabricating perceptually-varying surface texture qualities [67, 117, 271]. While these approaches are effective, they cannot change their physical properties without requiring refabrication.

*Metamaterials, an
underexplored type
of material for proxy
design*

In the context of this dissertation, we are interested in the possibilities of metamaterials as a novel method for designing dynamic proxies. Metamaterials can be defined as: “...a novel class of complex composite materials [with the] ability to exhibit any desirable electromagnetic, acoustic, or mechanical property such as negative mass, stiffness, or Poisson’s ratio...” according to Valipour et al. [352]. There exist many categories of metamaterials; however, we are most interested in the ones that allow us to change the surface structure. Generally, metamaterials are composed of unique cells on a regular grid. Varying cell designs are used to achieve the desired behavior [114, 352]. An exciting class of metamaterials is the so-called auxetic materials. In contrast to regular materials, they can enlarge their surface area (negative Poisson ratio) when being stretched. Steed et al. [320] proposed a deformable auxetic material that is actuated by a small number of mechanical pistons (i.e., linear motors). As a result, it can be bent in multiple directions while feeling smooth and rigid. Thus, it can vary the curvature of the surface, resulting in different terrains.



Figure 2.23: Top left: 3D-printed metamaterial textures. (Source: extracted and cropped from Ion et al. [159]; ©2018). Top middle: 3D-printed microstructures to control elasticity. (Source: extracted and cropped from Schumacher et al. [297]; ©2015). Top right: A mechatronic shape display based on auxetic materials. (Source: extracted and cropped from Steed et al. [320]; under CC BY 4.0). Bottom left: Strings to compress metamaterials, creating shape-change. (Source: extracted and cropped from Neville et al. [252]; under CC BY 4.0). Bottom middle: 3D-printed hairs with visual texture overlays to create combined visuo-haptic perception. (Source: extracted and cropped from Degraen et al. [69]; ©2019). Bottom right: 3D printing objects with varying deformation properties. (Source: extracted and cropped from Piovarči et al. [271]; ©2016).

Ion et al. [160] explored the possibilities of 3D-printed metamaterials to embed mechanical mechasims in 3D-printed objects. In metamaterial textures, Ion et al. [159] presented 3D-printed metamaterial surface geometries that can perform a controlled transition between two or more textures. To design and manufacture metamaterial textures using FDM 3D printing, they provide a publicly available editor. The shape and structure of the metamaterial are inspired by *Origami* and surface wrinkling techniques. Each metamaterial can be compressed, forcing certain cells to deform. For example, a bicycle handle with an adjustable grip, i.e., more or less friction caused by a compressed and uncompressed metamaterial, respectively. Yang et al. [373] presented an approach to enable reconfiguration of these metamaterial designs without the need for reprinting. In *MetaSense*, Gong et al. [125] integrate sensing capabilities into 3D printable metamaterial structures. The authors use conductive filament to embed capacitive touch-sensing sensors into the design, allowing them to measure capacitance variation based on user input. This enables sensing of various interactions such as acceleration, binary state, shear, and magnitude and direction of applied force. To enable compression of metamaterial designs, Neville et al. [252] present actuated

Metamaterials can be 3D-printed with flexible filaments such as TPU

Conductive filaments add sensing capabilities to 3D-printed metamaterials

metamaterials inspired by *Kirigami*, i.e., the art of cutting and folding paper to get 3D shapes. They ran multiple strings through their metamaterial designs to achieve deformations that go beyond single-axis compression. In fact, this opens up an interesting design space for metamaterial actuation that we build upon in Section 3.2.

2.4.2.1 *Rapid Prototyping*

*Fabricating a proxy
is slow...is it?*

Researchers presented several ways how to fabricate proxies for VR, reaching from prototyping with plywood, foam-, and cardboard [158] to fully 3D-printed replicas [91]. However, they usually come with a trade-off between fabrication time and proxy resolution. Here, rapid prototyping techniques can help to quickly create physical artifacts. For example, Mueller et al. [247] presented *faBrickation*, a prototyping technique to overcome the generally slow fabrication process of FDM 3D printing. The authors proposed to build the core of the prototype with low-resolution LEGO® bricks to provide a robust basis while simultaneously printing the high-resolution parts of the physical artifact. They report faster fabrication times than traditional 3D printing by a factor of $\times 2.44$. Platener [33] applies a similar concept, with laser-cut sheets being used to create the basic structure of the object, while 3D-printed parts are used for complex shapes, achieving a speed-up of up to $\times 10$.

2.4.3 *Conclusion*

*Manual
reconfiguration of
multimodal haptic
proxies*

Proxies come in all shapes and sizes, but seeking or fabricating a suitable proxy that resembles as many physical properties of a virtual object as possible is challenging [256]. For simulation and training, close replicas of virtual objects are crucial for the most important virtual objects in the IVE, allowing Lisa to experience plausible scenarios. Ultimately, this can lead to training results comparable to the real world, allowing Lisa to experience situations that are impossible (due to safety and ethical reasons) to recreate in the real world [164, 222]. One way to create replicas is through manual reconfiguration [91, 216, 391]. This is an inexpensive approach to enable proxies to closely resemble virtual objects in form, feel, and function while overcoming the limitation of needing one proxy for each virtual object [158]. Especially, block-based construction has been demonstrated as an effective and intuitive way for users to construct proxy objects [16, 91, 183, 284, 298]. In Section 3.1.1, we combine modular building blocks with rapid prototyping techniques to overcome the limited haptic resolution of block-based proxies. Rather than focusing on one type of haptic feedback, we follow a holistic approach by combining the advantages of PHF and AHF.

How users interact with the proxy limits the type of haptic feedback it needs to support. For instance, if Lisa only locally touches the object's surface to assess its material, the proxy does not need to match the virtual object's weight, center of mass, or size. During such explorations of textures and materials through Lisa's fingers, abstract fabricated structures have been combined with visually overlaid textures to simulate material experiences [69]. However, to support the visuo-haptic perception of different virtual materials, still a large enough set of proxies needs to be fabricated. Our work in Section 3.2 aims to support the design and fabrication of novel proxy objects that can dynamically change their tactile properties upon actuation, enabling a single physical artifact to represent multiple virtual objects, according to the concept of DPHF. We take inspiration from the most recent work in the fabrication space on 3D-printed metamaterials, studying their potential to enhance tactile feedback in VR.

Dynamic metamaterials for changing tactile properties

2.5 HAND-BASED ILLUSIONS FOR VIRTUAL REALITY

Illusion techniques for VR exploit users' perceptual system similar to the optical illusions shown in Section 2.2.2. In VR, we have full control over the visual stimuli, allowing us to manipulate what Lisa sees from what she feels in the IVE, 'tricking' her perception into, e.g., believing she is in a virtual space larger than is physically available [279]. This can be achieved through *redirected walking* during natural locomotion. The technique continuously applies a slight (unnoticeable) offset to Lisa's viewpoint. She compensates for this offset, resulting in her walking in circles even though she thinks that she walks in a straight line [322]. Other techniques utilize change blindness to dynamically change room layouts, leading to *impossible spaces* that physically do or even cannot exist [330]. In the context of this dissertation, we are interested in illusions that involve the visual and haptic (kinesthetic and tactile) sensory modalities. We focus on hand-based illusions, primarily with proxies, to enhance haptics in VR. Thus, we will use the remainder of this section to discuss the related literature in this space. Parts of the text in this section have previously appeared in:

Overcoming the physically 'impossible'

[98] **Martin Feick**, Kora Regitz, Lukas Gehrke, André Zenner, Anthony Tang, Tobias Jungbluth, Maurice Rekrut, and Antonio Krüger. Predicting the Limits: Tailoring Unnoticeable Hand Redirection Offsets in Virtual Reality to Individuals' Perceptual Boundaries.

In Proceedings of ACM UIST 2024

[103] **Martin Feick**, André Zenner, Simon Seibert, Anthony Tang, and Antonio Krüger. The Impact of Avatar Completeness on Embodiment and the Detectability of Hand Redirection in Virtual Reality.

In Proceedings of ACM CHI 2024

- In Proceedings of ACM UIST 2023* [104] **Martin Feick**, André Zenner, Oscar Ariza, Anthony Tang, Cihan Biyikli, and Antonio Krüger. Turn-It-Up: Rendering Resistance for Knobs in Virtual Reality through Undetectable Pseudo-Haptics.
- In Proceedings of IEEE VR 2023* [97] **Martin Feick**, Kora Regitz, Anthony Tang, Tobias Jungbluth, Maurice Rekrut, and Antonio Krüger. Investigating Noticeable Hand Redirection in Virtual Reality using Physiological and Interaction Data.
- In Proceedings of ACM CHI 2022* [99] **Martin Feick**, Kora Regitz, Anthony Tang, and Antonio Krüger. Designing Visuo-Haptic Illusions with Proxies in Virtual Reality: Exploration of Grasp, Movement Trajectory and Object Mass.
- In Proceedings of ACM CHI 2021* [95] **Martin Feick**, Niko Kleer, André Zenner, Anthony Tang, and Antonio Krüger. Visuo-haptic Illusions for Linear Translation and Stretching using Physical Proxies in Virtual Reality.

2.5.1 Designing Illusions with Proxies

In the previous Section 2.1.2, we already introduced the famous *Rubber Hand Illusion* [35]. Such hand-based illusion techniques have been of central interest to the research community because they promise an inexpensive way to improve VR experiences. For example, changing the mapping between real and virtual hand movements, which, e.g., can be achieved by changing the Control/Display (C/D) ratio, introducing a gain factor g , that virtually amplifies ($g > 1.0$) users' real-world movements. The *Go-Go* interaction technique [272] uses this approach to dynamically scale hand movements, allowing users to interact with virtual objects that are out of reach. Although Lisa could clearly notice the applied offset, she maintains high control (i.e., agency) over her movements [113]. Such beyond-real interaction techniques [3] are effective, improve ergonomics, and thus facilitate interactions inside IVEs.

Hand Redirection is the foundation

HAND REDIRECTION Before discussing the different techniques to improve proxy-based interactions in VR, we must introduce the underlying method that builds the foundation for most of these techniques. The terms HR or *Redirected Reaching* are used interchangeably in the literature. For the remainder of this dissertation, we will use the term HR, which describes the phenomenon of users adjusting their movement trajectory to compensate for a visually induced offset during aimed hand movements (see Figure 2.24). To this day, three different HR approaches have been proposed:

- *Body warping* applies manipulations to the virtual hand by off-setting its position from the position of the tracked real hand.

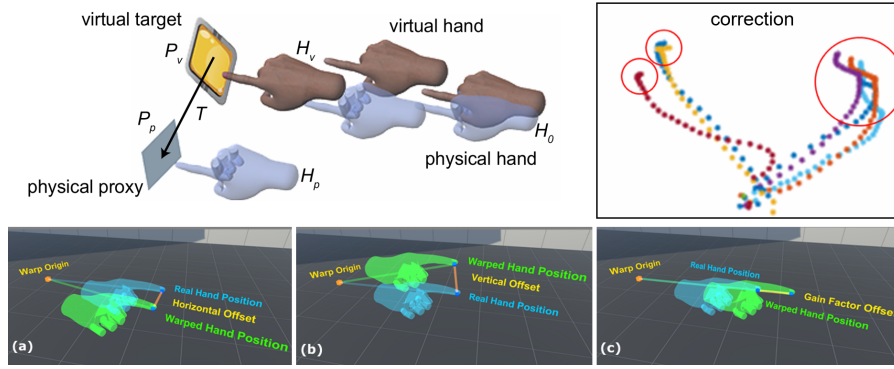


Figure 2.24: Top left: Gradual horizontal HR showing the position of the real and virtual hand during reaching. (Source: extracted and cropped from Cheng et al. [52]; ©2017). Top right: Shows correction phase in hand trajectory under the influence of HR. (Source: extracted and cropped from Azmandian et al. [17]; ©2016 ACM). Bottom: Displays the effects of horizontal (a), vertical (b), and gain-based (c) HR. (Source: extracted and modified from Zenner and Krüger [384]; ©2019 IEEE).

- *World warping* does not manipulate the virtual hand, but the IVE, by, e.g., rotating or translating the virtual scene.
- *Hybrid warping* combines both body and world warping.

While these three approaches exist, most of the related literature and the techniques presented in the dissertation utilize HR based on body warping. Different redirection algorithms have been proposed, reaching from fixed positional offset [25, 47], mapping functions [113, 272, 361], space partitioning [245], functional optimization [387] to linear interpolations [17, 52, 236]. The latter is the most common method, gradually increasing or decreasing the absolute offset between the position of the virtual and the tracked real hand. For the method to work, the system typically requires a *warp origin*, a *virtual target*, and *physical proxy* in 3D space. If the coordinate systems are aligned in the warped origin, the virtual and real hands share the same location across both realities; thus, they have a one-to-one mapping. Once Lisa starts moving her hand toward the virtual target, the system gradually applies the calculated offset. As a consequence, Lisa adjusts her real hand trajectory to hit the virtual target. The difference in hand movement trajectory under the influence of redirection vs. no redirection can be seen in Figure 2.24 top right. The warping origin can be set to anywhere on the hand [236], depending on which part is expected to make contact with the target, often the fingertip of the index finger. Similarly, the virtual and the physical target can be set to anywhere in the 3D space, and the redirection algorithm computes the required offset along a vector to ensure visual-haptic matching upon object contact.

*Existing algorithms
for Hand Redirection*

*Users compensate
for the offset during
hand reaching*

Predicting users' interactions

Heuristics and machine learning can help with this

Determining a suitable physical proxy for the targeted virtual object can be achieved through systems such as *Annexing Reality* [150] or *Substitutional Reality* [304]. However, knowing the target object the user is going to interact with next is a challenging problem. Previously, this has been realized through users selecting the next object [238], fixed order of interactions [17], through predictions and heuristics based on eye gaze and hand trajectory [17, 52, 237] or hand trajectory predictions based on ML [57]. Another challenge is aligning the virtual and real hand after redirection without sudden virtual hand jumps. Typically, users need to return their hand to the *warp origin* [384] before they can continue reaching for the next target. Cheng et al. [52] presented an *on-the-fly* HR technique, allowing for seamless transitions between redirection offsets. In the following, we overview the field of hand-based illusion with proxies, categorized into addressing the challenge of *Similarity* or *Colocation* [256]. Please note that since the space has grown incrementally and rapidly over the past decade, there exist several ways how to classify these approaches. We aim to be as inclusive as possible but only discuss approaches that relate to our overarching research questions.

2.5.1.1 Similarity—Changing the Perceived Properties of Proxies

Hand-based illusion with proxies can simulate tactile and kinesthetic effects through two central concepts, *Pseudo-Haptics* [210] and *Redirected Touching* [193].

Haptics without 'real' haptics

PSEUDO-HAPTICS Proposed by Lécuyer et al. [210], *Pseudo-Haptics* creates the perceptual illusion of haptics through visual manipulations of users' interactions. It purposely decouples the visual from the haptic sensory modality, e.g., by applying dynamic C/D ratio manipulations to a 2D mouse cursor to simulate virtual texture such as bumps and holes [209] or changing the perceived elasticity of images [14] (see Figure 2.25: right). In Lécuyer et al. [210] original work on pseudo-haptics, they simulate virtual springs varying in their stiffness through a single passive input device. The technique works based on the visual-dominance phenomenon, where the manipulated visual information dominates over the proprioceptive sensory feedback during sensory integration. Ever since pseudo-haptics has been used to simulate a variety of haptic properties [227] with Xavier et al. [367] providing the most up-to-date review article. Most relevant in the scope of this dissertation are pseudo-haptic techniques for enhancing interactions with proxies.

For example, Samad et al. [288] developed *Pseudo-Haptic Weight*, a proxy-based illusion technique that can vary the perceived weight of a virtual object when lifting it (see Figure 2.25: left). This can be achieved by applying $C/D \ g > 1.0$ to simulate lighter objects, result-




Figure 2.25: Examples of Pseudo-Haptic Effects. Left: *Pseudo-Haptic Weight* changes the perceived weight of virtual objects by manipulating the C/D ratio. (Source: extracted and cropped from Samad et al. [288]; ©2019). Top right: Changing the perceiving local elasticity of a sponge. (Source: extracted and cropped from Argelaguet et al. [14]; ©2013 ACM). Bottom right: Changing the perceiving softness of a surface. (Source: extracted and cropped from Punpongsanon et al. [275]; ©2015 IEEE).

ing in the virtual object moving faster than the tracked real object. Vice versa, $C/D \text{ } g < 1.0$ can simulate heavier objects by slowing down users' real-world movements with the proxy. They also presented a model based on multisensory integration theory that explains the relationship between C/D gain and perceived virtual object mass. Rietzler et al. [282] showed that the pseudo-haptic technique also works with a traditional VR controller and does not require a proxy that has high shape and size similarity. Yu and Bowman [377] extend these findings by proposing three pseudo-haptic techniques that can simulate mass and mass distribution during proxy-based object rotations, and Dominjon et al. [75] for simulating object masses during object translations.

Manipulating the perceived weight of virtual objects

Simulating stiffness or compliance of virtual objects embodied by a physical proxy has been of great interest to the research community. In *FlexiFingers*, Achibet et al. [5] combine a passive multi-finger haptic device with a pseudo-haptic technique to simulate different levels of stiffness when grasping and pinching virtual objects. Weiss et al. [360]'s pseudo-haptic stiffness illusion, allows them to change the perceived stiffness of virtual objects by manipulating the C/D ratio of the virtual hand. Participants used their hand to press down on a passive haptic device with constant stiffness while being exposed to C/D manipulations of different magnitudes. Their results suggest that participants perceived the virtual objects to be up to 28.1% softer ($g > 1.0$) and 8.9% stiffer ($g < 1.0$). Changing the perceived stiffness of virtual objects can also be achieved by visual deformation of the fingertip upon making contact with an object, as recently shown by Morimoto et al. [246]. Their results suggest that the deformation of the fingertip shape was small, the object was perceived as hard, and vice versa.

Stiffness and compliance illusions

Texture and
deformation
illusions

Punpongsanon et al. [275] influence the perception of softness by augmenting surfaces using different projection-based pseudo-haptic effects (see Figure 2.25: bottom right), while Sato et al. [290] proposed different visual hand augmentations to modify the user's perception of surface textures in terms of unevenness, slipperiness, and softness. Other systems combine elements of pseudo-haptic and AHF. For example, Adilkhonov et al. [6] proposed *VibeRo*, a system that combines pseudo-haptic stiffness and voice-coil actuator on the contact surface to simulate squeeze forces at the fingertips. Heo et al. [148] developed *PseudoBend*, consisting of a rigid proxy rod equipped with a vibrotactile voice-coil actuator and 6-DoF force sensor. They render grain vibrations and pseudo-haptic deformation of the rod according to the applied forces by a user. As a result, they can effectively create compelling haptic illusions of twisting, rotating, and bending.

A single proxy can
act as a stand-in for
multiple virtual
objects

REDIRECTED TOUCHING Kohli [193] introduced *Redirected Touching*, a technique that applies continuous HR during tactile exploration. The central idea is to have a single proxy which shape can be mapped to multiple virtual objects. To achieve this, the tracked real hand follows the shape of the object, e.g., contours and edges, while the virtual hand is being offset to follow the virtual shape of the object—creating a visual-haptic illusion. Kohli [192] describes the idea as a visual distortion of a virtual object's shape or size relative to a real object, i.e., the proxy. In their early explorations, the author demonstrates how a single physical flat table can provide haptic feedback for a curved, tapered, and sloped virtual table. Further, he developed a system that maps different virtual geometries onto dynamically captured physical geometry. An example can be seen in Figure 2.26 right, where a user touches the corner of the physical foam board, which is then mapped to different virtual geometries. Kohli et al. [194] presents a method for space warping based on interpolation and studies how this affects task performance. Their results suggest that warped virtual objects show similar throughputs and error rates to unwarped virtual objects. Zhao and Follmer [387] extended the space warping algorithm through functional optimization to support continuous redirection on complex geometries.

Continuous
redirection on
physical proxy

Ban et al. changed the perceived position of edges on a virtual object [19] and surface shapes of virtual objects [18] by visually redirecting the user's finger on a passive proxy. By combining the insights from these two experiments, Ban et al. [20] created a *Visuo-Haptic* system with a proxy that can dynamically change the position of a haptic edge. As a result, they could create the illusion of continuous exploration of several different virtual geometries, although they always touched the same physical proxy (see Figure 2.26: left). Bergström et al. [30] proposed *Resized Grasping*, a technique allowing a single




Figure 2.26: Examples of redirected touching. Left: Simulating various types of surface geometries by decoupling the virtual from the real-world position of the hand. (Source: extracted and cropped from Ban et al. [20]; ©2014 ACM). Middle: Improving the perceived performance of shape displays by vertically offsetting the virtual finger position. (Source: extracted and cropped from Abtahi and Follmer [2]; ©2018). Right: The user's virtual finger position is warped to touch three different virtual objects embodied by a single physical foam board proxy. (Source: extracted and cropped from Kohli [193]; ©2010 IEEE).

physical proxy to represent multiple virtual objects of different sizes by redirecting the user's fingers. The authors apply C/D gains during pinch-type grasping gestures, creating a visuo-haptic illusion of larger and smaller physical objects. Abtahi and Follmer [2] applied angle, vertical, and horizontal redirection to a user's index finger to increase the perceived resolution of pin-based shape displays for proxy rendering in VR. Their results demonstrate that redirection increased the perceived resolution of the shape display, allowing it to expand its haptic rendering capabilities. Redirection has also been applied to enhance haptic interaction with tools. Strandholt et al. [326] introduced *Redirected Tool-Mediated Manipulation* to improve realism for interactions with a physical proxy tool (e.g., a hammer, a screwdriver, and a saw) making contact with another physical prop. For example, a virtual hammer's position can be offset to ensure that physical impacts accompany each strike of a virtual nail. *VRGrabbers* [372] is a chopstick-like passive VR controller designed for precise object selection and manipulation. By its design, the controller can only haptically support grasping virtual objects of the same size; however, by redirecting the chopsticks, i.e., they visually move faster or slower than physically moved, the device can create the visuo-haptic illusion of virtual objects differing in size.

Redirected touching can overcome missing resolution of passive and active haptic devices

2.5.1.2 Colocation—Changing the Perceived Location of Proxies

HR-based techniques can also be used to address the general problem when proxy and virtual object are not in the same location. This is of great importance, especially for touch and even grasp interactions with objects in IVEs. While the concept of ETHF already provides one solution to this problem by dynamically re-positioning the proxy to align with the virtual object, it typically relies on heavy actuation. In

Does the proxy align with the virtual object?



Figure 2.27: Examples of haptic retargeting. Left: Virtual sliders, buttons and knobs embodied by a proxy. (Source: extracted and cropped from accompanying paper video by Matthews et al. [237]; ©2023). Middle: Single proxy cube provides touch feedback for multiple virtual cubes. (Source: extracted and cropped from Azmandian et al. [17]; ©2016 ACM). Right: *Sparse Haptic Proxy* in combination with HR renders touch feedback. (Source: extracted and cropped from Cheng et al. [52]; ©2017).

this section, we give an overview of the software-based techniques that can be used individually but also alongside ETHF devices.

HAPTIC RETARGETING Here, the most well-known technique is *Haptic Retargeting* introduced by Azmandian et al. [17]. It allows the system to redirect the users' real hands to the same physical proxy while touching different virtual objects. As an example, the authors demonstrate that a single physical proxy cube can represent three virtual cubes placed at different locations in the IVE (see Figure 2.27: middle). In their work, they compare the three warping techniques (i.e., body, world, and hybrid) and conclude that hybrid warping provides the highest flexibility. Further, they note that the amount of world warping should be kept to a minimum because it requires heavy manipulation of the IVE. Cheng et al. [52] showed the potential of body warping HR and a *Sparse Haptic Proxy*, a passive haptic device consisting of a set of geometric primitives (see Figure 2.27: right). They redirect the user's hand to a matching primitive based on the predicted virtual target. Both HR techniques require three positions, the *physical proxy* (pp), the *warp origin* (wo) and the *virtual target* (vt). The algorithm computes the offset vector between pp and vt and incrementally applies it in relation to the traveled distance from the start of the hand movement, i.e., the wo . In Azmandian et al. [17]'s body warping algorithm requires users to start and end their hand movement in the wo to undo the applied warp and allow users to reach for the next virtual object. This limits the kinds of interactions that users can perform because during natural interaction, users seamlessly transition between virtual targets. To address this, Cheng et al. [52] presents an extension of the HR algorithm that allows *on-the-fly* retargeting between virtual objects without having to retract the hand. They modify the offset to interpolate between

Redirecting the real hand to the proxy while touching a virtual object in a different place

Core technique that we build upon in this dissertation




Figure 2.28: Left: Experiment for bimanual haptic retargeting. (Source: extracted and cropped from Gonzalez and Follmer [133]; ©2019). Middle left: In *Retargeted Self-Haptics*, users use their own body parts for haptics. (Source: extracted and cropped from Fang and Harrison [87]; ©2021). Middle right: *Interface warp* allows a single physical button (bottom) to represent multiple virtual buttons (top). (Source: extracted and cropped from Matthews et al. [238]; ©2019 IEEE). Right: *REACH+* improves the perceived performance of low-speed ETHF devices. (Source: extracted and cropped from Gonzalez et al. [132]; ©2020).

the current vt and the next vt . In parts of our work in Chapter 4, Chapter 5 and Chapter 6 we use Cheng et al. [52]’s redirection algorithm, as it is a very intuitive way to achieve redirection. Matthews et al. [236] presented *Shape Aware Haptic Retargeting* to support interactions between any part of the user’s hand and any part of the vt . Their algorithm considers hand and target geometry and computes the shortest distance to the point of contact.

While previous work only focused on single-hand interactions, Gonzalez and Follmer [133] demonstrated that bimanual haptic retargeting can be effectively used by testing several combinations of simultaneous left- and right-hand retargeting (see Figure 2.28: left). Matthews et al. [238] introduced *Interface Warp* and combined it with bimanual retargeting to support asymmetric bimanual interactions that can be used to map one physical button onto many virtual buttons, e.g., on a traditional VR controller (see Figure 2.28: middle right). They explore the effects of the novel retargeting technique, showing faster task response times and fewer errors for their interface warp technique. Matthews et al. [237] built on these findings, developing a proxy interface with physical knobs, switches, and sliders that can resemble virtual controls (see Figure 2.27: left). They apply haptic retargeting to enable dynamic mappings between virtual and physical controls and study how different synchronization methods, manual and automatic, affect the usability of the interface. Their work builds on the novel concept we propose in Chapter 4.

*Redirection to
enlarge the
interaction space of
existing interfaces*

Fang and Harrison [87] proposed *Retargeted Self-Haptics*, a proxy-free approach to haptics, where hand movements are retargeted to the user’s own body to deliver haptic touch feedback. The authors provide several examples, illustrating the design space of the tech-

*Timely topic with a
continuous stream of
innovation*

nique. For instance, a user places their left hand onto a palm scanner while entering a password via a number keypad with the right hand. The right hand is retargeted to the backside of the left hand to provide touch feedback—using the user’s own body (see Figure 2.28: middle left). Finally, haptic retargeting and ETHF both address the challenge of *Colocation*. As a result, researchers looked into ways to combine them to maximize the techniques’ strengths. In *REACH+*, Gonzalez et al. [132] propose hand retargeting to improve the perceived performance of low-speed ETHF devices during physical interaction in VR. They estimate the user’s contact time with their intended virtual target using Cheng et al. [52]’s heuristic and redirect the user’s hand to a point within the ETHF’s spatiotemporally reachable 2D space (see Figure 2.28: right).

2.5.2 Detectability of Illusions

*What I see is not
what I feel*

The presented techniques offset visual from proprioceptive sensory information, which can, in fact, remain unnoticeable to users. This can be explained by the inaccuracies in the SCs model [128] and the visual-dominance phenomenon [47, 80, 123] as outlined in Section 2.2.4. Here, the brain monitors the hand movement execution and corrects for sensorimotor discrepancies without the awareness of the user. Burns et al. [47] and many others showed that vision usually dominates over proprioception during hand movements during sensory integration. Nevertheless, if the offset between vision and proprioception becomes too large, it can result in a semantic violation [265], which in turn leads to a break in presence—disrupting the immersive VR experience [314]. Hence, researchers looked at how much offset the system may introduce without risking disruption, attempting to find a sweet spot between noticeability and effectiveness. One way to ensure that hand-based illusions remain unnoticeable is by determining DTs. In signal detection theory, an absolute DT has to be defined as the level at which a stimulus will be detected with a certain probability (often 50% or 75%) [181]; in our case, that means the offset between the visual and proprioceptive sensory inputs. There exist a variety of methods that allow researchers to determine DTs for interactions, which we describe in the next section.

*Breaks in presence
can occur*

The majority of the related work on hand-based illusions uses the real world, i.e., one-to-one mapping between the position of the real and the virtual hand, as the baseline stimulus. Participants are asked to report any noticeable difference from this stimulus. However, it is important to note that thresholds can be determined with different baseline stimuli, e.g., after visuomotor adaptation. For example, offsets such as those deployed by the *Go-Go* interaction technique [272] can, after adaptation and onset of embodiment,

be effectively used because interactions have in a way that aligns with users' expectations (in accordance with the SCs). While such techniques are very effective in improving ergonomics and usability [194], they require an initial adaption phase [172], and offsets cannot be dynamically adapted unless we know how much we can divert from the current mapping [90]. A DT can, therefore, be seen as the first noticeable divergence for an expected behavior—the current calibration of the user. As a result, they are relevant not only for determining the maximum difference to the real world but to any VR experiences utilizing sensory discrepancies. Our work in Chapter 4, Chapter 5 and Section 6.2 uses the real world as a reference because it is the easiest and most common to measure, while Section 6.1 looks into adaption to visuo-proprioceptive offsets.

Users can adapt to offsets and perform tasks efficiently

Burns et al. [47, 48] were the first to explore the role of visual-dominance over proprioception during HR in VR. They offset the virtual from the real hand position to avoid penetration of virtual objects. By investigating users' DTs for visuo-proprioceptive offsets, they found that users are much more sensitive to visual interpenetration than to visuo-proprioceptive sensory conflicts. Zenner and Krüger [384] report DT estimates for the amount of unnoticeable gradual HR during mid-air reaching tasks, allowing designers to apply HR-techniques without risking detection. The reported HR DTs were later confirmed by Lebrun et al. [207] with Gonzalez and Follmer [133] extending the results to bi-manual HR, and Hartfill et al. [144] to different movement directions. However, these thresholds were established in very conservative settings, yielding relatively low DTs that often lack the ability to create strong perceptual effects. Researchers recently looked into two ways to improve this: (1) by deploying strategies that utilize blinks [385] and saccades [381] (co-author publication), distraction factors [83, 384] or tendon stimulation [261] to cover the presence of HR, allowing for greater offsets, or (2) looked into factors that impact the potential unnoticeable offset, a string of research that we extensively contribute to in Chapter 5.

Detection thresholds to provide limits for hand redirection

Techniques to enlarge the detection thresholds

2.5.2.1 Potential Factor Influencing Noticeability

By far, the most apparent factor is participants' awareness of HR being applied and the task to report its presence as soon as they noticed it, leading to very conservative estimates, potentially way below the 'true' noticeability level. This is supported by Debarba et al. [64] results that show participants perform poorly in detecting hand offsets when they are unaware of their presence. Benda et al. [26] recently followed up on this by quantifying the difference between DTs determined with participants being aware vs. unaware of the presence of a gain-based HR technique. Their results suggest that unawareness of HR result in an about $\times 2.3$ increase in DTs than traditional conserva-

Many factors impact the detectability of HR

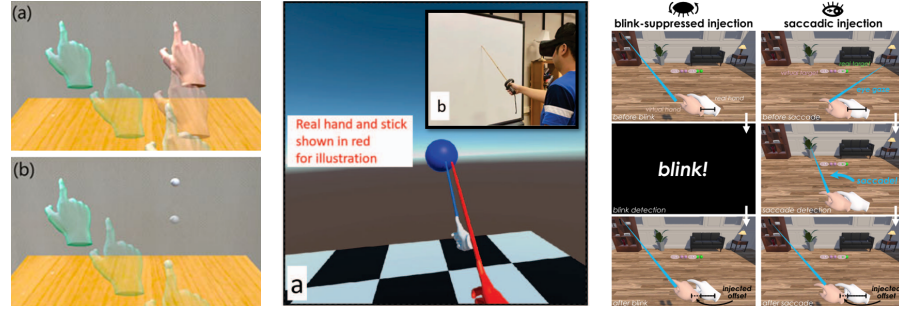


Figure 2.29: Left: Investigating the effects of realistic and abstract hand visualizations on the DTs. (Source: extracted and cropped from Ogawa et al. [262]; ©2021 IEEE). Middle: HR DTs threshold with passive haptic feedback. (Source: extracted and modified from Zhou and Popescu [390]; ©2022 IEEE). Right: Saccadic and blink-suppressed HR to increase DTs. (Source: extracted from Zenner et al. [381] (co-author publication); ©2024).

*DTs cannot be seen
as fixed values*

tive experiments. There also exists strong evidence for significantly different sensitivities to visuo-proprioceptive offsets, depending on movement direction [144]. Esmaeili et al. [83] investigated the influence of cognitively demanding tasks on the detectability of HR, reporting anecdotal evidence for larger DTs with increased mental task load. Experiments on redirected touch already suggested an influence of haptic feedback on the robustness of the illusion during interaction. Abtahi and Follmer [2] investigated the DTs for redirected index finger movements with a congruent haptic cue. Their results show much higher DTs in the presence of haptics, which may be a result of greater weighting of the sensory inputs during integration. Zhou and Popescu [390] recently studied DTs with a handheld stick with which the user can tap virtual objects, receiving PHF through a wall upon contact. Their presented technique enlarges the DTs, exploiting both concurrent haptic feedback and clever interaction design that only requires minimal hand movements to achieve larger offsets in tapping location (see Figure 2.29: middle). Ogawa et al. [262] showed that the hand visualization itself can affect the detectability of HR. They found that realistic avatars (i.e., human hands) can foster ownership of the redirected movements better than abstract avatars (i.e., spherical pointers), thus making the redirection less noticeable. Similar to the haptic stimuli, this can be attributed to the higher weighting of the visual modality during multisensory integration (see Figure 2.29: left). Very recently, Lebrun et al. [206] proposed that the optimal magnitude of hand-based illusions depends on users' proprioceptive and visual sensory sensitivities. The precision of the sensory modalities may serve as an input for computational modeling of the consequences of hand-based illusion at different magnitudes. In Chapter 5, we contribute to this ongoing research by exploring several factors related to the user, interaction, and proxy (RQ3).




Figure 2.30: Psychophysical methods. Left: Psychometric sigmoid function showing discrimination results for gain-based HR. (Source: extracted and cropped from Zenner and Krüger [384]; ©2019 IEEE). Right: The staircase procedure toolkit to determine DTs. (Source: extracted and cropped from GitHub repo by Zenner et al. [379]; ©2023).

2.5.2.2 Detect Noticeability through Psychophysics

The previously discussed literature typically applies psychophysical methods to determine DTs for various types of hand-based illusions with and without proxies. Our work also applies methodologies for estimating DTs in Chapter 4, Chapter 5, and Chapter 6. Since this is a central aspect of this dissertation, we outline the three key methods we use in the following section. A comprehensive overview can be found in the book “*Psychophysics: A Practical Introduction*” by Kingdom and Prins [181]. An absolute DT is defined as the point at which participants have the probability ψ_{target} to detect the stimulus when presented to them. Two probabilities have received special recognition in the scientific community, 75% and 50%. The latter is also known as the Point of Subjective Equality (PSE) [322] or the Conservative Detection Threshold (CDT), i.e., users have a 50% chance to detect the stimulus. This point is of interest because it is where users perceive a stimulus to be equal to the baseline. On the other hand, the DT marks the point where users can discriminate the stimulus from the baseline in 75% of the time. The target probability depends on the research goal, and our choice varies between projects. To determine DTs, we applied three scientifically established methods: *the method of constant stimuli*, *the adaptive staircase method*, and *the method of adjustments*. Next, we describe the theory behind these procedures. The exact configuration of these methods in our experiments can be found in the corresponding sections.

*Definition of
perceptual limits*

In the **method of constant stimuli**, participants are repeatedly exposed to a set of pre-selected stimuli of different intensities, i.e., in our case, illusions of different magnitudes. They range from undetectable to apparent stimuli [215] and are presented to participants in a randomized order. Each stimulus is repeated x times, with x being a parameter of the study design. Participants are instructed to detect the absence or presence of a stimulus after each trial and report this to the experimenter. As a result, we can compute the percentage of

*Methods to estimate
detection thresholds*

correct detection per stimulus intensity, i.e., the probability of detecting the illusion. Next, we fit a psychometric sigmoid function to the data, modeling the discrimination performance of the participants. The intersection between the psychometric function and the target probably marks the DT. Thus, the psychometric function describes the relationship between all stimuli and their probability of detection.

In contrast, the **adaptive Up(Δ^+)/Down(Δ^-) staircase method** does not model the discrimination performance across all stimuli but instead provides a streamlined way to estimate a DT at a target probability ψ_{target} . This can be achieved by selecting the next stimulus based on the previous responses of the participants. For example, an ascending staircase sequence starts with an unnoticeable stimulus and increases the stimulus intensity at a predefined step size Δ^+ as long as the participant does not detect it. Once the participant detects the stimulus, it decreases the stimulus by step size Δ^- . A directional change within a sequence is noted as a reversal point. Typically, the number of reversal points is used as a termination criterion and for computing the DT at a target probability ψ_{target} . Equation 1 can be used to compute the required step size Δ^+ and Δ^- for $\psi_{\text{target}} = .75$:

$$\psi_{\text{target}} = \frac{\Delta^+}{\Delta^+ + \Delta^-} \Rightarrow \frac{\Delta^-}{\Delta^+} = \frac{1 - \psi_{\text{target}}}{\psi_{\text{target}}} = \frac{(1 - .75)}{.75} = \frac{1}{3} \quad (1)$$

Note that if $\Delta^+ = 1$ and $\Delta^- = 1$, the staircase is called symmetric and targets $\psi_{\text{target}} = .50$ or the CDT. There exist three different types of staircases: ascending, descending, and interleaved. An ascending staircase starts with the minimum stimulus and increases the intensity until a participant detects it. Vice versa, a descending staircase starts with the maximum stimulus and decreases the intensity of the stimuli. Both approaches are prone to biases. Therefore, the interleaved staircase uses one descending and one ascending sequence and randomly assigns the next trial to one of the sequences. The procedure increases the next stimulus intensity if a participant fails to detect the current stimulus and decreases the next stimulus if the user detects the manipulation.

Types of staircases

Finally, we also applied the **method of adjustments** or sometimes called the average error method. Here, participants adjust, i.e., increase or decrease the stimulus intensity to match a comparison stimulus. For example, the participant is presented with a comparison color and is asked to adjust another color to match the comparison. The CDT is the point at which participants detect a difference from the comparison stimulus in 50% of the trials.

In this dissertation, we used three different task structures in combination with psychophysical experiments. A 1-Alternative-Forced

Choice ($1AFC$), a 2-Alternative-Forced Choice ($2AFC$) and a 2-Interval-Forced Choice ($2IFC$) method [181]. In a $1AFC$ or ('yes', 'no') task, participants perform a single trial at a given stimulus intensity and are asked to report if they noticed the stimulus or not. Typically, this task structure is used to determine an absolute DT. During $2AFC$ or $2IFC$, participants experience two stimuli and are asked to compare them, reporting if they differ or not. During a $2AFC$, the two stimuli are being presented simultaneously, while in the $2IFC$, participants experience them consecutively. By using these methods, we target the difference DT, which is the minimum amount of change in stimuli required for a human to notice it.

Task structures used in combination with psychophysical methods

2.5.2.3 A Novel Approach Towards Detecting Noticeability

In **RQ4**, we set out to establish a novel method to tailor hand-based illusions to individuals' perceptual boundaries in a continuous fashion. We base our explorations on related literature, looking at physiological, interaction, and eye gaze data.

PHYSIOLOGICAL DATA Physiological data such as RSP, EDA, Heart Rate Variability (HRV), and so on, have been used as part of many VR systems. This type of data can be tracked noninvasively by, e.g., attaching electrodes to the human body. For example, only two electrodes at the fingertips are needed to track EDA. Using physiological data as an HCI research tool is well established and has received considerable attention [205]. It is often used for detecting emotional states [154], stress, and cognitive load [72, 122] without the need to directly ask participants about their experience. Additionally, researchers constantly improve hardware and software to enable ubiquitous and high-quality data acquisition in our daily lives through devices such as smart rings, watches, and neglects. More recently, we have also seen HMDs with integrated sensing capabilities, such as Galea⁵⁰ that offers EDA, ECG, EEG, and eye tracking, demonstrating commercial interest and the potential of collecting physiological data with consumer-grade devices.

Beyond subjective user feedback

In VR, physiological data have been used as part of adaptive systems that aid relaxation and reflection. For example, Amores et al. [9] developed *DeepReality*, a system which adapts the VR environment based on users heart rate, EDA and brain activity. Another domain is (affective) games, where we would like to highlight two systems, starting with *BreathVR* [318], which uses breathing as an input, providing a unique and engaging way to interact in single and multiplayer VR games, increasing presence and enjoyment. And secondly, the infamous *Brainball* system [152], in which a pair

Adapting systems through monitoring physiological responses

⁴⁹ Wikipedia webpage: <https://tinyurl.com/26rvhuk6>. Last accessed: Nov 1, 2024

⁵⁰ Galea webpage: <https://galea.co/>. Last accessed: Nov 1, 2024




Figure 2.31: Simplified representation of brain areas responsible for processing visual and haptic information (left) and international 10–20 EEG electrode setup (right). (Source: Raw images taken from Wikipedia⁴⁹ under public domain license. Annotations based on Mutschler et al. [249] p.688 and p.704).

controls a steel ball with their brain activity. Here, the goal of the game is to relax because the ball rolls towards the more mentally tensed player who, as a result, loses the game. Most relevant to our work are studies assessing users' experience inside IVEs.

*Metric for assessing
system immersion*

*Reponse to VR
sickness and stress*

For example, Egan et al. [78] investigated heart rate and EDA as an objective evaluation metric to assess the system immersion of a VR experience. Here, participants experienced the same virtual scene through an HMD vs. a 2D monitor. Differences in participants' heart rate and EDA were correlated to the display condition, which was also correlated with their subjective assessments of the experience. Heart rate and HRV have also been used to study motion sickness within the context of long-term immersion in VR [140, 231]. Marchiori et al. [232] observed an increase in HR in response to virtual scenes that were perceived as less realistic according to participants' questionnaire responses. Guna et al. [140] studied changes in SCR and Skin Conductance Level (SCL), a measure of EDA, to assess VR sickness, demonstrating a correlation between SCL and participants' subjective responses in a Simulator Sickness Questionnaire (SSQ) [177]. Jordan and Slater [169] investigated SCR during the onset of presence in a gradually forming environment, reporting SCR responds realistically to the scenario shown in the virtual environment. Here, participants standing on top of the tall column exhibited a significant increase in the number of SCR, whereas participants at the ground level showed no increase in SCR. On the other hand, respiratory rate changes with respect to humans' perceived stress, and therefore, has been studied when experiencing stressful scenarios such as flight or roller coaster simulations in VR [86]. These studies aim to establish an explicit objective metric, which is in line with our eventual goal of detecting noticeable hand-based illusions and consequently may allow us to tailor them to individuals' perceptual boundaries (RQ4).



Figure 2.32: Event-Related Potential (ERP) in VR. Left: Hand movement violations of avatar. (Source: extracted from Padrao et al. [265]; ©2015 Elsevier Inc.). Right: Haptic mismatches trigger ERPs of different magnitudes. (Source: extracted and modified from Gehrke et al. [118]; ©2019).

Most closely related to our work are EEG studies [118, 119, 253, 303, 306], concerning ERPs caused by errors and mismatches. In the HCI and VR community, ERPs have become a useful measure, allowing researchers to detect if participants experience an error without directly asking them [197]. For example, Si-Mohammed et al. [303] showed that ERPs can be used to detect system errors in VR such as background anomalies or tracking errors while interacting with virtual objects. They further demonstrated that single trial classification can achieve average prediction accuracies of 59.3% and 84.9% depending on the type of error. Gehrke et al. [118] used ERPs to detect visuo-haptic mismatches by comparing three levels of haptic immersion when touching a virtual cube: (1) visual but no haptic feedback, (2) visual feedback + vibrotactile feedback on the fingertip and (3) visual feedback + vibrotactile feedback + electro-muscle stimulation. Their results show that it is possible to distinguish matching and mismatched trials using ERP's error peak negativity (see Figure 2.32: right). In another study by Yazmir et al. [374] ERPs were used in a visuo-haptic error-induced task, where participants used a *PHANTOM* haptic device to move a sphere horizontally. To get to the target location, one must pass an obstacle that momentarily obscures the sphere. In about 60% of the trials, they introduced a disturbance, i.e., the sphere was offset, horizontally and/or vertically, 'behind' the obstacle. In a small study, they collected evidence for a strong ERP shortly after the error. Padrao et al. [265] studied the difference between self-generated and externally imposed errors on the sense of agency when users were embodied by a virtual avatar. Their work provides a strong foundation for our investigations in Chapter 6 because the externally imposed errors were provoked by moving the virtual avatar hand in the opposite direction from the participant's real hand (see Figure 2.32: left). The results showed a strong similarity to ERP signatures related to semantic or conceptual violations (central cortex area).

EEG as a direct correlate of user experience

Semantic and haptic mismatches in VR

Body movement violations show similar results

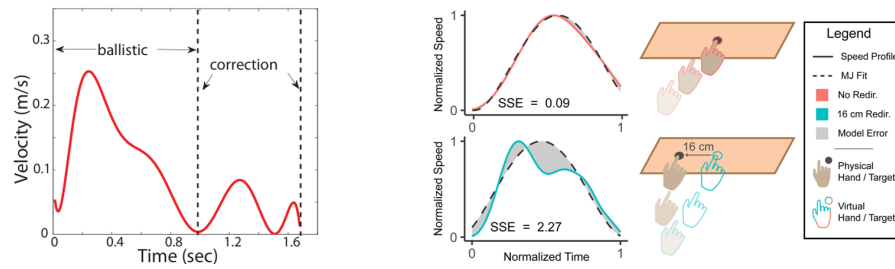


Figure 2.33: Hand movement velocity in VR. Left: Velocity profile during aimed hand movements in VR. (Source: extracted and cropped from Liu et al. [221]; ©2009 IEEE). Right: Velocity profile during aimed hand movements under the influence of horizontal HR. (Source: extracted from Gonzalez et al. [131]; ©2019).

INTERACTION & EYE GAZE DATA Aimed real-world hand movements have two distinct movement phases, ballistic and (an optional) correction [221], with predictable hand speed and arrival time using models such as the Minimum-Jerk model. Gonzalez et al. [131] found that horizontal HR of noticeable magnitudes significantly worsens Minimum-Jerk model fit, suggesting that prediction models should be adjusted for redirection (Figure 2.33: right). Not only hand speed and arrival times are affected by HR, but also by reach trajectory as shown in Figure 2.24 top right. In their follow-up work, Gonzalez and Follmer [134] incorporated a variety of distinct properties of redirected movements and presented an approach to predict a user's hand trajectory during redirected movements based on stochastic optimal feedback control. While these studies inform the selection of our movement features in Chapter 6, they did not incorporate the aspect of noticeability into their investigations. Notably is the work by Lebrun et al. [207], who present an empirical model predicting the hand trajectory as a function of the redirection. The presented model achieves robust prediction of detection probabilities for individuals. This string of research was executed parallel to this dissertation, demonstrating the importance and relevance of this research problem in the scientific HCI/VR community.

*Do hand movements
under the influence
of HR reveal their
noticeability?*

*Eye gaze as an
implicit metric for
HR DTs*

Eye gaze is commonly used in VR systems to select targets or interact with virtual content [201]. In the context of HR it has been used in haptic retargeting applications. For example, for target prediction to seamlessly redirect users' hands between physical proxies [52, 237]. This works because findings from interaction studies showed that during targeted movements, participants predominantly fixate their gaze on objects that are relevant to the task while devoting minimal attention to their own hand when reaching for an object [203]. Jordan and Slater [169] used eye scan path entropy as a correlate of feeling present in virtual environments, demonstrating the potential of this measure as a response variable.

2.5.3 Conclusion

Hand-based illusions that build on HR are an inexpensive and effective way to enhance haptic feedback and interactions in VR. The scope of presented techniques in the literature still leaves room for innovation, especially in combination with proxy objects. Most relevant to this dissertation are *Visuo-Haptic* [20] and *Pseudo-Haptic* [227] techniques that can change proxies perceived properties without requiring physical modification. Similar to AHF, the techniques promise high control over the provided stimuli and can dynamically adapt to the situation due to their inherent software-based nature. They have the potential to address the challenges of *Colocation* and *Similarity* [256], with the latter being the focus of Chapter 4.

Hand-based illusions to overcome physical limitations of proxies

However, the techniques come at a risk of disrupting the immersive nature of the experience if the applied offsets exceed users' perceptual limits. As a result, the research community established CDT for HR [384], providing estimates for the amount of sensory discrepancy between vision and proprioception that can be applied. While these conservative lower bounds can be safely used without risking disrupting the VR experience, they severely limit the types of illusionary (haptic) effects that can be created. There is an inherent trade-off between the detectability of applied offset and the magnitude of the illusion that can be achieved. Thus, researchers set out to explore DTs in various settings relevant to any interaction. Together with the work in this dissertation, particularly in Chapter 5, DTs manifest themselves as highly dependent on the interaction, the environment, and the individual user.

Introduced offsets depends on several unknown factors

The previously outlined studies, as well as in the upcoming experiments in this dissertation, use psychophysical methods to determine DTs to compare factors and investigate perceptual effects. However, these DTs only work in the specific study setups and lack practical relevance, given the complexity of constantly changing demands of IVEs. Users' body responses and their interactions have been shown to correlate with their subjective experiences of an IVEs. Inspired by this, we set out to explore a novel method for continuously monitoring and adaptation of the undetectable visuo-proprioceptive offset during hand-based illusions in VR in Chapter 6.

Current methods lack continuous adaption to dynamic IVEs

CHANGING PROPERTIES OF HAPTIC PROXIES

We open-source the *VoxelHap* toolkit to the community¹. Videos about the work and prototypes presented in this chapter are available online and can be accessed through the QR code. Images and parts of the text in this chapter, as well as the presented figures, tables, ideas, concepts, implementations, applications and uses cases, studies and experiments, results, discussions, and conclusions, have been published previously in:

[92] **Martin Feick**, Cihan Biyikli, Kiran Gani, Anton Wittig, Anthony Tang, and Antonio Krüger. *In Proceedings of UIST 2023*. VoxelHap: A Toolkit for Constructing Proxies Providing Tactile and Kinesthetic Haptic Feedback in Virtual Reality.



VoxelHap video

[93] **Martin Feick**, Donald Degraen, Fabian Hupperich and Antonio Krüger. *In Frontiers in Virtual Reality 2023*. MetaReality: Enhancing Tactile Experiences using Actuated 3D-printed Metamaterials in Virtual Reality.

3.1 CONSTRUCTING MULTIMODAL HAPTIC PROXIES



Figure 3.1: Lisa wants to use a power drill for training in VR, requiring a dedicated proxy that replicates form, feel, and function.

3.1.1 Introduction

Remember Lisa, our medical student who wants to improve her surgical skills in a safe and responsible manner. Lisa can use the VR training program, but she would need a dedicated controller for each of the tools she wants to practice with. A single reconfigurable device could address this problem. It could provide realistic haptic feedback to be a scalpel or a power drill because it can change its shape and function to adapt to the particular needs of the moment. This is where *VoxelHap* comes in, allowing Lisa to build a functional proxy

*Motivation for
VoxelHap*

¹ VoxelHap GitHub: <https://tinyurl.com/mp44kkpz>. Last accessed: Nov 1, 2024




Figure 3.2: BaseVoxel consists of a PCB in a block design, a translucent 3D-printed shell assembled with M3 screws, and conductive touch pads.

What is VoxelHap?

power drill using reconfigurable blocks (Voxels), approximating size and shape to allow for embodied interaction (see Figure 3.1). As she wants to practice working with the power drill, the proxy needs to support the desired functions: for example, provide tactile feedback when pressing the trigger button or kinesthetic haptic feedback when securing the drill bit. Lisa can add this functionality by using two special types of Voxels: Vibration and Rotation. Finally, she increases the shape fidelity of the proxy by adding 3D-printed ShapePlates with the aim of increasing overall similarity and usability. In summary, *VoxelHap* is a block-based construction kit that enables users to build fully functional proxy objects that deliver tactile and kinesthetic haptic sensations. In its current implementation, *VoxelHap* combines and supports a range of haptic modalities while enabling realistic input and output controls. Our main goal with *VoxelHap* proxies is to allow users to construct proxies with great expressibility, combining multimodal kinesthetic and tactile haptic feedback.

3.1.2 *VoxelHap Toolkit*

How does VoxelHap work, and what is it made of?

VoxelHap is a toolkit that gives users the ability to create haptic proxies based on a set of building blocks. These building blocks give the proxies different physical and interaction capabilities. *VoxelHap* is an end-to-end system that offers a software tool that supports the assembly process by providing visual guidance to efficiently construct the desired proxy. The system hosts an exact representation of each proxy in its current configuration in real time. At *VoxelHap*'s core, we have BaseVoxels, the most basic reusable unit. To increase haptic resolution, Voxels and Plates with special functionalities and properties can be added on demand. Below, we discuss the design rationale and fabrication of *VoxelHap*'s components.

VOXELS We use cube blocks (Voxels) as our basic building structure [91, 183, 258, 284, 294, 298, 357], because cubes are a well-established approach for construction kits, and have been

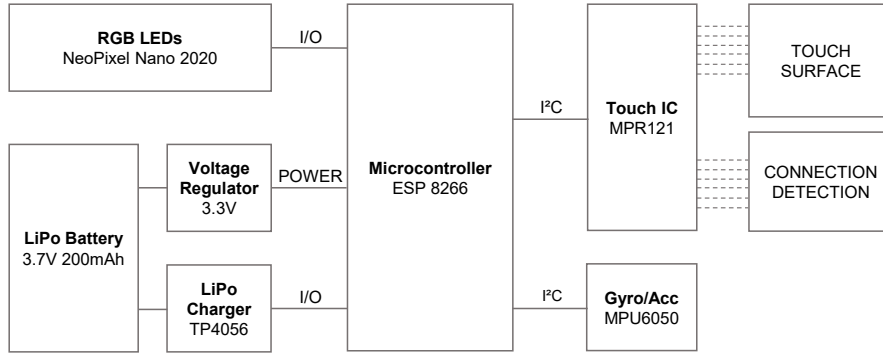


Figure 3.3: Schematics of a BaseVoxel.

demonstrated to be effective. Voxels can be connected to each other genderlessly using a custom-designed ConnectionPlate mechanism (i.e., they can be connected in any direction and in any orientation). Inspired by *ActiveCubes* [357], each Voxel is a self-contained unit and is designed for input and output. In this work, we present two additional types of Voxels that provide additional (corresponding) capabilities, VibrationVoxel and RotationVoxel. The toolkit can be extended to include new capabilities, e.g., twistable or bendable units or other types of sensors as previously shown in other works [16, 91, 239, 357]. Voxels are rechargeable and flashable through a custom docking station with pogo pin connections (see Figure 3.2 and Figure 3.4).

Blocks build the foundation of each proxy

BaseVoxel. The BaseVoxel offers the basis for all other types of Voxels. It measures 37×37 mm (outer shell) and weighs 38 g. Each BaseVoxel consists of six FDM 3D-printed translucent PETG shells of 3 mm thickness secured with M3 screws and has touch-sensitive areas on all sides, which have been 3D-printed with conductive filament. It hosts an Espressif ESP8266 (ESP12F Module), a TP4056 LiPo Charger IC, an MPU6050 six-axis gyroscope and accelerometer, a voltage regulator (LDL212D33R), a proximity capacitive touch sensor controller (MPR121), six NeoPixel Nano 2020 LEDs to illuminate the sides, and a 3.7 V 200 mAh LiPo for powering the Voxel. The PCB comes in a cube shape (33×33 mm) with a cut-out part for the antenna to ensure a proper fit inside the 3D-printed Voxel; the high-level schematics can be found in Figure 3.3.

Simplest self-contained unit

VibrationVoxel. This type of Voxel enables the rendering of tactile feedback when holding or touching the proxy, weighing 46 g. We implemented this by extending the BaseVoxel by a voice-coil actuator (HAPTIC™ Reactor ALPSALPINE AFT14), a Bluetooth 5.0 low latency (45 ms) audio receiver and an audio class D amplifier (PAM8403). Our PCB design ensures that the actuator can be placed and affixed inside the PCB cube. For ease of use, the vibration

Can render tactile output




Figure 3.4: Displaying how the PCB fits inside the 3D-printed shell. The ConnectionPlate adapter of the RotationVoxel. The pin layout, LED and touch sensing pin, and right next to it, the charging and flashing unit with pogo pins (left to right).

patterns (audio files) are sent directly from the PC via Bluetooth, i.e., once the VibrationVoxel is running, it can be used as any other audio output device.

*Provides kinesthetic
force feedback*

RotationVoxel. This type of Voxel enables proxies to have rotational parts [16, 91, 218], but also renders kinesthetic force feedback for rotational movements. To achieve this, we added a brushless DC Motor2204 260KV, an ICTMC6300 motor driver, and a magnetic position sensor (MAQ473) to the BaseVoxel's PCB. The RotationVoxel can be used passively, e.g., for sensing, or as an active component to provide resistance when turning it or even blocking involuntary rotations. To enable a solid connection, one side of the BaseVoxel and the PCB was removed to offer space for the ConnectionPlate adapter (see Figure 3.4). Total weight is 65 g.

In summary, block-like structures enable rapid prototyping of proxies within several minutes [91], creating functional low-fidelity versions of virtual models that may be sufficient in some situations, for example, to better understand the depth and scale of virtual scenes [248]. However, many interactions with objects require higher shape similarity [184].

*Connect Voxels to
each other*

PLATES Plates are optional passive components that can increase the haptic resolution and functionality of proxies. These allow Lisa to achieve high shape fidelity, which can be used for realistic training and simulations. We took inspiration from Mueller et al. [247]'s work on *faBrickation*, because similar to *VoxelHap*, the authors use basic block structures (LEGO®) to improve fabrication and prototyping time. Additionally, they 3D-printed high-fidelity parts for accurate shape approximation. Within *VoxelHap*, we currently provide 5 different types of Plates as depicted in Figure 3.5 and Figure 3.6: Connection-, Shape-, Texture-, Weight- and TrackingPlates.

ConnectionPlates. Required to connect Voxels to each other. They are 3D-printed using conductive filament, creating a closed circuit



Figure 3.5: 3D-printed ShapePlates of the power drill that can be attached to the low fidelity power drill to increase shape resolution.

upon connection, allowing the system to determine the side and the type of Voxels that got connected to each other (see Figure 3.6).

ShapePlates. Can be used to increase proxy fidelity if needed (see Figure 3.5). For instance, any shape can be 3D-printed onto the Plates to increase shape resolution. ShapePlates can be generated by first applying voxelization and second, projecting anything below a 45° angle onto the ShapePlate because that is the maximum overhang that can be reliably 3D-printed without a support structure. If 3D-printed with conductive filament, we can also detect touch input on ShapePlates, since they make contact with Voxel's touch pads, e.g., the power drill's trigger button shown in Figure 3.1.

Increase shape resolution

TexturePlates. Plates that are augmented with material textures such as fabric or rubber. They can be used to simulate local differences in textures to, e.g., improve the grip when holding a tool or, depending on the use case, create various touch sensations (see Figure 3.6: left). TexturePlates can also be combined with other fabrication techniques, such as 3D-printed hair structures by Degraen et al. [69] or metamaterials [93], to create different material texture sensations—the advantage being that it can be 3D-printed all together.

Renders tactile features

WeightPlates. The overall weight and center of mass are important properties of proxies [382] and hence, *VoxelHap* offers Plates to 'balance' them. To do so, we embedded lead into the Plates, allowing users to increase the overall weight. Moreover, the location of the WeightPlate will shift a proxy's center of mass. A single WeightPlate adds up to 45 g to the proxy (see Figure 3.6: middle).

Enhances kinesthetic feedback

TrackingPlates. Voxels sense their orientation and acceleration, but for accurate mapping between real-world proxy and virtual model,




Figure 3.6: TexturePlates, WeightPlate and ConnectionPlate.

a shared coordinate system is required. Here, TrackingPlates offer mounts for systems such as an HTC Vive tracker or Optitrack marker to enable robust 6-DoF tracking. Overall, we fabricated 13 Voxels (nine BaseVoxels, three VibrationVoxels, and one RotationVoxel) and several Plates of each type.

*Facilitates the
assembly process*

HAPTIC PROXY DESCRIPTION FORMAT While haptic proxies can be constructed haphazardly, we need a structured format to describe the appearance, behavior, and functionality of block-like proxies. For this purpose, we introduce the Haptic Proxy Description Format (.hpdf), a semi-generalizable description format for proxy design, which includes a proxy's kinematics, a functionality log, and a construction plan. To this end, users need to manually create the .hpdf once and thence, they can re-use or share it.

Kinematics. 3D models for different parts of the object as .fbx files, especially for functional parts. This includes physical properties of object parts such as their texture, material, and mass, but also mechanical features to allow realistic approximation and animations.

Functionality Log. Contains a list of available parametric input and output controls for the proxy such as button press and the resulting visual and haptic renderings.

Construction Plan. A visual construction plan where a user can decide between three levels of proxy fidelity: (1) a low fidelity approximation only requiring BaseVoxels, (2) high shape fidelity by adding ShapePlates, and (3) functionality through the use of Vibration- and/or RotationVoxels.

*How to construct
VoxelHap proxies*

CONSTRUCTION WALKTHROUGH In the following, we illustrate the construction process (Figure 3.7), where Lisa downloads the .hpdf file for the desired power drill. As soon as she powers the Voxels, they connect wirelessly and appear in the virtual environment. As she engages with the assembly process, two Voxels light up to indicate which sides should be connected. As Lisa follows the step-by-step instructions, different sides light up until she is finished with the low-fidelity proxy. She wants to increase the shape resolution, so she selects the ShapePlate option. As a result, she needs to 3D print the




Figure 3.7: *VoxelHap*'s supported construction process. *VoxelHap* provides a 3DUI, allowing users to select the desired fidelity. The basic construction plan only uses BaseVoxels, visualizing the Voxels' sides that should be connected through LED indication.

missing pieces to achieve the desired shape resolution. Similarly, she can select the functional features for the proxy. Lisa opts for the trigger button and three haptic renderings: vibrotactile feedback when using the trigger button, as well as for the power drill running in idle and power drill mode. For the latter, she can map interactions with virtual objects to their corresponding haptic feedback pattern; for instance, the haptic feedback when power drilling wood is different from that of steel or concrete. After Lisa has practiced using her power drill proxy, she wants to try a different type of power drill, e.g., one with two handles, to improve stability. Thus, she simply goes through the steps outlined before, reconfiguring her proxy.

*Offline
reconfiguration*

3.1.2.1 Constructing *VoxelHap* Proxies

In this section, we showcase *VoxelHap* proxies' expressiveness power across three dimensions: types of functionality, haptic feedback, and geometry. We did this as a first evaluation of the toolkit as proposed by Ledo et al. [212]. A video of how users interact with all presented *VoxelHap* proxies can be accessed through the QR code at the beginning of this Section 3.1.1.

TYPES OF FUNCTIONALITIES One of the *VoxelHap* toolkit's main contributions are proxy replicas that combine multimodal tactile and kinesthetic haptic feedback. To illustrate this, we built two functional proxies, a rattle, and a coded dial lock.

*Examples of
constructed proxies
using VoxelHap*

Rattle. We chose a matraca toy (see Figure 3.13) because it nicely demonstrates tactile and kinesthetic haptic feedback within a single interaction. When rattling, a user can feel the moving parts (inertia) and the distinct clacks of the cogwheel. This is an example of how




Figure 3.8: *VoxelHap* rattle. A matraca rattle toy, constructed using only Voxels (cubes) and the same proxy with ShapePlates to improve ergonomics and realism. Low shape fidelity, high shape fidelity, and corresponding virtual model (left to right).




Figure 3.9: *VoxelHap* dial lock. Low shape fidelity, high shape fidelity, and corresponding virtual model (left to right).

Combining tactile
and kinesthetic
haptic feedback

VoxelHap can support dynamic touch [349]. We assembled this proxy using four BaseVoxels, a VibrationVoxel, and the RotationVoxel. The BaseVoxels allow a rough approximation of the rattle, but more importantly, the RotationVoxel enables the rotating parts and thus renders the kinesthetic feedback (inertia). Tactile sensations of the cogwheel are then added at each cog (35°) using the VibrationVoxel. By 3D printing ShapePlates, the handle becomes more realistic and ergonomic (see Figure 3.8). This is an ungrounded haptic device.

Coded Dial Lock. A coded dial lock is another example where tactile feedback (tick marks) and kinesthetic feedback (resistance) are needed to provide compelling haptic feedback. The *VoxelHap* proxy consists of six BaseVoxels to create a supporting structure, approximating the height of a stationary dial lock, a VibrationVoxel to haptically render the tick marks, and the RotationVoxel for sensing and for providing basic resistance when turning (see Figure 3.9). This is an example of a grounded haptic device.

TYPES OF HAPTIC FEEDBACK Using the VibrationVoxel and RotationVoxel with touch/force sensing capabilities gives us the ability to create a variety of highly synchronized haptic impressions and hand-based illusions. Here, we present a set of functional proxies that we constructed using *VoxelHap*.




Figure 3.10: *VoxelHap* salt shaker. Low shape fidelity, high shape fidelity, and corresponding virtual model (left to right).




Figure 3.11: Pseudo-haptic button, stiffness and size variations. We can haptically render a virtual button press or change the perceived stiffness of a rigid Voxel by coupling the touch input to tactile output and the corresponding pseudo-haptic effect.

Pseudo-Haptic Illusions. Combining pseudo-haptic effects [227] with vibrotactile actuation [148] can result in realistic haptic sensations. As shown in Figure 3.11, we re-implemented Park et al. [267] simulated physical button press using *VoxelHap*, coupling touch sensor input (pressure) to the corresponding pseudo-haptic visualization and tactile feedback. Similarly, we can alter the perceived stiffness of the VibrationVoxel when pinching it. This can be achieved by mapping the force input to a mesh deformer² and the corresponding vibrations analogously to [6, 37, 179]. For this type of interaction, we can also alter the size of the virtual model while using the same proxy. Bergström et al. [30] showed that virtual objects can be up to 50% larger or 10% smaller compared to their physical counterpart. These illusions are depicted in Figure 3.11.

Adding existing illusion techniques to improve haptic rendering

Rendering Textures. Interacting with surface textures results in unique haptic feedback. For instance, the feeling when drilling wood differs significantly from that of drilling metal. Our *VoxelHap* power drill proxy (see Figure 3.1) allows users to perceive this difference through varying vibrotactile feedback [328]. Finally, TexturePlates can be used to provide tactile feedback when touching proxies in an

Vibrotactile rendering of mid-air textures

² Unity assetstore: <https://tinyurl.com/ykhcpevh>. Last accessed: Nov 1, 2024




Figure 3.12: *VoxelHap* mug. BaseVoxel, VibrationVoxel, and a mug handle using ShapePlates. Full shape resolution using ShapePlates. Added RotationVoxel and WeightPlate for moving mass simulation (left to right).

encounter-type fashion [12, 328], i.e., rotating to the corresponding textures when needed.

Simulating Mass. We can simulate coffee sloshing around in a mug, e.g., through vibrotactile feedback [335]. To demonstrate this, we built a mug consisting of a VibrationVoxel, a BaseVoxel, and ShapePlates (see Figure 3.12). Even more realistic sensations can be achieved by adding the RotationVoxel and a WeightPlate. Synchronizing the moving WeightPlate with the interaction creates a realistic feeling of inertia [287, 346]. This effect may be strengthened when combined with pseudo-haptic weight techniques [288].

Kinesthetic effect of moving mass

TYPES OF GEOMETRIES *VoxelHap's* core capability is limited by the Voxel's size. To increase shape resolution, one may add ShapePlates; however, this requires additional fabrication time. Moreover, the block-based approach favors certain types of geometries. Our current implementation benefits convex object geometries but only for objects bigger than the BaseVoxel size of 3.7 cm. Any arbitrary shape can be 3D-printed onto the ShapePlates to resemble smaller features of the virtual object, for instance, a trigger button, handles, or even other geometric primitives e.g., spheres, triangles or pyramids. One limitation is that the construction of high-fidelity objects such as pens or screwdrivers [218] may be 3D-printed onto a ShapePlate, but interactions will often be limited to touch and exploration rather than function. On the other hand, concave geometries can also be achieved. For instance, a ring may be built with BaseVoxels and ShapePlates, but this may quickly become bulky, depending on the use case.

Shapes can be attached to Voxels to increase resolution

3.1.2.2 Implementation

VOXEL FIRMWARE The Voxel software is implemented in C++ and uses an ESP-WIFI MESH³ to communicate with a master (ESP32), which is tethered to a host machine using serial port communica-



Figure 3.13: Recording vibrations from real-world object interactions using a condenser vibration pickup microphone.

tion. ESP-MESH has low power consumption, low latency, and great scalability, theoretically up to 1000 Voxels. There exist two template commands, for sensing and actuation, which are used to enable seamless communication between Voxels and the master. Our master application offers an exact representation of the current types of Voxels used, their functionality, and their status. Moreover, we can support and monitor multiple proxies at the same time. Once a user powers on a Voxel of any type, it automatically registers at the master with a unique ID and calibrates itself. Voxels can detect their neighbors when being connected to each other using conductive pins on the ConnectionPlates and pads on the underlying PCB. They constantly send updates to the master (baud rate 115200) to enable seamless interactions. The RotationVoxel was controlled using SimpleFOC [309].

VoxelHap's firmware consists of multiple parts

VIRTUAL ENVIRONMENT Our Voxel pipeline works with any machine and software that supports serial port communication. We implemented our virtual environment using Unity3D (v.2021.3.10f1). Sensing updates are received asynchronously and can be mapped to the desired visualization. Actuation commands can be triggered within the virtual environment and are then forwarded to the corresponding Voxel. To this end, we provide a basic set of functions as illustrated before; however, additional functionality can be added as needed. The virtual model was voxelized using the mesh voxelizer asset⁴. Currently ShapePlates have to be manually post-processed in CAD software. When initiating the construction process, the system sends two actuation commands (i.e., turn LED on) to a pair of Voxels. This lights up one side on each Voxel, indicating which sides should be connected [183]. Connecting to the wrong side will prevent users from continuing. Once either Voxel has established a connection, they send a confirmation event. The construction plan retrieves information from the manually generated .hpdf file.

How to render VoxelHap proxies input and output

RECORDING VIBROTACTILE FEEDBACK Vibrotactile feedback can offer rich haptic feedback but is challenging to program [66]. Inspired

³ Espressif: <https://tinyurl.com/4xf7694d>. Last accessed: Nov 1, 2024

⁴ Unity assetstore: <https://tinyurl.com/2km4bsua>. Last accessed: Nov 1, 2024

*Haptic feedback can
be recorded from
real-world
interactions*

by Tanaka et al. [335], we attached an AKG C411 high-performance miniature condenser vibration pickup microphone to the real-world objects, recording the vibrations caused by the interactions on a Zoom H4n Pro, as depicted in Figure 3.13. Through early pilot testing, we found that simply replaying the recorded sounds on the actuator did not result in convincing haptic feedback. The sensations felt rather weak, and pilots struggled to distinguish them. Therefore, we modulated the recorded sound using a square-shaped wave with an amplitude corresponding to the recorded volume (in dB). We synchronized the vibrotactile feedback using discrete key points in the interaction.

3.1.3 Technical Evaluation

*How does it compare
to existing toolkits?*

We evaluated the technical capabilities of *VoxelHap* with respect to the following aspects. The table below summarizes the results: (1) according to ESP-MESH up to 1000 Voxels; however, we can only guarantee 13 Voxels; (2) measured weight of the proxies (excluding trackers), except the dial lock, because it was stationary; (3) measurements were taken using a PCE-FM 50N Series force gauge, and we report average values after three repetitions. (4) and (5) are hardware limitations; (6) fabrication costs, including all hardware components. (7) was estimated by combining the latency of the different components. (8) is runtime under 50% load when fully charged. To further evaluate the toolkit, we conduct two user studies, comparing *VoxelHap* against the current state of the art, following Ledo et al. [212]’s proposed holistic evaluation strategy for toolkits. First, we want to understand the potential benefits of coupling function with the corresponding tactile and kinesthetic haptic feedback provided by *VoxelHap*. Second, we study *VoxelHap*’s high shape approximation feature ShapePlates.

3.1.4 Experiment 1—The Impact of Functionality

*VoxelHap vs.
controller*

*The impact of a
function*

We designed the first experiment to compare low-fidelity BaseVoxel approximations with combined functional multimodal haptic feedback against the current state of the art, a standard Vive controller that provides basic vibrotactile feedback. We allowed participants to customize the strength of the vibrotactile feedback for both devices until it felt realistic to them. We included this to collect more insights, potentially informing future research on designing vibrotactile feedback for VR. Here, we only use Base and Functional Voxels. Please note that the Vive controller uses a rumble motor instead of a voice-coil actuator. Therefore, we could not play the recorded sounds. We opted for the three *VoxelHap* proxies because they combine tactile and kinesthetic haptic feedback, differ in dimensions, and cover two types of devices, grounded as well as ungrounded. We decided against proxies that utilize hand-based illusions because they need to

1	Maximum # of Voxels	13
2.a	Shaker Weight	85 g
2.b	Rattle Weight	339 g
2.c	Shaker Weight With ShapePlates	105 g
2.d	Rattle Weight With ShapePlates	368 g
3.b	ConnectionPlate Strength Vertically	36.56 Nm
3.c	ConnectionPlate Strength Horizontally	7.87 Nm
3.d	ConnectionPlate Req. Connection Force	1.26 Nm
3.e	RotationVoxel Rotational Torque Range	0.3 – 17.2 Nmm
4	RotationVoxel Rotational Range	∞
5	RotationVoxel Max. Rotational Velocity	260 RPM
6.a	BaseVoxel Fabrication Cost	€19.87
6.b	RotationVoxel Fabrication Cost	€58.28
6.c	VibrationVoxel Fabrication Cost	€37.99
7.a	BaseVoxel & RotationVoxel Est. Latency	48 ms
7.b	VibrationVoxel Est. Latency	56 ms
8.a	BaseVoxel Avg. Battery Power	82 min
8.b	RotationVoxel Avg. Battery Power	26 min
8.c	VibrationVoxel Avg. Battery Power	63 min

Table 1: Summary of technical evaluation.

be carefully calibrated to an individual’s perceptual limits in order to remain unnoticeable [99]. Otherwise, this might affect participants’ assessments. The duration of the vibrations was tailored to fit the interaction, and strength was adjusted by changing the amplitude.

3.1.4.1 Task & Proxies

We implemented three mini-games, highlighting *VoxelHap*’s functional capabilities. We used the Rattle, Dial Lock and Shaker proxies shown above. Since we only fabricated a single RotationVoxel, we built a second stationary device with the RotationVoxel’s hardware, except the PCB. This was used for the safe dial lock. In the following, we describe the three mini-games.

DIAL LOCK The goal of this task is to unlock the safe. To do so, participants had to rotate the dial to different positions highlighted in red. They had to remain in that position for 500 ms before the next target appeared. Six combinations had to be solved to unlock the

safe. The order was randomized; however, the total rotational travel distance required to complete the task remained equal.

*Tasks and proxies
used in our
experiments*

RATTLE Here, participants were asked to cheer for their favorite team by rattling with a specific rotational speed in order to fill the progress bar. The progress bar filled quicker if they stayed within the correct speed range. This was done to limit participants' maximum rotational velocity while still allowing them to focus on the haptic feedback provided.

SHAKER Participants were asked to shake salt into the highlighted pots. The amount of salt required to complete one pot was set to three successful shakes, i.e., at least 75% of the salt needed to land in the pot. To finish the game, six completed pots were required, resulting in $3 \times 6 = 18$ interactions. The order of the highlighted pots was randomized to make the game more unpredictable.

3.1.4.2 *Design*

We used a within-subjects design. We had three conditions: Dial Lock, Rattle, and Shaker, each performed once with a *VoxelHap* proxy and the standard Vive controller. The conditions were counterbalanced using a Latin square, and we alternated the order of *VoxelHap*/controller with each participant. This way, we ensured that each participant was able to directly compare devices.

3.1.4.3 *Participants*

*Diverse set of
participants*

We recruited twelve right-handed participants (four females, eight males), aged 18–27 (mean = 24.15; SD = 3.02) from the general public and the local university. Participants had a range of different educational and professional backgrounds, including computer science, linguistics, and data science. All participants reported normal or corrected-to-normal vision and did not report any known health issues that might impair their perception. Three participants had never used VR before, six had used it a few times (1–5 times a year), two reported using it often (6–10 times a year), and one participant used it on a regular basis (more than 10 times a year). Participants not associated with our institution received €10 as remuneration for taking part in the experiment. The study was approved by the Saarland University's Ethics Board.

3.1.4.4 *Apparatus*

We used an apparatus consisting of an HTC VIVE Pro Eye tracking system with SteamVR (v.1.22) and OpenVR SDK (v.1.16.8). The simple virtual scene was developed in Unity3D (v.2021.3.10f1). We used an Acer Predator Orion 5000 PO5-615s offering an Intel® Core i9 10900k

CPU, 32 GB RAM and an Nvidia® GeForce RTX 3080 for running the experiment. *VoxelHap* proxies were tracked using Vive trackers (v.3). In order to avoid potential issues with battery power, we connected *VoxelHap* proxies to a power source for the study.

3.1.4.5 *Experimental Protocol*

Participants were given a general introduction to the study. Following this, we obtained consent, and asked them to fill in a demographics questionnaire. Next, we showed and explained the *VoxelHap* proxies to the participants to ensure that they were familiar with the devices. Then, they entered the VR environment and were guided through a practice round before they started with the main task. During the warm-up round, they were asked to adjust the strength of the vibrotactile feedback for both devices until it felt realistic to them. Participants were instructed to perform a mini-game with each of the devices. After they had completed the mini-game, a questionnaire appeared, assessing their experience with the device. After each condition (i.e. they finished a mini-game with both devices) participants were asked to name their favorite and the most realistic device and, if possible, explain why they preferred it over the other. The total experiment took about 45 min.

Participants received the following instructions

3.1.4.6 *Data Collection*

We collected data from five sources: a pre-study questionnaire for demographic information; field notes and observations; configured strength of haptic feedback; a questionnaire after each task and condition in VR using our *VRQuestionnaireToolkit* [94] and a semi-structured interview to better understand participants' experiences with the system. The questionnaire items were adapted from prior work [91, 382]:

- (1.1) My interactions felt realistic.
- (1.2) The vibrations were in sync with my interactions.
- (1.3) The vibrations matched my visual impression of the game.
- (1.4) I enjoyed playing the game.
- (1.5) Overall impression of the game.
- (2.1) Which of the two devices did you enjoy more?
- (2.2) Which of the two devices felt more realistic?

Questionnaire items were taken from related work

3.1.4.7 *Analysis*

Statistical tests were chosen based on whether the data satisfied parametric test assumptions at $\alpha = .05$ using Shapiro–Wilk tests and QQ plots. For outlier removal, we used the box plot method. We corrected pairwise comparisons using Bonferroni–Holm adjustments. Semi-structured interviews were coded for qualitative analysis.




Figure 3.14: Experiment 1 results. *VoxelHap* proxies against Vive controller. * = $p < .05$, ** = $p < .01$, *** = $p < .001$. *VoxelHap* received significantly higher ratings for realism than the Vive controller for all three proxies.

3.1.5 Results

Overall, our three *VoxelHap* proxies, Dial Lock, Rattle, and Shaker, received better scores compared to the standard Vive controller across all five questions regarding realism, synchronization, haptic feedback, enjoyment, and overall impression (see Figure 3.14). Differences were most clear for the Rattle, which saw the most significant differences, followed by the Dial Lock and the Shaker. Participants reported significantly higher levels of realism and enjoyment with the *VoxelHap* Rattle. This is supported by our analysis on which device participants associated with the highest realism. As depicted in Figure 3.14, participants overwhelmingly answered *VoxelHap* proxies. When asking participants which device they enjoyed the most, the *VoxelHap* Rattle and Dial Lock received comparable results, despite the Shaker being on par (see Figure 3.17). Here, some participants stated that the ergonomics of the Vive controller and the basic haptic feedback were “good enough” (P8) to play the game. Further, they expected the Dial Lock to have “more friction” (P6), which we intentionally kept low. Finally, participants chose *VoxelHap* proxies to feel more realistic: for example, the *VoxelHap* Rattle, due to its “realistic and smooth motion” (P1) and the “weight balance [feeling] so nice” (P5), because “you can really feel the weight moving” (P7). The *VoxelHap* Shaker was selected because “shape and vibration were more realistic” (P6) and had a “similar shape” (P3) than a real salt shaker. Finally, participants preferred the *VoxelHap* Dial Lock because it “was more intuitive” (P3) and “you could actually feel it rotating” (P11). Our study demonstrates that combining low fidelity approximation and function significantly increased realism in VR and that *VoxelHap* is an effective toolkit to achieve this.

VoxelHap
outperforms
controller

Participants
comment on
VoxelHaps' pros and
cons

3.1.5.1 Vibrotactile Feedback

We asked participants to configure the strength of the vibrotactile feedback during the warm-up phase in 10% increments (0%–unnoticeable; 100%–strongest vibration possible). They were instructed to select a strength that felt realistic to them. The results for both Vive controller, ($r(34) = -.60$, $p < .001$), and *VoxelHap*, ($r(34) = -.49$, $p = .002$), suggest a strong negative correlation

between object and configured vibration strength, meaning that for the Shaker participants selected weaker vibrations than for the Dial Lock and the Rattle, respectively (see Figure 3.17: right). This is in line with our assumption that the three study objects differ in their ‘expected’ feedback. Together with the findings above, receiving high ratings in synchronization and greater average scores for the provided vibrotactile feedback, this demonstrates that *VoxelHap* proxies effectively produce realistic haptic sensations.

Voxels can effectively render tactile feedback

3.1.5.2 Summary

Our study showed that *VoxelHap* proxies outperform the current standard VR controllers. This demonstrates that *VoxelHap*’s concept and implementation is effective and robust. However, we frequently observed uncanny valley of haptics effects [29]—where small differences affected participants’ assessments, e.g., the low friction of the Dial Lock. Even though participants preferred the direct mapping between *VoxelHap* proxies, the virtual model and their interactions, they often highlighted the Vive controller’s better ergonomics. This is not surprising because block-like structures can be difficult to hold (e.g., the handle of the rattle). With this in mind, we designed a second experiment, hypothesizing that adding *VoxelHap*’s ShapePlates significantly improves user experience.

VoxelHap proxies are robust and facilitate interactions in VR

3.1.6 Experiment 2—The Impact of Shape Fidelity

Informed by the results of experiment 1, we conducted a second experiment. Our central interest lies in whether adding shape resolution through ShapePlates improves the experience, especially how it influences realism, tactile feedback, and synchronization that may be affected by adding an additional 3D-printed layer to the proxy. By doing so, we could study if our three Voxeltypes and ShapePlates work effectively together. We kept the Task, Design, Experimental Protocol, Data Collection, and Analysis the same. For the Shaker, this meant that the proxy’s size exceeded the virtual model’s size by more than recommended [30]. We included this (1) to study trade-offs of ShapePlates and (2) to better understand the practical effects of *VoxelHap*’s geometric limitations. Finally, we also asked participants to customize the resistance of the Dial Lock to be as realistic as possible.

Effects of adding ShapePlates to VoxelHap proxies

3.1.6.1 Participants

We recruited a new set of twelve right-handed participants (five females, seven males), aged 19–32 (mean = 23.92; SD = 3.73) from the general public and the local university. Participants had a range of different educational and professional backgrounds, including psychology, computer science, law, cybersecurity, economics, biology, and




Figure 3.15: Experiment 2 results. *VoxelHap* proxies with ShapePlates against Vive controller. * = $p < .05$.




Figure 3.16: Experiment 1 and 2 descriptive comparison. *VoxelHap* proxies with vs. without ShapePlates.

Background of our participants

visual computing. All participants reported normal or corrected-to-normal vision and did not report any known health issues which might impair their perception. Four participants had never used VR before, six had used it a few times (1–5 times a year), one reported using it often (6–10 times a year), and another participant used it on a regular basis (more than 10 times a year). Participants not associated with our institution received €10 as remuneration for taking part in the experiment. The study was approved by the University’s Ethics Board.

3.1.7 Results

VoxelHap has clear advantages

Our results confirm the findings from experiment 1. *VoxelHap* proxies received greater average scores than the Vive controller. We found significant differences in realism for the Rattle and the Dial Lock but not for the Shaker (see Figure 3.15). Participants’ ratings on the most realistic device clearly favored *VoxelHap*. Interestingly, for the Shaker, the Vive controller came out as more enjoyable in the forced-choice question (see Figure 3.17). This is in line with participants’ comments often stating that the *VoxelHap* proxy “feels too bulky for a salt shaker” (P5), “is difficult to hold because of its size” (P10) or “is heavier than I would expect it to be” (P2). Participants were able to configure the resistance of the Dial Lock ($M_{\text{torque}} = 1.0 \text{ Nmm}$; $SD = 0.4 \text{ Nmm}$). Since all selected values lie well within the possible Nmm range, we conclude that the RotationVoxel’s technical capabilities to provide resistance is suitable for many rotational interactions, e.g., knobs.




Figure 3.17: Left: Participants' forced-choice votes regarding their preferred and most realistic device in both experiments. Right: Participants' configured strength of vibrotactile feedback in both experiments using *VoxelHap*.

3.1.7.1 Experiment 1 & Experiment 2

Here, we descriptively compare both experiments because a statistical comparison between them is not possible due to the iterative nature of our study design. When directly looking at participants' responses, it becomes evident that *VoxelHap* ratings were already quite high in experiment 1 (see Figure 3.16). Therefore, a distinct separation between ShapePlates vs. no ShapePlates was unlikely, especially given our sample size. Nevertheless, the questions regarding the most realistic and enjoyable device showed a clear increase in favor of ShapePlates, except for the Shaker. Here, adding ShapePlates led to a proxy that exceeded participants' expectations, both in terms of proxy size and weight. However, functionality, haptic feedback, and shape still seemed to be convincing to receive relatively high realism. ShapePlates are additional 3D-printed parts that can be attached to the Voxels. Yet it was unclear how this design choice might affect the perception of vibrotactile feedback. Therefore, we compared the configured strength (see Figure 3.17: right), synchronization, and vibration scores between the two experiments (see Figure 3.16), but we observed no clear difference that would suggest a decrease in haptic resolution, potentially caused by damped vibrations.

Comparing our results with and without ShapePlates

3.1.7.2 Summary

The results of experiment 2 confirm participants' positive responses towards *VoxelHap* proxies. Further, we found that the RotationVoxel's resistance feedback appeared to be in a reasonable range for hand interactions and that attaching ShapePlates does not seem to affect vibrotactile resolution. Finally, we gathered insights into how the different levels of shape fidelity provided by *VoxelHap* affect users' perception. *VoxelHap* proxies with ShapePlates were selected as more realistic, even though the Shaker exceeded a reasonable size and weight. This may suggest that shape contributes more strongly to the over-

all perception of realism than size [199] and weight. Nevertheless, it comes down to users' preferences and the demands of the task, which is in line with the core objective of *VoxelHap*, allowing users to approximate the (for them) crucial aspects, given the VR experience.

3.1.8 Discussion of *VoxelHap*

3.1.8.1 *VoxelHap* Types of Haptic Feedback

*Many capabilities,
but we could only
evaluate a subset*

As demonstrated above, *VoxelHap* proxies enable various types of haptic feedback such as simulating moving mass [54, 287, 335], stiffness [6, 20, 37], resistance [104] and size [30] variations or texture feedback [12, 69, 93]. In this work, we use application cases to evaluate the capabilities of *VoxelHap*, which, together with our technical and user evaluation, aligns with the proposed methods for toolkit evaluation according to Ledo et al. [212]. Yet, our user evaluation only included a subset of the demonstrated proxies. Thus, the created haptic effects need to be evaluated in future work. To this end, we only built three types of Voxels, but *VoxelHap* is not limited to this. We also imagine integrated hinges [218], bendable [91, 148, 239], twistable [391] or stretchable [91, 95] parts to unlock more functionality and combinations of haptic feedback. Moreover, Voxels could also be equipped with wheels [333], allowing them to re-position themselves, acting as an encounter-type device [240]. There is also a possibility to include rich electrotactile feedback as suggested by Groeger et al. [139]. The authors embedded conductive pads and wires directly into the 3D prints, providing tactile cues on various geometries. Since we use the same materials for fabrication, their technique could be adapted to further increase the tactile resolution of *VoxelHap*.

3.1.8.2 *VoxelHap* Beyond Virtual Reality

*VoxelHap also works
beyond VR interfaces*

VoxelHap was designed to build proxies that can be used in VR, but any other XR technology could benefit from it. For example, functional physical visualizations are still used in many domains, such as design or urban planning, because it is easier to understand spatial dimensions and facilitate discussions [156]. Users could collaboratively create functional drafts and the corresponding virtual model is generated automatically. In addition, *VoxelHap*'s rich input and output controls could be used for interactive music production or 3D animations using a paradigm of programming by demonstration [168]. We also envision *VoxelHap* proxies to be effective in remote collaboration, acting as a shared physical artifact. For example, the local novice could feel the manipulation of the proxy and could be guided through a task by a remote expert [96].

3.1.8.3 *VoxelHap's Limitations*

Given the technical capabilities of our Voxels, they do have a competitive size [91, 294, 357]. Nevertheless, they would benefit from an even more compact design. Both BaseVoxel and VibrationVoxel leave room for optimization; however, the RotationVoxel mainly determined the final dimensions. This is a result of the chosen motor because it needed to be large enough to produce human-noticeable resistance. We tested several different motor versions and opted for the smallest possible motor. This remains a problem of any active haptic device or component—they are often bulky, because of the forces they need to produce. Our construction pipeline only supports Voxels and passive ShapePlates. We decided against active ShapePlates, since it would have required us to equip ShapePlates with hardware, slowing down the fabrication and construction process. Depending on the use case, e.g., designers who re-use their ShapePlates in 3D modeling applications [183, 216], might benefit from active components. In addition, the construction process relies on the manually generated .hpdf which we aim to automate in the future. To this end, we did not evaluate user performance in constructing *VoxelHap* proxies but leave this for future work. Each Voxel is a self-contained unit and, therefore, comes with its own power source. This limits the scalability and usability to a certain extent. Voxels cannot share or distribute power and need to be recharged individually. One of our earlier prototypes used wireless charging, but this only works when the proxy is frequently placed and left on the table, for example, encounter-type or grounded proxies. Another interesting approach would be to utilize users' mechanical manipulations of functional parts, i.e., rotation, to harvest energy [340].

Current proxy size limitations

Alternative design choices

3.1.9 *Conclusion & Contributions*

In this section, we contribute to **RQ1** by developing *VoxelHap*, a block-based construction toolkit, allowing users to build functional handheld size VR proxies that provide multimodal tactile and kinesthetic haptic feedback. *VoxelHap* proxies can change their haptic properties through manual reconfiguration by using Voxels, blocks with special haptic functionalities, and Plates, which can be attached to increase the haptic resolution of proxies. To illustrate the potential of the toolkit, we presented a range of fully functional proxies for various use cases and applications (**C5**). In two experiments, we evaluated a subset of the constructed proxies and studied how they compare to a traditional VR controller. First, we investigated *VoxelHap's* combined haptic feedback, showing that it can provide a variety of tactile and kinesthetic haptic sensations while allowing dynamic touch and exploratory interactions (**C1**). To this day, most controllers are limited

Provide multimodal tactile and kinesthetic haptic feedback

Presented use cases and applications

*Two experiments
show its
effectiveness*

*Rapid fabrication
techniques for proxy
design*

*Insights that point
to the uncanny
valley of haptics*

*Open-source hard-
and software*

to rendering one type of haptic feedback and lack full embodiment, which is crucial to enable natural and intuitive interaction [91]. Together, with ongoing research in the field [16], we highlight the need for more holistic approaches to haptic feedback and not solely focus on the rendering quality of a single haptic modality. Second, we investigated the trade-offs of using additional shape resolution by applying *VoxelHap*'s ShapePlates. Here, our findings show that proxies with higher shape fidelity outperform traditional controllers and were generally favored by participants. However, we also provide further evidence for the uncanny valley of haptics theory [29]. *VoxelHap* proxies achieve greater shape and size similarity than a generic controller. As a result, participants seemed to expect all of the proxies' properties to match the virtual object, whereas the generic VR controller did not trigger such expectations (C3). Here, participants mostly preferred the great ergonomics over proxies. This could be a result of our simple study tasks, where shape and size approximations and whole-hand interactions were not a crucial part of the interaction. Thus, we contribute towards a better understanding of proxy design, suggesting that VR designers need to carefully balance task, interaction, and proxy. As presence research suggests, realism seems to be directly bound to users' expectations and is not necessarily correlated with real-world similarity. Finally, by open-sourcing *VoxelHap*'s hardware and software, we hope to ease access, save resources and encourage the community to contribute to haptic interfaces—even beyond VR (C6). *VoxelHap* is the first VR toolkit that incorporates existing hand-based illusion techniques into the design process, but without a formal investigation, it falls short of its potential. To this end, *VoxelHap* suffers from two major limitations, it requires manual reconfiguration and it lacks haptic resolution in the light of uncanny valley. Therefore, we use the next section to investigate an approach that supports dynamic reconfiguration (i.e., shape-change) of the proxy.

3.2 ENHANCING TACTILE PERCEPTION THROUGH VISUO-HAPTICS




Figure 3.18: Metamaterial with visual material textures overlays. Six visual material textures: concrete, wood, plastic, fabric, glass, and metal. In addition, an example of how a user interacts with a 3D-printed metamaterial.

3.2.1 Introduction

Touch sensations are an essential part of understanding the physical world around us [283], and remain crucial in creating realistic and plausible virtual experiences wherein Lisa can feel present and act accordingly [310]. To successfully render haptic properties, proxies need to provide physical sensations with a sufficiently similar feeling in terms of its material, e.g., texture, and geometrical properties, e.g., shape [256]. *VoxelHap* achieved this, but resulted in scaling issues, as the number of manual reconfigurations required to represent large amounts of virtual objects, each with varying haptic properties, rapidly increased. Therefore, we utilize digital fabrication to create more scalable and flexible proxy objects. Recent advancements in fabrication technologies support the manufacturing of highly detailed physical artifacts with varying haptic properties. Examples include methods for designing objects with desired mechanical behavior, such as elasticity or deformation through varying internal microstructures [34, 297], or for fabricating perceptually-varying surface texture qualities [67, 117, 271]. When combining such abstract structures with visual texture overlays in IVEs, they are able to support the visuo-haptic perception of different virtual materials. However, still a large enough set of objects needs to be produced [69]. This section aims to support the design and fabrication of novel proxy objects that can dynamically change (i.e., DPHF device) their tactile properties upon actuation and combine this with visual overlays. We take inspiration from the most recent work in the fabrication space on 3D-printed metamaterials [159], studying their potential as proxies to enhance tactile experiences in VR.

Dynamic reconfiguration helps overcome scaling issues

Combining digital fabrication with visual overlays

3.2.2 *Metamaterials for Tactile Texture Perception*

Here, we outline the design, implementation and fabrication of our actuated 3D-printed metamaterial prototypes. We provide valuable insights and lessons learned to ensure that our approach is reproducible. Our interest lies in changing roughness and hardness properties of a metamaterial by actuating (compressing) it, allowing it to act as a proxy for several virtual material textures.

3.2.2.1 *Design*

*We present five
metamaterial
patterns*

In total, we designed five different metamaterial patterns that can be actuated. To influence their tactile perception, we focus on the properties of hardness and roughness, which have been found to be the main contributors of texture and material perception [263]. Below, we discuss the design choices for our two perceptual dimensions of interest, hardness and roughness.

*Designing for
hardness*

HARDNESS To change the hardness of the metamaterial, we utilize the concept of porous materials—solids that contain (penetrating) pores. A measurement for how porous a material is, is called porosity, which is defined by the fraction of pore volume to the total volume of a material. If a material has a porosity level of 0.2–0.95, it counts as porous, according to the definition of Ishizaki et al. [161]. This criterion is fulfilled for our designed metamaterial patterns. Lu et al. [226] investigated strength, the elastic modulus and the hardness of porous materials. They found that all of these material properties substantially decrease when porosity increases. We build on their results by using porosity as our key design variable to change the hardness within the same metamaterial. To summarize, when we decrease the porosity within 3D-printed metamaterials, we increase their hardness. In Figure 3.19, we illustrate a compressed cell of our metamaterial 3D prints, leading to a reduced fraction of the pores. Thus, the porosity decreases, and therefore, we would expect an increase in the perceived hardness of the metamaterial. Please note that compressing the metamaterial will ultimately lead to emerging features such as bumps (see Figure 3.20). In fact, this adds a second ‘haptic’ layer, potentially affecting hardness sensations. To investigate this potential issue, we ran a preliminary experiment.

*Designing for
roughness*

ROUGHNESS To achieve changing roughness sensations within one metamaterial, we based our designs on two core principles. First, the emerging features from our metamaterials should be in the range of 2 mm to 3.5 mm, because Klatzky et al. [185] found that the perceived roughness increases consistently with inter-element spacing up to approximately 3.5 mm. Second, the features should still be perceived as one surface. Therefore, we incorporate findings from two-point




Figure 3.19: Metamaterial deformation simulation. (A) Fixating metamaterial arches creates a uniform periodic deformation. (B) By not fixating uneven arches, the metamaterial deforms upwards.

discrimination experiments. Dellon [70] describe this as the distance between two perceivable points that can be distinguished from one another. Depending on age, this distance usually varies between two and three mm. Hence, we use a feature distance of a maximum of 3 mm, ensuring that our metamaterials are perceived as one surface.

METAMATERIAL PATTERNS Finally, we present our resulting metamaterial patterns based on the previously outlined design principles and many iterations. MAT₁ and MAT₂ are closest to the proposed feature spacing range of 2 mm to 3.5 mm when they are compressed (see Figure 3.20). They contain the smallest bumps and spikes emerging that we were able to produce with our fabrication technique described below. Furthermore, their patterns aim to achieve the smallest feature size achievable with the way our metamaterials are structured—without the features being spaced too far apart. The main difference between MAT₁ and MAT₂ is the design of the spike shape emerging. MAT₂ has more rounded edges. We include this slight variation, as it reassembles a worn-off version of MAT₁, because Myers [250] found that worn-off materials can result in different roughness sensations. Adding a worn-off state of MAT₁ enables us to collect insights, helping us to better understand longevity and potential perceptual drifts caused by material deterioration. We included MAT₃ design as a pattern that does not use ‘spiky’ features. This comes with the sacrifice of the metamaterial features becoming larger—exceeding the 3.5 mm discrimination threshold. Hence, this design does not meet the criteria to achieve the feeling of one surface, as the features become clearly distinguishable. Yet, exploring a non-spiky pattern was an interesting design variation. MAT₄ is the spiky counterpart to MAT₃, allowing us to investigate how much ‘spikiness’ plays a role with larger surface features. Finally, metamaterial MAT₅ was specifically designed for simulating hardness. This material has the greatest change in porosity among all metamaterials presented. With this, we aim to study whether the change in hardness found by Lu et al. [226] also holds true for our metamaterial patterns. When looking at the structure of this material, the emerging features appear to be straight bumps and, thus, will most likely not be associated with roughness.

*Designing
metamaterial
patterns for fingertip
exploration*

*MAT₃ uses
non-spiky features*




Figure 3.20: Metamaterial designs MAT1 to MAT5, each in three different actuation states 0 mm, 2 mm, and 6 mm. 4 mm actuation state is missing.

3.2.2.2 Actuating 3D-printed Metamaterials

In this section, we describe our actuation approach. Our ultimate goal was to implement a prototype that allows for the linear compression of metamaterials, resulting in various different states. As a result, we can control surface changes as shown before.

*Design requirements
to achieve the
desired deformation*

INITIAL PROTOTYPING & DESIGN REQUIREMENTS Ion et al. [159] demonstrate how 3D-printed metamaterials behave when being compressed. An example of a non-actuated as well as an actuated metamaterial can be seen in Figure 3.21. In order to actuate our metamaterials, resulting in the desired equal deformation, every second pillar inside the metamaterials must freely move vertically upwards upon compression. For instance, in Figure 3.19 these are the uneven pillars 1, 3, 5 and 7; even pillars, i.e., 0, 2, 4, 6, and 8 are prevented from going upwards. The triangular shapes in the segments between the pillars encourage the behavior of moving only every second pillar upwards. However, one key limitation is the flexibility of the metamaterials, which causes the mechanism to not work properly when metamaterial structures become larger. For instance, for two rows of cells, triangular shapes work properly, but once we add a third or more rows to the metamaterials, the deformation behavior changes. Figure 3.19 (B) illustrates the effect when compressing a larger metamaterial, only relying on the deformation mechanism 3D-printed directly into them. The metamaterial starts not deforming row by row but instead deforms in total and escapes as an arch unevenly upwards. Based on our initial experiments and several prototyping iterations, we formulated the following requirements for the actuation mechanism:

R1: Our actuation approach should force the metamaterials to move vertically upwards row by row instead of the whole metamaterial arching as one unit.

R2: Our actuation approach should only allow every second pillar within the metamaterials to freely move vertically upwards and downwards.

*Requirements for
our metamaterials
actuation device*

R3: Our actuation approach should result in a uniform distribution of emerging features throughout the entire actuation spectrum.

ACTUATION APPROACH Based on the requirements stated above, we developed a simple actuation approach inspired by Neville et al. [252]. The authors fed strings through their metamaterials and used a pulling force on the strings to achieve the deformation. Applying this to our designs comes with the need to further modify the 3D prints with holes running through the materials and cut-outs in the inner metamaterial pillars. We started testing with one hole through each metamaterial cell, which we directly 3D-printed into them. Initially, we used standard fishing line fed through the metamaterials, allowing us to compress the material by pulling on the strings. However, the metamaterials deform unevenly, as illustrated in Figure 3.19 (B). This is caused by the strings because they do not restrict the metamaterial to a single plane of motion. Therefore, we replaced the fishing line with small metal pipes that can be seen in Figure 3.21. The stiff metal pipes restrain the metamaterials from staying in one plane—solving the problem of non-uniform deformation (**R1**). To achieve the desired deformation state, we switched to a pushing/pulling approach and added two fixation blocks, where either end of the metal pipes can be inserted. The cranks are moved by a Nema 17 stepper motor and controlled using an Arduino Uno and a DRV8825 driver. The setup can be seen in Figure 3.21. As defined above, every second pillar within the metamaterial must be allowed to move vertically upwards and downwards (**R2**). With only the holes for the pipes inside the materials, the pillars inside are now restricted from moving because of the metal pipes. As a result, we changed the holes in every second pillar into full openings. For the material next to the holes to not rub at the stiff metal pipes, we also cut out the bottom layer of the metamaterial cells, which can be seen in Figure 3.21 (right).

*Compression to
achieve the
deformation*

*Adding small metal
pipes*

In the last step, we specifically looked at (**R3**) because we observed differences in how the features emerged due to the way we achieved the material compression, i.e., the metamaterial cell row nearest to the crank will compress more than the second one and so on. This was unproblematic for achieving binary actuation states, compressed and uncompressed; however, it does not support consistent ‘infinite’ states in between. Hence, we attempted to create a more uniform force distribution while actuating. To do so, we first embedded small metal springs around the metal pipes. If metal springs are sufficiently

*Altering the design
to allow equal
emerges of features*




Figure 3.21: Metamaterial actuation setup. Left: A stepper motor and two cranks are used to compress the metamaterial. Center: Detailed view of a compressed metamaterial showing periodic deformation. Right: Embedded tubes guide metal support pipes to ensure equal distribution of applied force, creating a uniform compression across the surface.

stiff, they will transfer the applied pressure to the neighboring cell before they fully compress. Nevertheless, we were unable to find metal springs that had the correct stiffness given our form factor; most of them were not stiff enough or too large. Instead of springs, we then 3D-printed small tubes out of flexible filament that we fitted around the metal pipes shown in Figure 3.21. They offer suitable stiffness, which can be adjusted by changing their wall thickness. For our 3×3 metamaterial designs, we used tubes with a wall thickness of 2 mm and a length of cell space minus one mm. This allowed us to reach four visibility distinct actuation states, with 2 mm steps of compression added for each new state. Currently, our metamaterial patterns have to be manually post-processed in CAD software after they have been procedurally generated. In addition, they require additional hardware embedding, e.g., metal pipes to serve as support structures. More advanced multimaterial 3D printing may reduce the need for such additional support, as this would be able to embed internal structures directly into the design using rigid filaments such as PLA or carbon compounds. Additionally, the currently embedded metal pipes influence hardness. To address this, we experimented with different designs and found that up to a metamaterial size of 3×3 or 4×4 cells at a cell size of 15 mm, designers can rely on only two metal pipes in the outermost cell rows. This creates an area in the center of the metamaterials that can be explored without the influence of the support structures.

3.2.2.3 Fabrication Process

Our objective was to create a 40 mm \times 40 mm intractable surface area that can be explored with a user's index finger [69]. Similar to Ion et al. [159], we aimed to fabricate our metamaterials with a low-cost off-the-shelf FDM printer. According to the authors, they used filament with a shore hardness rating of 85A. The lower the shore hardness rating (A-value), the softer the material. However, since our use case differs from their application, we 3D-printed metamaterial

Four compression states can be reliably achieved

Recommendations for alternative fabrication approaches

Experimenting with hardness ratings of 3D printing filament

samples with multiple filaments: Prusa PLA (95A), Recreus FilaFlex (90A), Recreus FilaFlex (82A) and Recreus FilaFlex (60A). We used the most basic cell shape available in the editor and set the cell size to 15 mm. All samples were 3D-printed on a Prusa i3 MK3S using the recommended settings from the corresponding filament providers. We investigated two properties: (1) how easily the metamaterial can be deformed and (2) how well the prints return to their original shape after deformation. Both, the 82A and the 60A shore hardness filaments showed promising results, satisfying the two criteria. In next step, we 3D-printed varying cell sizes: 15 mm, 12 mm and 9 mm. The latter being the smallest cell size we could reliably print with the Prusa i3 MK3s. When testing the deformation of these prints, we found that the smaller the movable features of the metamaterial, the softer the filament needs to be. This appears to be an interesting trade-off between cell size and filament hardness and an important aspect when fabricating metamaterial structures. While experimenting with the printing settings, we observed that many variables, such as slight variations in nozzle temperature or tightness of the printer screw, significantly affect the hardness and quality of the resulting metamaterial prints. Therefore, we provide our final 3D printing settings below. This should be a solid starting point for anyone who wants to fabricate and experiment with their own 3D-printed metamaterials.

Setup for technical exploration

Fabrication process and the resulting 3D prints are sensitive to many variables

- Basic settings: the basic profile and settings should be taken from the manufacturer's website.
- Print quality: 0.15 mm.
- Infill: 10%. May be increased for larger metamaterial prints.
- Printing speed settings: we set printing speed settings to 20 mm/s.
- Nozzle temperature: 212 °C.
- Printer screw position: unscrew 2 – 3 rotations.
- Extruder: set to zero in the slicer software.

Finally, our final prototypes use a 3×3 grid with 15 mm cell sizes. This results in a 4.5 cm \times 4.5 cm surface in the uncompressed state. The average minimal actuation of the metamaterials for the first features to emerge lies around 2 mm. Further actuation states are visibly distinguishable when actuating linearly in 2 mm steps. With each metamaterial design (MAT1–MAT5), we achieved four distinct actuation states, where the maximum compression will be reached at 6 mm. As a result, even maximal compression still results in a large enough surface of 4.5 cm \times 3.9 cm to be explored.

Final setup for our user study

3.2.3 Experiment 1: Tactile Metamaterial Perception

To understand the tactile perception of our metamaterials, we conducted an initial user study.

3.2.3.1 Design

The goal of this preliminary experiment was to gather first insights into how our actuated 3D-printed metamaterial patterns are haptically perceived in terms of roughness and hardness. Additionally, we wanted to understand whether certain metamaterial patterns are more suitable for simulating roughness and hardness. To investigate this, we used a within-subjects design consisting of a baseline and similarity assessment phase.

BASILINE ASSESSMENT In the baseline condition, we presented each of the 20 material states once to the participant. They were asked to explore them with the index finger of their dominant hand. Then, they rated each material in terms of roughness and hardness on a 9-point Likert scale. The order in which the metamaterials were presented to the participants was counterbalanced using Latin square. We asked the following questions:

Q1.1: How hard would you rate the inspected material? (1 = extremely soft; 9 = extremely hard)

Q1.2: How rough would you rate the inspected material? (1 = extremely smooth; 9 = extremely rough)

SIMILARITY ASSESSMENT To directly compare the different 3D-printed metamaterial patterns and their actuation states to each other, we included a second condition. This resulted in $\frac{20 \times 19}{2} = 190$ material combinations participants had to assess. Here, we presented two different metamaterial patterns simultaneously to participants. We asked them to explore both samples and consequently rate them on a 9-point Likert scale with respect to their similarity in terms of roughness and hardness:

Q2.1: How similar does this material feel in terms of hardness? (1 = extremely different; 9 = extremely similar)

Q2.2: How similar does this material feel in terms of roughness? (1 = extremely different; 9 = extremely similar)

*Study questions for
baseline assessment*

*Study questions
similarity
assessment*

3.2.3.2 Hypothesis

We formulated the following two hypotheses:

H1.1: Increasing the compression of our 3D-printed metamaterial patterns leads to an increase in their perceived roughness.

Hypothesis for our preliminary experiment

H1.2: Increasing the compression of our 3D-printed metamaterial patterns leads to an increase in their perceived hardness.

3.2.3.3 Participants

We recruited six participants (two females and four males), aged between 23 and 59 (mean = 25.5; SD = 15.2) with backgrounds in law, engineering, accounting, computer sciences and the medical domain. Regarding hand dominance, five participants indicated being right-handed, while one participant noted they were ambidextrous with a preference for left-handed interaction. Participants stated they did not have any impairments that could influence both their visual or tactile perception. All participants indicated on a 5-point Likert scale they never worked with textiles (mean = 2.0; SD = 1.0), and noted they did not often perform precise handwork (mean = 1.1; SD = 0.3). Participants received candy for taking part in the experiment. The study was approved by the Saarland University's Ethical Review board.

3.2.3.4 Apparatus

The experiment took place in a quiet room. Our set of metamaterials consisted of a total of 20 surfaces, i.e., five specific metamaterial patterns, with each having four fixed compression states, see Figure 3.20. To avoid visual bias, participants were not allowed to see the surfaces during exploration. To this aim, they were seated in front of a table with a cardboard screen. The opening at the bottom of this screen ensured enough room for the hand of the participants to reach through. Behind the screen, the metamaterial patterns were manually placed by the experimenter in a fixed location in a laser-cut hold out of wood. A digital rendering of the setup is depicted in Figure 3.22.

Metamaterial exploration study setup

3.2.3.5 Experimental Protocol

Before starting the experiment, participants provided signed consent and were briefed regarding the upcoming course of events. A short introduction illustrated the setup and explained the experiment's purpose. Participants were asked to complete a questionnaire to collect their demographic information. During the experiment, participants were instructed to explore the physical samples using only the index finger of their dominant hand. For each task, we demonstrated how to

Participants ran through the following procedure




Figure 3.22: Digital rendering of the experimental setup. The experimenter prepared the samples for exploration, while the participant rated the tactile perception of the samples hidden behind the screen.

explore the materials to ensure consistency between samples. Specifically, participants were allowed to use circular lateral exploration to assess each structure's surface information and could lightly press or tap each structure to assess hardness properties. Participants first completed the baseline task before continuing with the comparison task. The observer noted the responses for each trial and activated the next sample. Short breaks were scheduled between tasks and when the participant noted a feeling of numbness in their finger. Per participant, the total experiment took about 60 min.

3.2.3.6 *Data collection & Analysis*

We collected data from four sources, i.e., demographic information for the pre-study questionnaire, ratings for hardness, roughness and similarity for each trial on a 9-point Likert scale. Statistical tests were chosen based on whether the data satisfied parametric test assumptions at $\alpha = .05$ using Shapiro–Wilk test and QQ plots. Reported p-values are corrected using Greenhouse–Geisser when the sphericity assumption was violated. For outlier removal, we initially applied the box plot method. Post-hoc pairwise comparisons were corrected using Bonferroni–Holm adjustments.

3.2.4 *Results*

In this section, we present the results of our preliminary experiment. The analysis is split into two parts corresponding to our study conditions. First, we report the roughness and hardness assessments of the metamaterials by our participants.

3.2.4.1 *Analysis of Baselines*

Here, we consider average ratings for roughness and hardness for each metamaterial and its compression states. We only ran our statis-




Figure 3.23: Participants separate assessments for haptic roughness and hardness for the different metamaterial actuation states 0, 2, 4, and 6 mm (left to right).

tical analysis on each metamaterial, as we were primarily interested in the change of roughness and hardness within the same material.

ROUGHNESS All metamaterials MAT1 ($F_{(3,15)} = 69.0$, $p < .0001$, $\eta^2 = .828$), MAT2 ($F_{(3,15)} = 97.8$, $p < .0001$, $\eta^2 = .828$), MAT3 ($F_{(3,15)} = 31.7$, $p < .0001$, $\eta^2 = .725$) and MAT4 ($F_{(3,15)} = 36.8$, $p < .0001$, $\eta^2 = .763$), except MAT5 ($F_{(3,15)} = 1.80$, $p < .0001$, $\eta^2 = .121$) showed a main effect when analysing their roughness scores. As stated above, it is not surprising that MAT5's actuation states did not show a main effect because its features are well above the 3 mm threshold [70]. We ran post-hoc analysis on MAT1, MAT2, MAT3 and MAT4 using Bonferroni-Holm adjustments. The results are depicted in Figure 3.23 (top). All significant differences are in line with our **H1.1**, i.e., showing an increase in perceived roughness with increased compression of the metamaterial.

Compressing metamaterials significantly increases the perceived roughness

HARDNESS MAT2 ($F_{(3,15)} = 20.1$, $p < .0001$, $\eta^2 = .760$), MAT3 ($F_{(3,15)} = 16.9$, $p < .0001$, $\eta^2 = .578$), MAT4 ($F_{(3,15)} = 6.82$, $p = .004$, $\eta^2 = .491$), and to our surprise even MAT5 ($F_{(3,15)} = 14.2$, $p = .002$, $\eta^2 = .515$) showed a main effect, when analyzing their hardness scores. Only MAT1 ($F_{(3,15)} = 2.0$, $p = .157$, $\eta^2 = .233$) did not reveal a main effect. We ran post-hoc analysis using Bonferroni-Holm adjustments. The results are depicted in Figure 3.23 (bottom). For hardness, our results are mixed. Even though we found significant differences between the actuation states, they do not seem to follow a

For hardness perception results are mixed

clear trend, an exception being MAT₄ which follows our **H1.2**. For all other metamaterials, our preliminary study fails to provide evidence.

3.2.4.2 Analysis of Similarities

*How similar are the
selected
metamaterials to
each other?*

Here, we consider the similarity ratings obtained through paired comparisons of our metamaterials. We combined the similarity ratings of roughness and hardness for each metamaterial pair by averaging them. To determine the consistency of the obtained data, we used Spearman's rank correlation tests on the averaged similarity assessments. Here, we found the averaged similarity ratings for each participant to be correlated with those of every other participant ($M_r = 0.52$, $p < .001$). Given this result, we note that averaged similarity assessments remain consistent across participants. For further analysis, the similarity assessments (1–9) were converted to normalized dissimilarity ratings (0–1). With these ratings, we created a symmetric dissimilarity matrix containing the perceptual distances between all metamaterials and compression states. Using an analysis of similarities (ANOSIM), we compared different groups within our distance matrix. We found a significant difference when comparing groups of different compression levels ($R = 0.476$, $p < .001$). However, we did not find a significant difference when comparing groups of different metamaterials ($R = 0.156$, $p = .076$). We conclude that compression levels have the greatest impact on users' perception of the material.

3.2.4.3 Summary & Metamaterial Selection

*We selected MAT₁
and MAT₄ as the
most promising
candidates*

The main goal of the experiment was to collect early feedback on the prototypes and select the most promising metamaterials for our main study. Our preliminary findings suggest that MAT₁, MAT₂, MAT₃, and MAT₄ appear to be promising candidates for increasing the perceived roughness upon actuation. In contrast, for hardness, only MAT₄ followed our hypothesis, showing an upward trend in hardness upon compression. Consequently, we decided to include MAT₄. In addition, we also used MAT₁ in the main experiment as it created the most distinct levels of roughness according to participants' assessments.

3.2.5 Experiment 2: Visuo-Haptic Perception

Using our selected metamaterials MAT₁ and MAT₄, we conducted a user study in VR to investigate how our designs were able to simulate visuo-haptic materials.

3.2.5.1 Design

The goal was to study how adding visual material textures atop our 3D-printed metamaterials affects their perception. Therefore, we combined our metamaterial patterns with virtual material textures in VR and asked users to rate how well they matched. Moreover, we were interested in the range of possible material textures that a single actuated metamaterial may simulate. We again used a within-subjects study design. We focused on the two most promising metamaterial patterns from our preliminary experiment, MAT₁ and MAT₄, at the compression rates of 0, 2, 4, and 6 mm. This makes a total of eight metamaterial patterns that participants ought to explore. Due to our concerns regarding the longevity of the metamaterials, we again fixated the different compression states by screwing our actuation mechanism with the metamaterials onto small wooden plates (see Figure 3.20). The virtual material textures were chosen based on previous research by Degraen et al. [69]. We added two additional virtual materials, resulting in the following selection: metal, wood, glass, plastic, concrete, and fabric, depicted in Figure 3.18. We believe that our selection covers a wide range of everyday materials and, therefore, provides a solid starting point. Below, we describe the three conditions in the order participants completed them.

*Adding visual
texture overlays*

*Preparing the
samples*

HAPTIC BASELINE The participants were asked to explore all eight metamaterial patterns once and rate them regarding their perceptual dimensions, roughness, and hardness on a 9-point Likert scale. Additionally, we asked: “Which material do you think this is?” to better understand whether participants associate the 3D-printed metamaterials with known textures or objects from their everyday environments. To ensure that the participants only relied on their haptic senses, the VR headset was turned off during this condition, only showing a dark screen. The order in which the participants explored the eight metamaterial patterns was counterbalanced using a Latin square design.

*Our two baseline
conditions*

VISUAL BASELINE Next, participants rated all visual material textures regarding their roughness and hardness, i.e., they rated how they would expect the materials that they saw to feel on a 9-point Likert scale. We also used the question “Which material do you think this is?”, to make sure all materials were properly recognizable by the participants. The order in which the virtual material textures were presented was again counterbalanced using a Latin square.

VISUO-HAPTIC COMBINATION Finally, we looked at the combination of visual material textures and haptic exploration of the metamaterial patterns. Each metamaterial pattern was paired with all six virtual material textures, resulting in 48 pairs. These pairs were presented to the participants in a randomized order. Repeatedly,

one of the material spots in the VR environment was highlighted, indicating which material textures should be explored next. The participants then proceeded to explore the highlighted material both visually and haptically at the same time. Consequently, they gave a separate rating for roughness, hardness, and similarity (between the visual and haptic material texture) on a 9-point Likert scale and responded to the “Which material do you think this is?” question.

*Questions assessing
visuo-haptic
combination*

Q1.1: How hard would you rate the inspected material? (1 = extremely soft; 9 = extremely hard)

Q1.2: How rough would you rate the inspected material? (1 = extremely smooth; 9 = extremely rough)

Q1.3: How does the visual and the haptic perception match? (1 = not at all; 9 = extremely similar)

3.2.5.2 Hypothesis

Besides confirming our preliminary findings from experiment one, we formulated two additional hypotheses for our main experiment:

*Hypothesis for our
main experiment*

H2.1: Combining our 3D-printed metamaterials with visual material overlays leads to a shift in perception.

H2.2: Our 3D-printed metamaterial patterns provide matching haptic sensations for visual material textures.

3.2.5.3 Participants

*Participant pool for
our main experiment*

We recruited 16 participants (six females and ten males), aged between 22 and 59 (mean = 25.5; SD = 11.1) with various backgrounds, including computer sciences, education, economics, biology, law, electrical and mechanical engineering. Twelve participants stated that they were right-handed, and three were left-handed. All participants confirmed that they have no impairment in their index finger which may affect their perception. Additionally, we asked how often participants perform precise handwork, e.g., sewing or stitching, on a 5-point Likert scale. Four participants rated themselves with a four or five; all others were three or below (mean = 2; SD = 1.3). Participants' answers about how often they work with textiles, e.g., clothing, design, or tailoring, on a 5-point Likert scale showed that only two answered with a higher value than four. All others are around one and two (mean = 1; SD = 1.1). Out of the 16 participants, only three had prior experience with VR. Therefore, we highlighted very carefully to all participants that they can pause or stop the experiment at any time, e.g., in case they feel uncomfortable. Participants received

unlimited candy as compensation for taking part in the experiment. The study was approved by the Saarland University's Ethics and Hygiene Board.

3.2.5.4 Apparatus

Our apparatus consisted of an HTC VIVE Pro tracking system with SteamVR (v.1.22) and OpenVR SDK (v.1.16.8). The simple virtual scene was developed in Unity3D (v.2021.1.0). We used a Raubtier NBB01471 offering an Intel® Core i7 9700kf CPU, 32 GB RAM and an Nvidia® GeForce RTX 2080Ti for running the experiment. We included an androgynous hand representation without noticeable characteristics. To enable hand tracking, an HTC Vive tracker was attached to the back of the participant's hand. We used the calibration method from Zenner and Krüger [382], asking participants to touch the touch-sensitive trackpad of an HTC Vive controller with the tip of their index finger. This was used to properly align the virtual hand model with the real hand. We rendered either a right or a left virtual hand, corresponding to a participant's handedness. To present the eight different metamaterial patterns to the user, we laser-cut a mount where the prepared metamaterial plates fit in. The mount was placed on a table in front of the user, allowing participants to explore the metamaterials while, at the same time, enabling the experimenter to quickly rearrange the metamaterial patterns throughout the study.

Hardware setup in the experiment

3.2.5.5 Experimental Protocol

First, participants received a general introduction to the study. Next, we gathered their consent and asked them to fill in a demographics questionnaire. We explained the task, including a demonstration of how to explore the metamaterial patterns with the index finger of their dominant hand. Then, we attached the Vive tracker to the participant's hand, they entered the virtual environment, and we performed the hand calibration. Participants were guided through a practice round, showing them how to respond to the questions displayed in VR. Finally, they went through our three study conditions described above. Participants were prevented from seeing the metamaterial patterns before and during the study to avoid biases. Overall, the experiment took about 70–80 min.

Demographics hand calibration and practice round

3.2.5.6 Data collection & Analysis

We collected data from six sources: a pre-study questionnaire for demographic information; hardness, roughness and similarity ratings on a 9-point Likert scale; participants material assessments and a semi-structured interview to better understand participants' experiences with our prototypes. Questionnaire responses were collected inside the virtual environment using our *VRQuestionnaireToolkit* [94].

Collected data from six sources




Figure 3.24: Haptic baseline. Participants separate assessments for roughness (top) and hardness (bottom) for the different metamaterial actuation states.

Interview responses and participants' comments were coded for analysis. Analog to experiment 1, we chose statistical tests based on parametric test assumptions at $\alpha = .05$, and we used outlier removal with the box plot method. Post-hoc pairwise comparisons were corrected using Bonferroni-Holm adjustments.

3.2.6 Results

3.2.6.1 Baseline Results

In this section, we present the results from our main experiment. The analysis is split into two parts corresponding to our study conditions. We start by reporting our results from the haptic and visual baseline conditions. Finally, we present the results from the visuo-haptic combination.

Haptic baseline largely confirms preliminary results

HAPTIC BASELINE Participants rated the different actuation states of MAT1 and MAT4 significantly different from each other, in terms of roughness MAT1 ($F_{(3,15)} = 2.025$, $p = .023$, $\eta^2 = .578$) MAT4 ($F_{(3,15)} = 2.025$, $p = .023$, $\eta^2 = .164$) and hardness MAT1 ($F_{(1.92,28.7)} = 72.1$, $p < .0001$, $\eta^2 = .774$) MAT4 ($F_{(1.9,28.6)} = 182.0$, $p < .0001$, $\eta^2 = .853$), across various metamaterial actuation states (see Figure 3.24). This confirms the results from our preliminary experiment and thus provides evidence for **H1.1** and **H1.2**. One exception is that MAT1 in its maximum actuation state (6 mm),




Figure 3.25: Visual baseline. Participants separate assessments for roughness (left) and hardness (right) for the different visual material textures.

was rated significantly softer than its predecessor 4 mm, which contradicts with our hypothesis that increasing actuation leads to a perceived increase in hardness. This might be caused by the fully emerging features, adding up to 2 mm in height. This feature layer has fewer support structure; thus, causing the material to deform until the ‘base’ layer by acting like a cushion. Considering that all other actuation states worked as expected, we believe that there is an interesting trade-off between feature design and compression rate that requires further exploration.

Participants’ responses when being asked to identify the material without visual information. 64% of the collected answers were rubber or plastic variations for MAT1. Interestingly, we mostly observed other answers such as fish scales, tar, grass, tree branching or fine stone, in the actuation states and not in the non-actuated state. This suggests that there exist materials that are closely reassembled by the actuated metamaterials. For MAT4 however, substantially fewer alternatives to rubber and plastic (only 15%) were given. For example, tree bark, sandstone or metal. While the non-actuated metamaterials, MAT1 and MAT4, were mostly associated with rubber or plastic, the higher compression rates are potentially capable of providing more complex material sensations. Moreover, participants often used adjectives such as “rough” and “damaged” to describe metamaterials with greater actuated states.

Shifting material perception without a visual stimulus remains challenging

VISUAL BASELINE The results confirm our selection of visual material textures vary in expected roughness ($F_{(6,30)} = 4.025$, $p = .003$, $\eta^2 = .653$) and hardness ($F_{(6,30)} = 2.025$, $p = .023$, $\eta^2 = .853$) (see Figure 3.25). For example, glass and concrete have significantly ($p < .0001$) different roughness ratings, comparable to fabric and metal for hardness ($p < .0001$). Participants were able to identify all textures without prior information, purely upon visual inspection. We only observed three mismatches, metal and concrete were perceived as marble twice, and glass was perceived as ice once. Please

Visual material renderings in VR could be easily identified

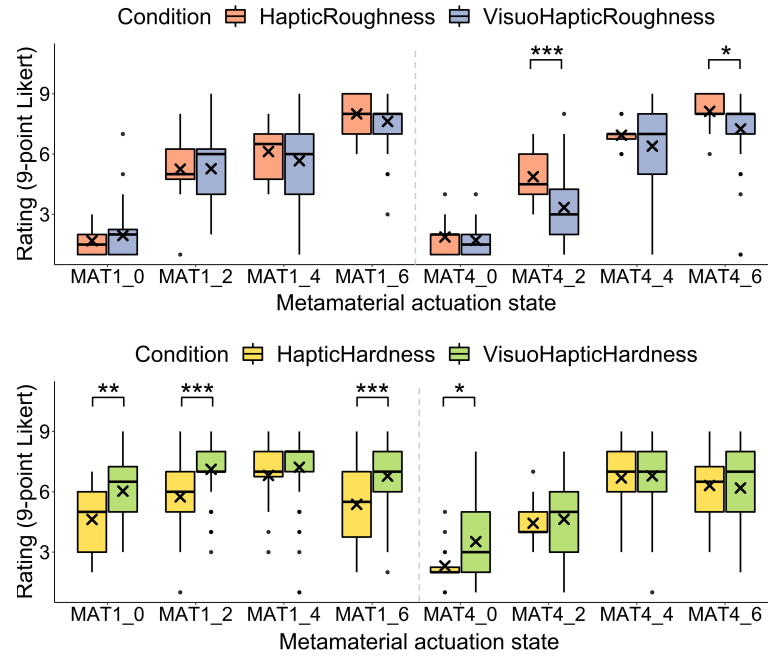


Figure 3.26: Haptic baseline vs. visuo-haptic comparison. Combined assessments for roughness (top) and hardness (bottom).

note that the fabric material was described as “jeans” in 81% of answers, which was not concerning to us. This is a confirmation that our visual materials were accurately represented in VR.

3.2.6.2 Mixed Perception Results

Visual material overlays significantly affect tactile perception

VISUO-HAPTIC PERCEPTION First, we analyze how adding visual material overlays affected the perceived roughness and hardness of our metamaterials. To do so, we compared participants’ roughness and hardness assessments from the haptic baseline to the visuo-haptic condition using multiple pairwise t-tests (Bonferroni-Holm adjusted) at each actuation state. The results are depicted in Figure 3.26 and Figure 3.27, suggesting that displaying a visual material texture on top of our metamaterials significantly affects perception.

Visual material overlays affect perceived hardness

For hardness, MAT₁, three out of four actuation states are perceived as significantly harder when a visual material was added. For MAT₄, only zero actuation shows a significant increase in reported hardness ($p < .05$). When comparing the visual baseline to the visuo-haptic condition, we found similar effects. Expected hardness based on vision was significantly different to the visuo-haptic hardness score for four out of six visual material textures: concrete ($p < .0001$), fabric ($p < .0001$), glass ($p < .0001$) and metal ($p < .0001$). Fabric was predicted to be softer, but felt harder in contrast to concrete, glass

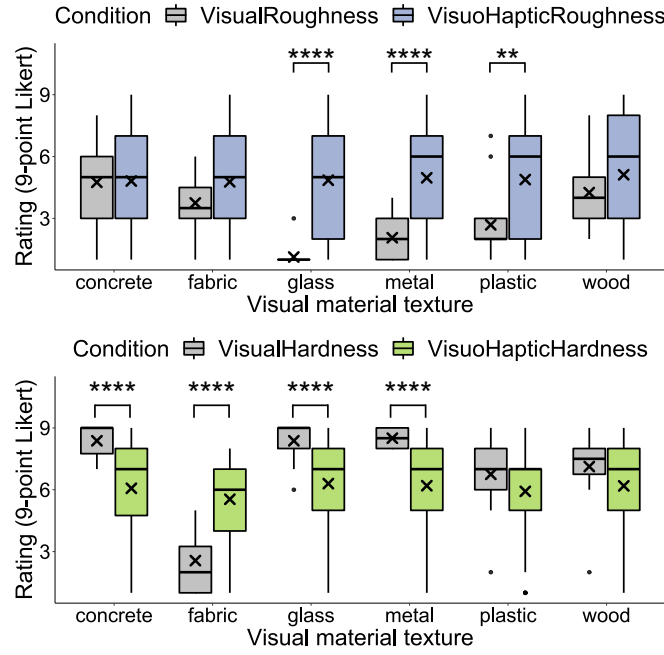


Figure 3.27: Visual baseline vs. visuo-haptic comparison. Combined assessments for roughness (left) and hardness (right).

and metal which were softer than anticipated.

For roughness, only MAT₄ at the two actuation states 2 and 6 mm were perceived as significantly less rough when a visual material was added compared to the haptic baseline. However, visual baseline and visuo-haptic condition, significantly differed for three out of six visual material textures: glass ($p < .0001$), metal ($p < .0001$), plastic ($p < .01$). Here, adding a metamaterial significantly increased roughness for the three materials. Together with our previous findings, we conclude that participants mostly relied on their haptic sense to assess roughness and hardness of a metamaterial, and that adding visual information can lead to a perceptual shift.

Visual material overlays affect perceived roughness

PERCEPTION OF SIMILARITY For each visuo-haptic combination, we recorded the similarity rating of participants' visual and haptic sensations. Figure 3.28 shows the average matching rate for each combination. For both metamaterials, we compared the ratings between the four compression levels using Friedman tests ($\alpha = .05$) with post-hoc analysis using Wilcoxon signed ranks tests and Bonferroni-Holm correction. Overall, the similarity ratings were found to significantly differ depending on the level of compression the material was in (MAT₁: $\chi^2(3) = 55.590$, $p < .001$, $W = .191$; MAT₄: $\chi^2(3) = 48.791$, $p < .001$, $W = .169$). Pairwise comparisons found significant differences between all compression states excluding 2 mm and 4 mm for MAT₁, and 0 mm and 2 mm, and 4 mm

	MAT1				MAT4			
	0 mm	2 mm	4 mm	6 mm	0 mm	2 mm	4 mm	6 mm
concrete	6.00	5.75	5.50	4.31	4.94	5.88	4.81	5.25
fabric	5.13	4.69	4.44	4.75	5.38	6.69	4.81	4.06
glass	6.81	3.56	3.56	2.13	6.13	4.5	2.88	2.06
metal	6.75	5.13	5.44	3.94	4.75	5.94	4.81	4.5
plastic	7.19	4.25	3.88	3.56	6.75	5.63	3.56	2.94
wood	6.00	6.31	5.94	4.94	5.38	5.63	5.50	5.56

Figure 3.28: Average similarity ratings for each visuo-haptic combination per metamaterial.

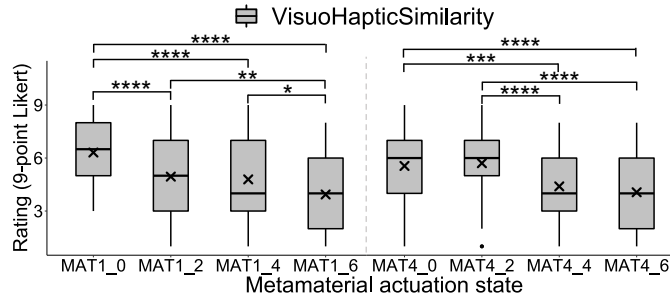


Figure 3.29: Similarity ratings in the visuo-haptic combination for MAT1 and MAT4 at each actuation state.

Similarity rating
affected by
metamaterial
compression

and 6 mm for MAT4 (see Figure 3.29). Furthermore, Kendall's tau-b correlation tests indicated a medium, negative correlation between the rating of visuo-haptic similarity and the compression state of a metamaterial ($\tau = -0.273$, $p < .001$, $N = 768$). These results indicate that as the metamaterial was compressed, the visual and haptic impressions deviated from each other, showing an increase in perceived mismatch.

Robustness of
visuo-haptic illusion

For each metamaterial, we additionally compared the similarity ratings between the six virtual textures using Friedman tests ($\alpha = .05$) with post-hoc analysis using Wilcoxon signed ranks tests and Bonferroni-Holm correction. Overall, the similarity ratings were found to significantly differ depending on the visual material overlaid (MAT1: $\chi^2(5) = 28.614$, $p < .001$, $W = .089$; MAT4: $\chi^2(5) = 25.844$, $p < .001$, $W = .081$). Pairwise comparisons indicated significant differences between the pairs of glass-concrete, glass-metal, glass-wood, and wood-fabric for MAT1, and between the pairs of glass-concrete, glass-fabric, glass-plastic, and glass-wood for MAT4 (see Figure 3.30). From these results, we see that the similarity ratings were significantly lower for virtual glass, indicating that visuo-haptic mismatches were easily perceived for this texture.




Figure 3.30: Similarity ratings in the visuo-haptic combination for MAT1 (left) and MAT4 (right) for each visual material overlay.

SUBJECTIVE MATERIAL PERCEPTION The anecdotal data of the perceived materials was further analyzed by manually extracting the materials and objects participants identified. We characterized both abstract and concrete depictions into a set of 28 distinct perceptions. Here, we found eight recurring categories, namely stone-like (stone, tile, concrete, rock, marble, pebbles), wood-like (wood, bark), glass-like (glass), fabric-like (fabric, cotton, silk, wool, jeans, pad, car seat fabric), plastic-like (plastic, laminate, epoxy, plastic doormat, toy), metal-like (metal, sheet, splinter, tank armour), rubber-like (rubber, tire), and other (water). From this, we note that our set of metamaterials combined with visual textures was able to elicit a wider range of material perceptions. To explain their perception, participants used adjectives to detail on the materials and objects in 28.52% of all indications. We further analyzed adjectives by grouping them into 4 categories, i.e., tactile, visual, material, and geometric. In most cases where adjectives were used (66.52% of all adjectives), tactile properties indicated sensations such as roughness, smoothness, or the presence of a texture. Material properties (32.16%) were used to indicate the state of a perceived object, e.g., damaged, broken, or worn. Furthermore, a small amount of adjectives referred to either visual properties (0.88%), e.g., shiny, or geometric properties (0.44%), e.g., small. From this, we see that adjectives were mostly referring to users' tactile perception. We further investigated the relationship between the rating of visuo-haptic similarity, the compression state of the metamaterials, and the subjectively perceived materials. Kendall's tau-b correlation tests indicated a small but significant positive correlation between the use of an adjective and the compression state of a metamaterial ($\tau = .110$, $p < .001$, $N = 768$). However, there was no correlation found between the visuo-haptic similarity rating and the compression state of a metamaterial ($\tau = .038$, $p = .226$, $N = 768$). From this, we see that adjectives were likely not used to explain mismatches between users' visual and haptic perception, but more likely to elaborate on the increased amount of surface features upon compression of the metamaterials.

Participants described material states

Participants used adjectives to describe their experiences

3.2.6.3 Summary

*Metamaterials
enhance tactile VR
experiences*

Our main study supports the findings and hypothesis **H1.1** and **H1.2** from our first preliminary experiment. Hence, we conclude that metamaterial patterns MAT₁ and MAT₄ can create distinct levels of roughness, and especially MAT₄, also for hardness. When combining our metamaterials with visual material textures, we often observed a drift in perception **H2.1**. For example, glass was rated with the lowest roughness score; however, when combined with our metamaterials received significantly higher scores. Thus, participants relied more on their haptic sense in case of increasingly conflicting sensory (haptic and visual) information. Our analysis of the similarity ratings suggests that as the metamaterial was compressed, the visual and haptic impressions deviated from each other, showing an increase in mismatching perception—most likely caused by the lack of adaption of the virtual materials. This was supported by the high similarity ratings for the 0 and 2 mm actuation states, contrary to the significant decline at 4 and 6 mm (**H2.2**).

3.2.7 Discussion of Metamaterials

Motivated by recent advancements in the field of fabrication, we investigated the use of metamaterial textures to influence tactile perception. Next, we discuss the potential impact of our investigation and its current limitations and provide directions for future work.

3.2.7.1 Metamaterials for Touch Experiences

*Compressing
metamaterials affects
perception*

Our work started out with the design and fabrication of five different metamaterials. Through lateral compression, the surfaces were able to dynamically change their physical configuration, which influenced their compliance and surface texture. From the results of our preliminary perceptual user study, we show these different physical configurations of our metamaterials translate to variations in their perceived tactile hardness and roughness. While we found a clear relationship between the compression state of a metamaterial and its tactile perception, not all designs consistently influenced perception in the same manner. Our results contribute to the field of fabrication and underline that metamaterials are able to extend existing methods for fabricating ‘feel’ aesthetics [344]. The careful construction of metamaterials allows for the design of different tactile states that are able to communicate different sensations. Our psychophysical investigation bridges the gap between fabrication of metamaterial textures [159] and their perception. In a further step, a perceptual modeling approach would enable procedural design of different tactile states in single fabricated samples [271].

3.2.7.2 *Metamaterials for Passive Haptic Feedback*

In a second study, we evaluated the use of a subset of our metamaterials in a virtual setting. The baseline analysis shows that the visual properties of our set of virtual textures clearly communicated different impressions. For our physical metamaterials, the haptic baseline indicated significant differences in roughness perception between compression states, with limited differences in hardness perception. By combining visual and haptic stimuli, we see that our metamaterials can serve as passive haptic proxies for virtual textures. Depending on the virtual texture and the metamaterial used, visuo-haptic integration was able to serve a consistent tactile experience. This is underlined by the similarity ratings of the visual and haptic stimuli. However, discrepancies between visual and tactile features drove perception to align with the haptic stimuli, causing haptic dominance. A similar effect shown by Iesaki et al. [157], stated that although tactile impressions can be intentionally changed by providing appropriate visual stimulation, the coarseness of the visual and tactile textures has to be close to each other. This is further illustrated by the indicated material perceptions. As highly matching sensations were able to elicit clear material perceptions, discrepancies were explained through the use of adjectives. Our work builds upon previous work in combining visual and haptic impressions to influence material perception [69, 182], and fabrication technologies for designing varying psychophysical impressions [66, 139, 271]. From our results, we see that haptically-varying metamaterial surfaces enable material perception in combination with visual textures. As users actively try to make sense of sensory input, further work investigating the thresholds of visuo-haptic discrepancies is needed to understand how consistent material perceptions can be simulated. As underlined by the field of haptic design, there is a pressing need for intuitive and comprehensive approaches to enable designers to create and share convincing and immersive experiences [68, 292, 293]. Rather than taking an active feedback approach [66, 366], we envision the use of fabrication technologies for designing haptic experiences. Specifically in virtual environments, the visuo-haptic perception of fabricated artifacts is able to drive on-demand creation of haptic experiences. Our work presents an initial investigation into the use of metamaterials for tactile experiences through passive haptic feedback in VR. Our results are bound by several limitations that should be improved.

Adding visual feedback shifts tactile perception

3.2.7.3 *Study Limitations*

In this first investigation, we asked participants to explore our metamaterials through circular lateral movement using their index finger, which was monitored by the experimenter. Since the perception of textures is highly motion dependent [290], the lack of rigid control

*Motion dependence
and magnitude of
perceived change*

might have influenced our results. This needs to be addressed in future psychological experiments by, e.g., only allowing passive exploration through a robotic device, controlling for various interaction parameters such as velocity, pressure, and trajectory. Future work also needs to improve the understanding of how the compression rate of metamaterials affects the magnitude of, e.g., perceived roughness. For example, by determining the Just-Noticeable-Difference (JND) between actuation states to more deeply investigate their relationship.

3.2.7.4 Visuo-Haptic Matching

*Interaction with
materials was not
visually supported*

A main limitation illustrated by our results is the lack of appropriate adaptation of the used virtual models. In our naive approach, we did not alter the virtual model overlaid onto the fabricated samples. As the lateral compression of our samples influenced their physical configuration, participants clearly noted discrepancies between their tactile and visual perception. For high levels of compression, this visuo-haptic mismatch led to haptic dominance, i.e., the scenario where participants relied on their haptic sense to guide their entire perception. Future iterations need to ensure that the physical metamaterials are correctly aligned with the virtual object they are paired with [256]. To adapt virtual models appropriately, an accurate approach is needed to record the state of the metamaterial in question while tracking the user's interaction with the sample. Recent advancements in multi-material fabrication have allowed for the integration of conductive materials that can be used for capacitive sensing. This would allow tracking of a user's finger during interaction. Furthermore, recent work has shown that sensing mechanisms can be directly embedded into the fabrication process to track the state of the metamaterial [125]. Through such sensing approaches, the virtual representation can be adapted to emphasize the presence of more features.

3.2.7.5 Actuating Metamaterials

*Actuation approach
needs further
refinement*

Our metamaterials changed their physical configuration through lateral compression. For this, we designed and created a setup that included an electromotor for actuating linear gears, controller by a microcontroller. Internal support for the surfaces was provided through embedded metal pipes that controlled unwanted deformations by stabilizing the samples. During our perceptual studies, we used pre-compressed samples to ensure consistency and to reduce the study duration. As our current setup remains bulky and lacks precision, there is certainly room for improvement. For example, the actuating elements can be embedded directly in the metamaterial design.

3.2.7.6 *Fabrication of Perceptual Features*

One of the key limitations that we encountered while prototyping was the metamaterial's feature size that we could reliably print with an off-the-shelf FDM 3D printer. We believe that there is still potential to improve upon our results and create other metamaterials with different surface texture information. Perhaps other fabrication methods such as stereolithography (SLA) 3D printing, can be used to fabricate metamaterials, including smaller features, which may benefit the haptic resolution. In addition, approaches such as *Tactlets* [139] which enable rich electro-tactile feedback on various 3D-printed geometries, may help to overcome such limitations. To further increase the resolution, our approach can be combined with methods that capture real-world surface information [67]. Finally, this can be used in perceptual modeling to ensure the user's perception aligns with the designer's intention [271]. At this early stage, we did not conduct a technical evaluation investigating longevity and potential issues caused by material deterioration. We observed that our actuation approach suffers from inaccuracies that need to be overcome in future iterations, and therefore, we decided to use fixed samples in our experiments.

*Alternative
fabrication methods*

3.2.7.7 *Visuo-Haptic Integration*

Our method builds upon passive haptic feedback approaches that pair physical objects with virtual representations to create visuo-haptic experiences in VR. As pointed out by related work, such an approach has the potential to improve the user's sense of presence, positively impacting their experience [158]. The idea of metamaterials builds upon the concept of fabricating abstract objects that influence tactile dimensions in order to build texture or material impressions in combination with visual information [69]. By capturing the baseline assessments, we were able to receive an initial impression of what users might expect from both modalities. For example, when visually assessing a texture, some users might interpret the texture completely differently, leading to different tactile expectations. This might have caused a bias that influenced their further assessments. Through the visuo-haptic combinations, we aimed to understand how different combinations matched and how they were able to build consistent material impressions. Future work would need to explore how the presence of roughness and hardness guides material perception in a more general manner and how immersive environments benefit from these combinations in terms of realism and presence.

*Merging visual and
haptic sensory
inputs*

3.2.7.8 *Beyond Touch*

Expanding our approach to real-world use cases may create the demand for local and global controllability, rather than uniform compression as designers will need to integrate local texture variations.

*Not limited to form
factor*

We believe that more advanced metamaterials can fulfill these requirements because they can actuate themselves in a cell-by-cell fashion [368]. We are also interested in taking this approach beyond one-finger-touch interactions. For example, Ion et al. [159] presented a bike handle, differing in friction, depending on the compression of the metamaterial. Similarly, we could envision metamaterials that augment controller-based interfaces [363, 382]. Here, auxetics metamaterial designs might be worth investigating because their form factor remains unaffected upon actuation [320]; instead, our current metamaterial designs become smaller, reducing the interaction area.

3.2.8 Conclusion & Contributions

*Selection of most
promising
metamaterials*

In this section, we contribute to **RQ1** by presenting a novel concept to simulate tactile experiences in VR using dynamic 3D-printed metamaterial compression in combination with visual material overlays (**C1**). We outline the design and fabrication process of five metamaterials that can be fabricated with off-the-shelf hardware and we open-source our 3D models to the community (**C6**). We bridge the gap between fabrication and user perception of our metamaterial designs by conducting two experiments. As a first proof of concept, we investigated how participants perceive metamaterials and their compression states across the perceptual dimensions of roughness and hardness. In summary, MAT1–MAT4 showed promising results in influencing roughness, but only MAT4 for hardness, showing that not all designs consistently influenced perception in the same manner. Based on these insights, we selected MAT1 and MAT4, providing the widest range of tactile impressions.

*Metamaterials can
be an effective asset
in haptic proxy
design*

Next, we conducted our main experiment, exploring the use of these two metamaterial patterns as passive haptic proxies for simulating visuo-haptic material experiences in VR by combining them with six visual material texture overlays. Our results suggest that metamaterials can act as VR proxies and support tactile material perception through visuo-haptic stimuli. Our approach contributes to the field of fabrication and underlines that metamaterials can be an effective tool in designing haptic VR experiences through proxies (**C1**). The design and compression of metamaterials allow one to create different tactile states that can communicate different sensations. This significantly enhances the state-of-the-art as it allows dynamic shape-change, according to the concept of DPHF, of the metamaterial upon compression. Therefore, they can provide a wide range of tactile sensations without requiring additional fabrication. Together, this paves the way for truly multipurpose proxies as metamaterials may be applied to existing controllers and devices. Thus, metamaterials could also be used in combination with the *VoxelHap* toolkit presented in

*On the way to
multipurpose
proxies*

the previous section, augmenting Voxels and/or Plates. Because it acts as an add-on, it can be combined with approaches that allow for direct and embodied interaction. Finally, by introducing mismatches between what users feel and what they see, we also contribute to (C3). The process of multisensory integration between visual and tactile features drove perception to align more with the haptic stimuli, causing haptic dominance. This contributes to the theoretical knowledge and understanding of the robustness of hand-based illusions in VR, showing that although tactile impressions can be intentionally changed by providing appropriate visual stimulation, the coarseness of the visual and tactile textures has to be close to each other.

Conflicting sensory information

3.3 SUMMARY

In this chapter, we presented two novel approaches for proxy design, going beyond focusing only on a single haptic modality or severely limiting the type of interactions inside IVEs (RQ1). As a result, our approaches enhance the field of proxy-based haptic feedback while tackling the grand challenge of *Similarity* [256].

In Section 3.1.1, we presented the *VoxelHap* toolkit which enables Lisa to construct highly functional proxy objects (i.e., replicas) using Voxels and Plates. Voxels are blocks with special functionalities that build the core of each physical proxy. They enable functional proxy parts with sensing and actuation capabilities to ensure high visuo-haptic synchronization during object manipulation. Additional Plates increase a proxy's haptic resolution, such as its shape, texture, or weight, resulting in a holistic, multimodal haptic experience (C1). The toolkit comes as an end-to-end system, providing open-source hardware and software tools that support the construction and reconfiguration of proxies (C6). We demonstrate the capabilities of the *VoxelHap* toolkit through our technical evaluation and the construction of a range of fully functional proxies (C5). In two experiments with 24 participants, we evaluate a subset of the constructed proxies, studying how they compare to a traditional VR controller. First, we investigated *VoxelHap*'s multimodal haptic feedback, showing that it can provide a variety of tactile and kinesthetic haptic sensations while allowing dynamic touch and exploratory interactions (C1). Second, we investigated the trade-offs of using additional shape resolution by applying *VoxelHap*'s ShapePlates. Here, our findings show proxies with higher shape fidelity outperform traditional controllers and were generally favored by participants. However, the greater *Similarity* between the virtual object and the *VoxelHap* proxy led to uncanny valley effects where participants' sensitivity to visuo-haptic mismatches increased (C3). Finally, one of the main drawbacks of the approach is the manual reconfiguration required to

VoxelHap proxies for multimodal haptic feedback

change the properties of the proxy, slowing down its performance.

Therefore, in Section 3.2, we create more flexible and dynamic proxies by combining actuation with 3D-printed metamaterials. To this aim, we proposed metamaterial structures able to alter their tactile surface properties, e.g., their hardness and roughness, upon lateral compression. First, we designed and open-source five different metamaterial patterns based on features that are known to affect tactile properties (C6). In a preliminary experiment with six participants, we evaluated whether our samples were able to successfully convey different levels of roughness and hardness sensations at four levels of compression. While we found that roughness was significantly affected by compression state, hardness did not seem to follow the same pattern. In a second experiment with 16 participants, we focused on two metamaterial patterns that appeared to be promising for roughness and hardness and investigated their visuo-haptic perception in VR. Here, eight different compression states of our two selected metamaterials were overlaid with six visual material textures. Our results suggest that the visuo-haptic combinations enhanced tactile material perception in VR through which a single physical proxy can provide a wide range of tactile sensations (C3). Additionally, when asked which material participants perceived, adjectives such as “*broken*” and “*damaged*” were used. This indicated that metamaterial surface textures could be able to simulate different object states. Our results underline that metamaterial design is able to extend the gamut of tactile experiences of 3D-printed surface structures, as a single sample is able to dynamically reconfigure its haptic sensation through compression (C1).

*Dynamic
metamaterials can
enhance texture
perception*

Finally, both research projects can be situated in the field of haptic proxies for VR, and while they provide valuable contributions, they share the underlying limitations of any physical artifact that cannot yet be easily overcome.

Videos about the work and prototypes presented in this chapter are available online and can be accessed through the QR codes. Images and parts of the text in this chapter, as well as the presented figures, tables, ideas, concepts, implementations, applications and uses cases, studies and experiments, results, discussions, and conclusions, have been published previously in:

[104] **Martin Feick**, André Zenner, Oscar Ariza, Anthony Tang, Cihan Biyikli, and Antonio Krüger. *In Proceedings of ACM UIST 2023*. Turn-It-Up: Rendering Resistance for Knobs in Virtual Reality through Undetectable Pseudo-Haptics.



Video

[95] **Martin Feick**, Niko Kleer, André Zenner, Anthony Tang, and Antonio Krüger. *In Proceedings of ACM CHI 2021*. Visuo-haptic Illusions for Linear Translation and Stretching using Physical Proxies in Virtual Reality.



Video

4.1 TOWARDS PSEUDO-HAPTIC RESISTANCE

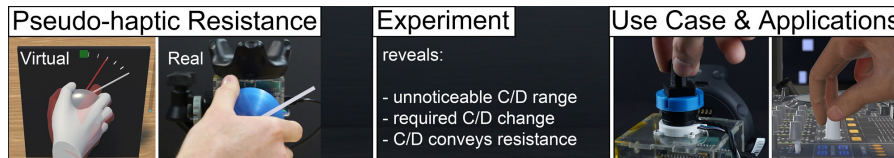


Figure 4.1: *Pseudo-Haptic Resistance* can be created by ‘*slowing down*’ a user’s real-world rotations of a knob. As a result, the red needle rotates slower than the white needle.

4.1.1 Introduction

Proxy objects underlie the physical constraints of fabrication, materials, and energy management that cannot be easily overcome. Many haptic VR controller and proxies, including our approaches developed in Chapter 3, suffer from this general limitation. However, hand-based illusion techniques for VR may help to address these limitations as they can ‘*trick*’ Lisa’s perception into experiencing, for instance, haptic feedback without the need of physically rendering it [227]. Yet, while some illusion techniques can render haptics, there still exists a high demand for unnoticeable hand-based illusion techniques for

*Undetectable
hand-based illusions
to enhance haptics*

Approaches to
haptically render
virtual buttons

Haptic rendering for
knobs is missing

VR (**RQ2**) that can be used to improve the haptic resolution of proxies and devices. Through *Pseudo-Haptics* [227], we can provide haptic feedback based solely on visual stimuli, taking advantage of the visual-dominance over proprioception [47]. Recently, this technique has been successfully used to simulate 3D button presses in mid-air [32]. Ariza Nunez et al. [15] presented the *Holitouch* technique, conveying the holistic sensation of stiffness, contact, and activation when pressing mid-air buttons. They used a wearable device combining tapping and vibrotactile sensations on the fingertip and C/D ratio manipulation to render various types of buttons. There exists a considerable amount of work on buttons, but knobs, i.e., rotational elements that are part of many interfaces with clear affordances in the real world, have received little to no attention. This prevents VR designers from incorporating such elements in VR as it is unclear how appropriate haptic feedback can be achieved through proxy interfaces, and how many proxies would be required to render virtual knobs that differ, for example, in resistance. To fill this gap, we designed a novel type of pseudo-haptic effect, called *Pseudo-Haptic Resistance*, which combines rotational C/D gain manipulations with a physical proxy knob.

4.1.2 Undetectable Pseudo-Haptic Resistance

Changing perceived
haptics while
remaining
unnoticeable

In this section, we introduce our *Pseudo-Haptic Resistance* and outline the design and implementation of the prototype that provides physical resistance. Finally, we show the most relevant findings from our preliminary experiment. Figure 4.1 illustrates *Pseudo-Haptic Resistance* as it is experienced by users. Here, the user has physically rotated the knob 60° , but the virtual knob has only rotated 30° . We create this effect by scaling the user's real-world movements (Control) with what is displayed in the virtual world (Display) with the C/D ratio. This is a common approach to create illusions in VR, for example, pseudo weight [288, 323]. Therefore, we opted for this method scaling down users' real-world movements, resulting in smaller virtual movements than physically performed, which can be achieved with a C/D gain factor ≤ 1.0 . We hypothesize that by combining this technique with functional proxies that restrict users' movements [99], it may be possible to suggest physical resistance through a pseudo-haptic effect [227]. In this work, we are only interested in *increasing* the perceived resistance since this would require embedding larger actuators into the hardware to achieve the same haptic sensations with alternative, active haptic approaches.

4.1.2.1 Proxy Knob

We built a proxy knob consisting of an ESP32, a brushless DC Motor2204 260KV, an ICTMC6300 motor driver and magnetic position sensor (MAQ473) that enables sensing and actuation (see Figure 4.2).




Figure 4.2: Proxy knob.

The motor was controlled using FoC [309] through a proportional controller, which after an immediate (unnoticeable) ramp, provides a constant level of resistance. We laser-cut an acrylic case for the hardware components, and designed and 3D-printed a proxy knob, measuring 7 cm in diameter, with a custom mount that a user can comfortably hold and rotate. In addition, we built a table mount using a Manfrotto 3D pan/tilt tripod head with individual axis control attached to a laser-cut wooden plate with adhesive tape on the bottom. The proxy knob can generate between 0.3–3.2 Nmm of resistance, measured with a PCE-FB 50N force gauge. A corresponding virtual replica of the knob and its functionality was implemented in Unity3D (v.2022.2.1f1). The proxy knob can be rotated to any desired position, which is directly streamed using serial port communication (baud rate 155200), ensuring direct coupling of the real and the virtual environment.

Hardware setup we built

The knob can render rotational forces

4.1.2.2 Preliminary Experiment

We conducted a small preliminary experiment with four participants to understand how users perceive our proposed pseudo-haptic technique and inform the design of our main experiment. To do so, we asked novices to rotate two knobs and compare them. One knob used a one-to-one mapping between the real and the virtual rotations; on a second knob, we applied C/D gain manipulations of different magnitudes (0.1–1.0). The physical resistance provided by the device was set to the minimum (0.3 Nmm). Then, we specifically asked participants about the differences between the knobs with respect to their properties, their general impression of the visualization and the provided haptic feedback. We did not specifically ask them any questions that would point them towards our pseudo-haptic effect. Below, we summarize the most relevant feedback and suggestions:

Perceptual effects of rotational C/D gains

Range of C/D gains. All participants reported that the virtual knob rotated noticeably slower than the real knob when applying a C/D

gain below 0.6. Therefore, in the following main experiment, we set the minimum C/D gain to 0.5 to ensure that we would stay within a reasonable range of C/D gain manipulations.

Pseudo-Haptic Resistance. P1 and P3 felt the virtual knob sometimes had more “friction”, while P4 described this sensation as feeling “resistance”. These sensations corresponded to when we applied a C/D gain below 1.0. This suggests that participants were able to perceive the pseudo-haptic effect.

We identified the following themes

Visuo-Haptic Integration. To assist participants during the task, we visualized tick marks, creating a dial-like layout (see Figure 4.1). P2 commented that s/he expected to also haptically “feel” these tick marks. We found this to be very intriguing, and because our proxy knob and FoC [309] supported the haptic rendering of tick marks, we included this variation, with and without tick marks, in our main experiment. Following multisensory integration theory [80], higher visual-haptic congruence provided by tick marks on the knob could potentially allow for more offset between real and virtual rotations to go unnoticed, similar to what has been found for hand realism [262].

Expected Physical Resistance. P2, P3 and P4 stated that the physical knob rotated “too loosely”, and suggested that it should have more physical resistance. As a result, we asked them to increase the physical resistance of the knob (C/D gain = 1.0) using the keyboard, until it matched their expectations. We took the average of the three values (1.1 Nmm) and used it as the default value in our main experiment.

4.1.3 Experiment

Three-stage psychophysical experiment

We designed an experiment to investigate: (1) the DTs for manipulating the C/D gain when rotating knobs (i.e., how much manipulation goes unnoticed), (2) the Just-Noticeable-Difference (JND) in resistance caused by C/D gain manipulations (i.e., how much manipulation is required to change the user’s perception of resistance), and (3) how C/D gain manipulations translate to perceived physical resistance. To do so, we conducted an experiment using different psychophysical methods described in Section 2.5.2.2. Below, we briefly summarize each method’s configuration. We formulated the following five hypotheses:

H1: Considerable offsets between rotations of the real and virtual knob can be introduced without users noticing.

H2: By increasing visuo-haptic synchronization, the range of unnoticeable offset becomes larger.

H3: Users associate C/D gain manipulations with changes in physical resistance.

H4: Within the range of unnoticeable C/D gains, designers can effectively manipulate the perceived physical resistance.

H5: With unnoticeable C/D gain manipulations, we can achieve multiple distinguishable levels of perceived resistance.

4.1.3.1 *How much C/D gain manipulation goes unnoticed?*

To determine DTs (**H1** and **H2**), we used an established adaptive psychophysical 1-up-1-down interleaved staircase procedure, exposing participants to different stimuli (C/D gains) repeatedly. Using a fixed step size, we target the CDT or point of subjective equality (see Section 2.5.2.2). To summarize, an interleaved staircase uses an ascending and a descending sequence, and randomly assigns the next trial to one of the sequences. The procedure increases the next stimulus in a sequence if a participant fails to detect the current stimulus in this sequence, and decreases the next stimulus if the user detects the offset. A directional change within a sequence is called a reversal point, which we used as a convergence criterion ($r = 4$). Since we are the first to explore rotations, we used our preliminary experiment to determine range and start values for the procedure. We found a fixed step size (i.e., changes in the C/D gain) of 0.05 to be appropriate and selected 1.0 and 0.5 as starting values.

Determine DTs for rotational C/D gains

TASK Participants were asked to rotate the (proxy) knob until the virtual knob matched a target position displayed in the virtual world (at 60°)—thus, the virtual distance remained equal (see Figure 4.3). After they successfully established the position, a one-second dwell-time indicator appeared, and they were required to maintain this position, before rotating the knob back to the start position. Next, a 1AFC ('yes', 'no') question appeared, and they responded to the following statement: "My real hand rotated further than my virtual hand" following Steinicke et al. [322]. Participants were instructed to perform a 'normal' rotation of the knob at a comfortable speed. Participants were not allowed to repeat the rotation. The physical resistance of the proxy knob was kept the same throughout this part of the experiment.

Rotate and report if visuo-proprioceptive offset is present

4.1.3.2 *How does changing the C/D gain affect the perceived resistance?*

We used a 2IFC procedure to determine the JND in resistance (**H3** and **H4**) [181]. We exposed participants to two successive stimuli (C/D




Figure 4.3: Experimental task. Participant rotates the knob to target location (A) and returns to the initial position (B).

C/D gain required to trigger the haptic effect of resistance

gains), one being the baseline (C/D gain = 1.0) and a second stimulus being randomly picked from the range from 1.0 to 0.5 in 0.05 steps. Both stimuli were presented to the participant in a random order. The stimuli were generated using the 1-up-1-down interleaved staircase procedure outlined above. We kept step size, target probability (here: 50%-correct) and start values the same to ensure comparability.

Rotate twice and compare both rotations in terms of their physical resistance

TASK Similar to the above, participants were asked to rotate the (proxy) knob until it matched the target position. However, this time they had to repeat this interaction and then compare both rotations of the knob in terms of their physical resistance, i.e., one baseline (C/D gain = 1.0) and one stimulus (C/D gain ≤ 1.0). Because the physical resistance of the proxy knob was kept the same, participants rated the pseudo-haptic effect created by the C/D gain manipulation. Next, a 1AFC ('yes', 'no') question appeared, and they responded to the following statement: "The physical resistance of two knobs felt the same". Participants were prevented from repeating the rotation and could only explore each of the knobs once.

4.1.3.3 How do C/D gain and perceived resistance relate to each other?

Relationship between pseudo-haptic resistance and physical resistance

In the last part of our experiment, we used the method of adjustments to investigate how C/D gain and perceived resistance relate to each other (**H5**). To this end, participants experienced two knobs, initially not differing in their physical resistance, but only in their C/D gain [288]. We asked them to adjust the physical resistance of the (red) knob, with a C/D gain = 1.0, until it matched the resistance of the (blue) knob, with a C/D gain ≤ 1.0 (see Figure 4.4). Participants had 25 sec to adjust the resistance following Samad et al. [288]'s methodology. After 25 sec, the configured resistance was logged and could not be changed anymore. However, we only continued with the next trial when participants confirmed that they felt ready, allowing them to take breaks in between. Please note that participants were informed that no physical difference between the two knobs is also a valid option, which was explained and shown to them during the warm-up. Nevertheless, each trial started with the base level, which is different




Figure 4.4: Method of adjustments. Participants can switch between the two knobs (only differing in C/D gain) using the right arrow key. They can adjust the physical resistance of the (red) knob so that it matches the (blue) knob. They had 25 sec to match the resistance of the two knobs.

from what was used by Samad et al. [288] and thus may have affected the results.

TASK In this part of the experiment, participants could rotate the (proxy) knob in whatever way they preferred. Yet, they could only switch between the two virtual knobs (red and blue) when the needle was pointing upwards. Initially, the physical resistance of the proxy knob was the same for both virtual knobs, as in previous parts of our experiment. A progress bar showed the possible range of physical resistance that could be configured in 0.1 Nm·m increments. We tested C/D gains ranging from 1.0 to 0.5 in 0.1 increments (i.e., 1.0, 0.9, 0.8, 0.7, 0.6, 0.5) in a random order. We decided on this after rigorous pilot testing, allowing us to collect more data points per participant to increase robustness. To do so, we repeated the procedure three times, resulting in $3 \times 6 = 18$ configurations each participant adjusted.

Configure resistance of one knob to match another

4.1.3.4 Design

In this experiment, we used a within-subjects design. We had four study parts: two times DTs for rotating knobs with and without haptic tick marks, JND in resistance and a condition that studied how C/D gain affects the perceived physical resistance. We fully counter-balanced the conditions using a Latin square ($n = 4$).

4.1.3.5 Participants

We recruited 20 right-handed participants (nine females, eleven males), aged 20–38 (mean = 26.42; SD = 3.65) from the general public and the local university. This excludes one participant (P12), in whose case we had to stop the experiment due to system failure. Participants had a range of different educational and professional backgrounds including media informatics, computer science, education, pharmacy, anglistics, neuroengineering, embedded systems,

*Healthy mix between
novice and expert
VR users*

data science and artificial intelligence. All participants reported normal or corrected-to-normal vision and did not report any known health issues which might impair their perception. Eight participants had never used VR before, six had used it a few times (1–5 times a year), no one reported using it often (6–10 times a year), and six others used it on a regular basis (more than 10 times a year). Ten participants reported that they had not played VR games before, five people responded sometimes or infrequently (1–5 times a year), one often (6–10 times a year), and four people on a regular basis (more than 10 times a year). Participants not associated with our institution received €10 as remuneration for taking part in the experiment. The study was approved by the Saarland University’s Ethics Board.

4.1.3.6 Apparatus

*The following
hardware setup was
used in the study*

Our apparatus consisted of a HTC VIVE Pro Eye system and our implemented knob prototype shown above. The virtual scene was aligned based on the position of the physical knob using a Vive tracker, ensuring one-to-one mapping of the virtual- and real-world setup. On the software side, we used SteamVR (v.1.17) with the OpenVR SDK (v.1.1.4). We used a simple virtual scene, consisting of a table, the virtual knob, and an instruction screen, which was developed in Unity3D (v.2022.2.1f1) and was running on an Acer Predator Orion 5000 PO5-615s offering an Intel® Core i9 10900k CPU, 32 GB RAM, and an Nvidia® GeForce RTX 3080. We used a simple and fixed hand representation to prevent unwanted effects [262]. The experimental logic was implemented using the Unity Experiment Framework (UXF v.2.4.3) [42] and the Unity Staircase Procedure Toolkit [379]. Participants remained seated on a chair throughout the experiment, and we supported their arm position with a pillow to reduce fatigue. Participants’ responses were collected using a keyboard. The default resistance of the knob was taken from our preliminary experiment.

*Tactile cues to
enhance visuo-haptic
integration*

HAPTIC TICK MARKS. To ensure high synchronization, the haptic ticks were timed with the virtual marks. To achieve this, we had to scale them up/down, depending on the applied C/D gain. The tick marks were rendered by changing the proportional controller’s set-point value when the knob reached the point halfway between two tick marks.

4.1.3.7 Experimental Protocol

After giving participants a general introduction to the study and obtaining their informed consent, they filled in the demographics questionnaire. Following this, we showed them the physical knob and explained how it should be used. Next, they were introduced to VR

and the task through an open-ended practice round. Since all study conditions significantly differed, participants were given enough time to familiarize themselves. Participants were instructed to grasp the proxy object as indicated by the virtual hand and were told to maintain this pose, not readjusting their grip, which was monitored by the experimenter. We instructed participants to respond to the question as quickly as possible by using a keyboard with their non-dominant hand. After completing all conditions, participants filled in a SSQ [177]. The total experiment took about 60 min per participant.

Study procedure for each participant

4.1.3.8 Data Collection

We collected data from six sources: a pre-study questionnaire for demographic information; the subjective responses to the forced-choice questions; the configured physical resistance; field notes and observations; a short post-study interview and a SSQ in VR using our *VRQuestionnaireToolkit* [94].

4.1.3.9 Analysis

We analyzed our data using a One-Way Repeated Measures ANOVA. First, we removed significant outliers using the box plot method and verified the normality assumptions at $\alpha = .05$, using a Shapiro-Wilk test. We checked the assumption of sphericity using Mauchly's test and applied Greenhouse-Geisser corrections when sphericity was violated. Post-hoc pairwise t-tests were corrected using Bonferroni-Holm adjustments.

Traditional frequentist analysis

To further investigate our data, we conducted a Bayesian analysis using the BayesFactors R package¹ with default priors (v.0.9.12+-4.4). ANOVA effects are reported as the Bayes factor for the inclusion of a particular effect (BF_{incl}), calculated as the ratio between the likelihood of the data given the model with the effect vs. without that effect [178]. Additionally, we performed paired Bayesian t-tests using default effect size priors. Results are reported as two-tailed Bayes factors BF_{10} and effect size estimates as median posterior Cohen's δ with a 95% credibility interval (95%CI) [178].

Complemented by Bayesian analysis

4.1.4 Results

In this section, we report our results. First, we look at the DTs, comparing knob rendering with and without haptic marks to investigate **H1** and **H2**. Next, we analyze the results of JND and contrast them with the DTs regarding **H3** and **H4**. Finally, we plot participants' perceived physical resistance against the tested C/D gains (**H5**). The SSQ

¹ BayesFactor GitHub: <https://tinyurl.com/mr3v3jyp>. Last accessed: Nov 1, 2024




Figure 4.5: Results from our threshold experiments show that only a small C/D gain is needed to produce the illusion of resistance. The required C/D gain is within the unnoticeable range.

results did not suggest simulator sickness caused by exposing participants to, sometimes, noticeable offsets above their thresholds (Total Severity (TS) score: mean = 18.04, SD = 7.02) compared to [2, 95]. For each participant, we computed DTs by averaging the last three reversal points within the ascending and descending staircase sequence. All staircase plots are available in the supplementary materials of the corresponding paper.

4.1.4.1 How much C/D gain manipulation goes unnoticed?

Figure 4.5 shows that we could introduce substantial offsets between real and virtual rotations of the Knob (mean = .67; SD = .08) and Knob + Ticks (mean = .67; SD = .08), without participants noticing it, confirming **H1**. Our statistical analysis showed a main effect ($F_{(2,38)} = 21.96$, $p < .0001$, $\eta^2 = .431$, $BF_{incl} = 773384$); however, post-hoc analysis did not suggest a significant difference between the Knob and Knob + Ticks condition ($p = .861$, $\delta = -0.05$). In fact, our Bayesian analysis revealed evidence for the absence of an effect ($BF_{01} = 3.21$, with median posterior $\delta = .038$, 95%CI = $[-.592, .509]$), suggesting that it is 3.21 times more likely to observe this data under the null hypothesis. Thus, we reject **H2**, because based on our collected sample, we could not identify an effect on the C/D gain tolerance caused by higher visuo-haptic congruence. Despite this not being reflected in the DTs, participants frequently commented: “*This feels so satisfying*” (P13), “*So nice..so nice*” (P19) or intentionally moved slower during the warm-up, noting “*the ticks are in perfect synchronization*” (P5).

Substantial amount
of rotational C/D
gain remains
undetectable

4.1.4.2 How does changing the C/D gain affect the perceived resistance?

Interestingly, the CDTs for C/D manipulations found in the first part of the experiment are significantly greater than the JND (mean = .81; SD = .06) in resistance found in the second part ($p < .0001$, $Knob_{\delta} = 1.96$, $KnobTicks_{\delta} = 1.80$) (see Figure 4.5). This was supported by the

Bayesian analysis, providing very strong evidence for the existence of an effect, both for Knob ($BF_{10} = 35600$, with median posterior $\delta = 1.824$, $95\%CI = [1.049, 2.606]$) and Knob + Ticks ($BF_{10} = 8460$, with median posterior $\delta = 1.662$, $95\%CI = [.907, 2.423]$) against the JND. Since the physical resistance did not change in this part of the experiment, we conclude that, similar to our preliminary experiment, participants in fact associate rotational C/D gain manipulations with changes in physical resistance (**H3**). Moreover, as a central finding, we could show that within the range of unnoticeable C/D gains, designers can effectively change the perceived physical resistance of the knob proxy (**H4**).

Only small C/D gain necessary to trigger haptic effects

4.1.4.3 How do C/D gain and perceived resistance relate to each other?

Finally, we looked at the resistances perceived when different C/D gains are applied. Initially, the resistance was set to 1.1 Nmm, which pilot study participants stated to be most realistic, given the visualization, form and feel of the knob, and the investigated interaction. In the last part of our experiment, we obtained three configured physical resistances per C/D gain, which were averaged for each participant and then statistically compared. The results can be seen in Figure 4.6, showing a robust baseline and an increase in configured physical resistance with decreasing C/D gains. We found very strong evidence for a main effect ($F_{(1.72, 29.29)} = 27.34$, $p < .0001$, $\eta^2 = .441$, $BF_{incl} > 10k$) of C/D gain on the configured resistance. Post-hoc analysis indicated that there exist many significant differences between C/D gains, demonstrating consistent and distinct configurations made by participants. In addition, we found a strong negative correlation ($r(351977) = -.68$, $p < .0001$) between C/D gain and configured physical resistance, meaning that lower C/D gains lead to higher physical resistances. This shows that participants were able to consistently assign a physical resistance to a C/D gain (**H3**) and that the effects of C/D gain manipulations were perceptually distinguishable for our participants (**H5**).

Rotational C/D gain trigger perceptual different levels of resistance




Figure 4.6: Participants adjusted physical resistance given the applied C/D gain. The orange line represents default resistance.

C/D	0.5	0.6	0.7	0.8	0.9	1.0
0.5	-	< .01	< .001	< .001	< .0001	< .001
0.6	137.10	-	.21	< .01	< .0001	< .001
0.7	2315.4	1.11	-	< .05	< .01	< .01
0.8	858.48	113.91	6.83	-	.349	< .05
0.9	3755.2	5719.7	105.47	0.36	-	.21
1.0	2139.1	1243.3	49.67	13.94	1.09	-

Figure 4.7: Bayes factors (BF) and p -values for comparing configured physical resistance between C/D gains.

4.1.4.4 Summary

Our experiment confirms that substantial differences between the real and the virtual rotations of the knob can be introduced without users noticing (**H1**). Participants interpreted the pseudo-haptic effect triggered by C/D gain manipulation as physical resistance (**H3**). The JND in terms of physical resistance stays within the unnoticeable range of C/D gain manipulations, enabling designers to increase the perceived resistance (**H4**) without users noticing that they are being manipulated. The range of tested C/D gains created distinguishable levels of *Pseudo-Haptic Resistance* and participants reliably assigned physical resistance to the presented C/D gains, with smaller C/D gains resulting in higher levels of perceived physical resistance (**H5**). However, we could not identify a shift in perceptual thresholds caused by higher visuo-haptic congruence (**H2**). In the following, we outline important aspects and lessons learned when incorporating *Pseudo-Haptic Resistance* in rotational interactions and discuss what could potentially break the illusion.

Undetectable at
different magnitudes
of perceived physical
resistance

RENDERING RESISTANCE FOR ROTATING KNOBS We found that selecting an appropriate base level of resistance, which is in line with participants' expectations, is crucial for our *Pseudo-Haptic Resistance* illusions to work. Notably, in the preliminary experiment, participants immediately commented on the missing base resistance, which we intentionally kept low (0.3 Nmm). We changed this for the main experiment, and not a single participant commented on this. Following the method of adjustments, the first significant change in participants' configured resistance occurs at a C/D gain of 0.8. This is in line with our JND results, showing that participants notice a change in resistance at mean = 0.81 compared to the baseline. However, it is unclear whether the C/D steps required to achieve perceptually different levels of resistance between two interactions follow a linear pattern [90]. With decreasing C/D gains, the amount of variance in the configured physical resistance became larger. This may be attributable to two things: (1) individual variances in participants' perceptual abilities

Perceptual effect of
resistance may
depend on base level

and (2) the fact that for some participants, the illusion broke when being exposed to C/D gains above their DT.

BREAKING THE ILLUSION We observed two different effects that occurred when the *Pseudo-Haptic Resistance* illusion broke. P8 stated, “my virtual hand moves slower than my real hand”, which is the effect of a C/D gain ≤ 1.0 . To our surprise, the participant did not change the default resistance of 1.1 Nmm to a C/D gain of 0.5, but selected 1.6 Nmm for C/D gain = 0.8. In the interview, they said that when the virtual hand moved slower they closed their eyes during the procedure, because it “was so confusing to compare the resistances, when the dials moved at different speeds”. In contrast, P18 reached the maximum physical resistance that could be configured with our device and stated, “Can’t I go higher with the resistance of the red dial?”, followed by “I would need more resistance to match the blue dial” (i.e., C/D gain = 0.5). Thus, the participant still ‘believed’ the pseudo-haptic effect but exhausted the physical scale. This happened all three times when the 0.5 C/D gain was presented to this participant. Please note that P4 and P18 were also identified as (the only) outliers, following the box plot method, and were therefore excluded from the analysis above.

Proprioceptive dominance occurs

Robustness of the illusion when reaching its limits

4.1.5 A Model of Pseudo-Haptic Resistance

Our goal is to provide a model for *Pseudo-Haptic Resistance* that can be applied to new and exciting VR applications. Thus, we demonstrate how our results can be applied. Since the provided resistance depends on the radius of the knob, we compute the torque that the DC motor can produce using the following equation $\tau = \mathbf{r} \times \mathbf{F}$. As a result, depending on the designer’s needs, they can derive the required size of the motor to produce the desired haptic resistance given the lever, or vice versa (see Figure 4.8). Data from the method-of-adjustments experiment was fitted following a forced fusion procedure as indicated in Samad et al. [288]. Figure 4.8 shows the experimental data for perceived torque according to the veridical values rendered by our device per C/D gain, and the orange fitting curve of the model ($R = 0.96377$). Additionally, the model in Equation 2 predicts the perceived torque in our VR knobs for a given C/D gain, which is parameterized with the variable CD. The resulting value unit is Newtons per millimeter (Nmm), representing the counterforce provided at knob rotation.

Relationship between rotational C/D gain and perceived resistance

$$\text{Torque(Nmm)} = \frac{6.684}{0.054 + (0.946 * CD)} \quad (2)$$

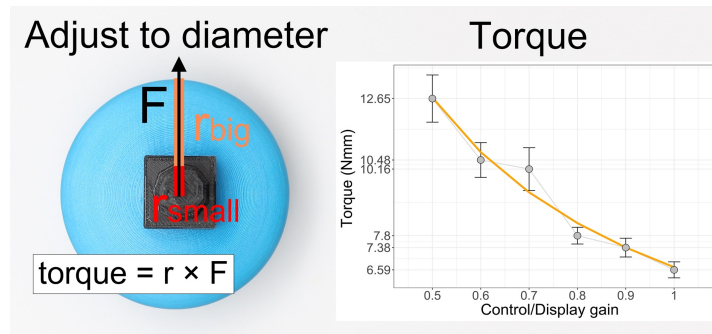


Figure 4.8: Left: Depiction of torque components on our 3D-printed knob. Right: Fitting for perceived torque by C/D gain.

4.1.6 Applications & Use Cases

To illustrate how VR designers can use our findings on rotational offsets and the introduced *Pseudo-Haptic Resistance*, we outline two potential use cases. We used the proxy knob hardware described above, and attached different 3D-printed objects using our custom mount.

4.1.6.1 Re-mapping 3D Interfaces

The ultimate goal of proxy-based haptics is that users only need a minimal set of physical props to simulate a large set of virtual objects. When also considering function, i.e., the different states that proxies can adopt, it becomes more challenging. To illustrate this, we implemented a DJ desk application where we redirect the virtual knobs on the DJ desk to a single physical proxy knob (see Figure 4.9). Depending on the selected virtual knob, the proxy knob resets itself to the corresponding state (e.g., the rotation limits or to align its haptic features). With only two functional proxies, slider (which we developed in the next section) and knob, we could, in fact, provide haptic feedback for the DJ desk's UI by seamlessly redirecting users [52], without them noticing (see Figure 4.9 B). This is possible because the

Redirected 3D user
interfaces



Figure 4.9: (A) displaying our knob proxy next to our slider proxy developed in the section below. (B) shows a user's hand is redirected. (C) indicates the start position, while (D) shows the position of the hand after rotation in the real and virtual world.




Figure 4.10: 3D-printed steering wheel (A) and door handle (B) connected to our prototype using the custom mount. *Pseudo-haptic Resistance* allows users to perceive changing levels of resistance during interaction. Red indicates displaced hand.

different hand poses, when interacting with the UI elements, do not interfere with HR [99]. Finally, interactions could even be improved by applying the *REACH+* technique [132].

4.1.6.2 *Pseudo-Haptic Resistance Beyond Knobs*

Our results may be used beyond knobs to create realistic haptic sensations without the need of large and expensive actuators. For example, other types of rotational manipulations could benefit from *Pseudo-Haptic Resistance*: for example, simulating different door handles and levers, or even bi-manual interactions with steering wheels. We also envision our technique at a smaller scale, e.g., single finger interactions with existing controllers. Often, they already host mechanical actuators to provide haptic feedback and their haptic resolution could be improved further. To showcase this, we implemented a set of these interactions, depicted in Figure 4.10.

4.1.7 *Discussion*

4.1.7.1 *Limitations of Pseudo-Haptic Resistance*

Our proposed *Pseudo-Haptic Resistance* technique can be easily integrated because there is no special hardware or knowledge required. However, as with any other vision-based approach, it only works when users directly look at the object or their hand during interaction. We also encountered this issue with one of our participants who closed their eyes in order to focus on the physical resistance. This poses some limits on the technique, as users sometimes solely rely on their proprioception, e.g., when interacting with devices while focusing on another primary task. This is especially true for experts familiar with an interface.

*Illusion requires
visual attention*

Conservative C/D
gain prevented
stronger effects

We included a condition with haptic tick marks in order to improve visuo-haptic integration, hypothesizing that it should increase DTs [80, 262] which would have enlarged the possible range of unnoticeable C/D gains for *Pseudo-Haptic Resistance*. However, our results suggest that it did not have a measurable effect on the DTs. Nevertheless, we believe it is an interesting area to explore in future work, since in this paper, we were focusing on the most conservative case, i.e., we told and showed participants the effect of C/D gain manipulation and their only task was to detect it [384]. VR experiences are much more complex and dynamic, including distracting factors or tasks that require more attention—likely extending the range of unnoticeable C/D gains [26, 83].

Properties of the
knob and the
interaction may
affect our results

In our preliminary experiment, pilots commented on the physical resistance of the knob being ‘too loose’. This suggests to us that there is a physical resistance, given the visualization, form and feel of the knob, that participants intuitively expect based on their everyday experience with similar interfaces in the real world. Hence, as with any other psychophysical experiment, our results may need to be adapted, depending on the VR application. To this end, it is unclear how different base levels of physical resistance might affect the presented *Pseudo-Haptic Resistance* model.

4.1.7.2 Utility of *Pseudo-Haptic Resistance*

Starting point for
VR designers

Our model of *Pseudo-Haptic Resistance* is a simplification and cannot replace a more comprehensive multisensory integration model such as proposed by Ernst and Banks [80]. However, it still provides designers with a validated range of perceived torque (i.e., from 6.683 to 12.684 Nmm according to our tested C/D gains), reasonable for common VR interactions with 3DUIs. Our technique is applicable to existing systems and controllers and can be used to adapt interfaces to different users, e.g., based on their individual ergonomic preferences, without changing the hardware. In this work, we were only interested in increasing the perceived resistance in order to improve the haptic resolution of new and existing devices. Still, we recommend that future work should investigate whether *Pseudo-Haptic Resistance* works bi-directionally, i.e., if C/D gains ≥ 1.0 result in less perceived physical resistance. Similarly, the method could be used not only for rotations but also for translations or even stretching. The latter could be achieved by gradually changing the C/D gain.

4.1.7.3 Extending the Approach

We are also interested in situations where proxies cannot be used, and hence, the overall multimodal feedback relies on wearable haptic feedback and/or pseudo-haptics: for example, mid-air knobs, where, simi-

lar to *Holitouch* [15], a holistic approach could be developed. Here, our pseudo-haptic technique provides the first puzzle pieces to achieve this, and we encourage the community to build upon it. Finally, we did not evaluate the presented use cases through a formal user study but left this for future work.

4.1.8 Conclusion & Contributions

This section contributes to **RQ2** by presenting a novel undetectable pseudo-haptic technique called *Pseudo-Haptic Resistance*, which grants the VR system control over the perceived resistance during proxy-based rotational interactions (**C2**). We implemented a physical proxy knob capable of sensing, i.e., knowing the position of the knob, and actuation, i.e., providing physical resistance at all times (**C1**). Our preliminary study with four participants suggested that rotational C/D gains of a knob may translate to perceived changes in physical resistance. Given our promising results, we performed a thorough investigation of our proposed technique using a three-staged psychophysical experiment with 20 participants. The results showed that we can introduce substantial offsets between real and virtual rotations of a proxy knob without users noticing. With this, we provide evidence in favor of the visual-dominance phenomenon by effectively exploiting the sensory dominance of vision over proprioception. To quantify this effect for this type of interaction, we report the previously unknown DTs (**C3**). Moreover, we demonstrate that the technique can effectively convey distinguishable levels of rotational resistance while staying within unnoticeable ranges of C/D gains (**C2**). Second, we found that higher visuo-haptic integration through our visual and haptic tick marks did not result in a measurable effect on the detectability of rotational hand offsets (**C3**).

Novel undetectable pseudo-haptic technique

Conveys varying levels of resistance

In addition, we provide a first multisensory integration model based on forced fusion that describes how C/D gains correspond to perceived physical resistances and outline how our results can be translated to other rotational manipulations by presenting two use cases and applications (**C5**). We believe that our results will help to overcome the current limitations of proxy-based VR systems. At the same time, our technique can easily be deployed in new and existing VR applications. To this end, we have contributed a thorough investigation of *Pseudo-Haptic Resistance* technique with a proxy knob. However, one central UI element is still missing, sliders, which would allow Lisa to experience haptically supported interfaces in VR.

Model of C/D gain to perceived resistance

Applications and use cases demonstrate its potential

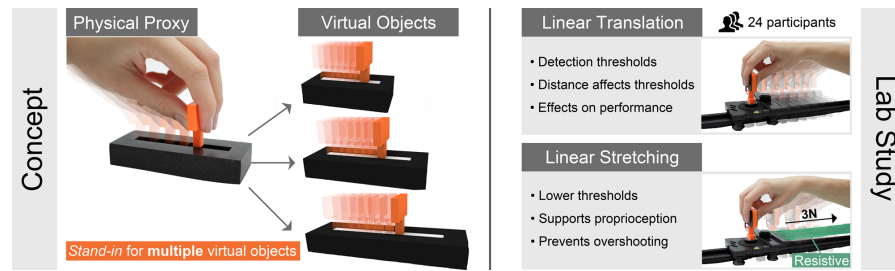


Figure 4.11: We create *Visuo-Haptic Translation* by using a single physical proxy slider that can represent multiple virtual sliders varying in their lengths. Our lab study reveals how much real-to-virtual offset can be introduced for linear translation and stretching while remaining unnoticed.

4.2 TOWARDS VISUO-HAPTIC TRANSLATION

4.2.1 Introduction

In this section, we focus on the missing fundamental object manipulation of linear translation (i.e., a 1D slider), which can be found in many domains, including traditional 2D UI elements such as scrollbars as well as physical interfaces (e.g., sliders and switches to control machines, vehicles, devices, and tools). Our central interest lies not in creating a pseudo-force but in understanding how much linear translation capability a physical proxy needs in order to function as a realistic stand-in for a broad spectrum of virtual objects with translatable parts. Specifically, we asked: “Does the physical slider need to be as long as the virtual slider?” If not, how much can this vary while remaining unnoticeable? (RQ2). In contrast to our previous investigation on the rotation of a knob at fixed travel distances, we also test sliders of varying lengths and bidirectional C/D gains. Through the latter, we hope to be able to design illusions that allow for simulating smaller and larger virtual sliders than physically available. Further, we explore linear stretching, where the object itself supplies force in the opposing direction of movement, to understand how this might affect users’ perception of the illusion and the process of multisensory integration.

4.2.2 Undetectable Visuo-Haptic Translation

Linear translation is a common object manipulation and plays an important role in our everyday life. For instance, humans can translate an entire object to a desired location, and many physical objects and interfaces offer translatable parts. Software GUIs also frequently make use of linear translation, scrollbars and sliders are used for input in video-, and photo-editing, design, office, programming, gaming and countless other softwares; physical sliders are often found to control machines, vehicles, devices and various other equipment.

Linear translation
and stretching




Figure 4.12: Left: Effect of using C/D ratio manipulations. Right: Force needed (Newton) to stretch the object across different distances and measured with a PCE-DFG Series force gauge.

Additionally, many tools make use of translational mechanical parts such as when measuring distances using calipers, opening and closing a containment or clothing through zippers, workshop tools such as screw clamps and mechanisms to extend a device, e.g., cutters, rulers, or even fishing rods to shrink their form factor. Linear stretching shares a lot of similarities with linear translation but adds relative force feedback to the translation. This force feedback can give users an understanding of the extent to which they manipulated the device, allowing them to better predict the resulting state or action. For instance, when using a slingshot, a user can estimate distance, trajectory, and velocity with which the object is going to fly based on the stretched distance and the perceived resistance force. Thus, resistive force actively contributes to the understanding of traveled distance.

Why is linear translation relevant?

4.2.2.1 C/D ratio

Similar to *Pseudo-Haptic Resistance*, we take advantage of the visual-dominance effect by manipulating the C/D ratio, with the goal of enabling functional proxy objects with translational parts to act as stand-ins for multiple virtual objects. Manipulating C/D ratios is well-known in the context of traditional 2D mouse interfaces [50] and can be easily adapted to VR [10, 75, 288]. By increasing the C/D ratio, the method scales up the performed physical interactions resulting in a larger virtual movement than the actual physical movement (see Figure 4.12 (C/D: 1.4)). Reducing the C/D ratio leads to the opposite effect: the real travel distance is further than the displayed virtual distance (see Figure 4.12 (C/D: 0.8)). Our goal was to study to what extent we can use this technique for linear translation and linear stretching without being detected by humans.

Technique behind our undetectable illusion

4.2.2.2 Form Factor & Interaction Technique

The various examples above illustrate the importance of linear translation as a widely used interaction method. Yet, the size of the object being manipulated, how it is handled during manipulation, and so on varies. Chapter 5 looks into potential effects on varying form factors and different types of interactions with a proxy. To control for these variables in this research project, we decided to focus on single-handed manipulations and a simple physical slider proxy that users can comfortably grasp and hold. The form factor was determined using pilot studies experimenting with different slider sizes and widths (see Figure 4.14). There are various ways to grasp and manipulate this simple slider [106]. The most common approach for this object size and width appears to be a pinching-type gesture using middle, index, and thumb fingers. We choose this interaction method based on pilot testing and previous work on illusions with proxies [30, 288]. Participants were asked to grasp the object using this pinching-type gesture, and this posture was replicated by a virtual hand (see Figure 4.14 and Figure 4.15). The experimenter monitored participants throughout the experiment to ensure that they maintained this grip.

*Rationals behind our
proxy setup*

4.2.3 Experiment

In this experiment, we study how much offset between the physical and the virtual representation can be introduced for linear translation and linear stretching without resulting in a semantic violation. We also investigate the effects that C/D ratio manipulations have on the interaction. We conducted a psychophysical threshold experiment to investigate the CDT of C/D ratio manipulations for both manipulation techniques, linear translation and linear stretching, and for two different travel distances, 7 and 14 cm. We chose 7 cm based on the travel distance offered by standard off-the-shelf slider potentiometers which are part of many UIs, and 14 cm to test twice the length. In the experiment, participants were seated on a chair while viewing a simple virtual environment (through an HMD) with their dominant hand being tracked. The virtual scene contained a table and the slider setup corresponding to the physical world. Participants used a thumb, index-middle finger pinch to grasp a virtual slider embodied by a functional physical proxy slider. They were told to translate the slider to a displayed position while being exposed to different C/D ratios repeatedly. Once they reached the target position, they were asked a 1AFC about whether they noticed a manipulation [322]. Specifically, after each manipulation, they were to respond either ‘yes’ or ‘no’ to the following statements: “*The virtual slider moved faster*” or “*The virtual slider moved slower*” depending on the condition, following Steinicke et al. [321]’s methodology. Participants were informed about the pro-

*Why did we choose
the distances*

*Using a 1AFC
question to assess
noticeability*




Figure 4.13: P#4's interleaved-staircase in the upper stretching 7 cm condition converging at CDT = 1.3.

cedure and had to report a manipulation as soon as they noticed it. The findings help to improve the understanding of how we can utilize the visual-dominance effect for linear translation and stretching across different distances aiming to support the design of rich physical proxies in VR.

4.2.3.1 Design

We used an adaptive psychophysical 1-up-1-down interleaved-staircase procedure with a 2×2 within-subjects study design. We had two independent variables: (1) manipulation techniques and (2) travel distances with two levels each. In total, we investigated four conditions. We measured three dependent variables: (1) participants' responses to the forced-choice question regarding the manipulation, (2) time needed to reach the target, and (3) movement profile for each trial. We used an interleaved-staircase exposing participants to different stimuli (C/D ratios) repeatedly. We chose a fixed step size 1-up-1-down design targeting the CDT. Here, either of the two possible responses is equally likely to occur (50% due to chance). Since the procedure can target different probabilities as described in Section 2.5.2.2, we can compute the required step-size (ψ_{target}) for the Δ^+ and Δ^- method and CDT = .5 as follows:

$$\psi_{\text{target}} = \frac{\Delta^+}{\Delta^+ + \Delta^-} \Rightarrow \frac{\Delta^-}{\Delta^+} = \frac{1 - \psi_{\text{target}}}{\psi_{\text{target}}} = \frac{(1 - .5)}{.5} = \frac{1}{1} \quad (3)$$

We used the number of reversal points ($r = 5$) in each sequence as a convergence criterion for the staircase procedure. For each condition, we utilized two separate staircase procedures to determine the upper (between 1.0 to 2.0) and lower (between 1.0 and 0.4) CDTs. Thus, the highest C/D ratio subjects could be exposed to was 2.0, and 0.4, respectively. These values were chosen based on previous findings in hand/finger-redirection [2, 83, 133, 384], however after pilot testing, we set the start values in the staircase procedure to upper (1.0–1.8) and lower (1.0–0.5) using a 0.1 fixed step size as this leads to quicker convergence [181]. We counterbalanced the ordering of the four conditions using Latin square and randomized the order for the upper and lower threshold procedures for each condition.

Psychophysical experiment to determine DT

Compute step-size for 50% target probability

Range of tested C/D gains

4.2.3.2 Participants

*Recruited
participants with
varying backgrounds*

We recruited 24 right-handed participants (eight females; 16 males) aged 22–34 (mean = 27.04; SD = 3.42) from the general public and the local university. This excludes one participant who was omitted from the analysis due to not reaching convergence in the study. This could have been due to system error or the participant not understanding the study which we could not determine in hindsight. Participants had a range of different educational and professional backgrounds, including media informatics, computer science, education, chemistry, pharmacy, biology, anglistics, economics and law. Five participants had never used VR before, twelve had used it a few times (1–5 times a year), three people used it often (6–10 times a year), and four other people on a regular basis (more than 10 times a year). Nine participants reported that they have not played VR games before, 14 people responded sometimes or infrequently (1–5 times a year), and one person on a regular basis (more than 10 times a year). Participants not associated with our institution received €10 as compensation for taking part in the experiment. The study was approved by the Saarland University’s Ethical Review and DFKI’s Hygiene Board.

4.2.3.3 Apparatus

*Hardware and
software setup*

In our study, we used the apparatus shown in Figure 4.14, consisting of an HTC VIVE tracking system (2PR8100); HMD, base stations, a VIVE controller and a VIVE tracker with SteamVR (v.1.13.10) and the OpenVR SDK (v.1.12.5). The virtual scene was developed with Unity3D (v.2019.2.17f) representing a small virtual world including the slider setup on a table running on an MSI P65 Creator 8RE with an Intel Core i7-8750H CPU, 16GB RAM and a NVIDIA GeForce GTX 1060. For hand tracking, we used a Leap Motion controller (SDK v.3.2) attached to the HMD to avoid augmenting the user’s hand. We used a simple hand representation in the VR world to avoid distractions




Figure 4.14: Study setup showing the camera slider hosting the 3D-printed slider with a conductive coating on both sides of the slider. The resistive band and custom 3D-printed quick-release mount are shown on the right.




Figure 4.15: Study task: Participant waits at the start position (A-a) and stops at the goal position (B-b) followed by the forced-choice question asking participants whether they noticed a manipulation

caused by, e.g., textures. The custom slider rig was built using an 80 cm camera slider, enabling a smooth translation movement across its length and forcing participants to translate the object in this direction. We 3D-printed a physical proxy slider and a custom mount for the VIVE tracker to ensure a fixed position while interacting with the device. Following initial testing, we chose a slider width of 1 cm, which users could comfortably grasp and hold throughout the experiment. To avoid hand-tracking issues, we snapped the virtual hand to the slider as soon as the participant grasped and held on to the physical proxy slider. To support this, we included capacitive touch sensing capability to the physical proxy slider using conductive 3D printing filament (composite PLA – Electrically Conductive Graphite Ø1.75 mm) on either side of the slider connected via wires to an Arduino Nano 3.x running capacitive touch sensing firmware. Once the participant grasped the proxy slider, the microcontroller sends a touch event to the VR machine using serial port communication. For the linear stretching conditions, we used a resistive rehabilitation band providing 3.05 N (7 cm) and 5.02 N (14 cm) resistance (see Figure 4.12: right C/D ratio: 1.0). We measured the resistance before and after the study and could not find any difference potentially caused by material fatigue. The components were secured on the standard camera mount, which also had a plug mechanism to quickly de-/attach the resistance band in between conditions. We carefully calibrated the setup for each participant.

Capacitive touch sensing to improve visuo-haptic synchronization

Physical properties of the slider setup

4.2.3.4 Experimental Protocol

The study was conducted in a quiet room to avoid distraction and ensure the same testing conditions. After a study introduction, informed consent and explaining the hygiene measurements in place, participants filled in a demographics questionnaire. Then, they were introduced to the virtual environment and the task. They were informed about the procedure and the goal of the study. Following this, they performed an open-ended practice round until they were familiar with the system and the task. Participants were instructed to hold onto the slider throughout one staircase round, followed by a break. They were not permitted to grasp the slider differently or repeat a trial. Once they reached the goal position, the 1AFC question ('yes'

Instructions to participants

*Controlling for
movement speed*

or ‘no’) about whether they noticed a manipulation appeared. Participants were instructed to answer as quickly as possible. The slider had to stay at the goal position within a 5 mm threshold for the question to remain visible. Participants responded using the VIVE controller in their non-dominant hand by pointing with a laser, before bringing the slider back to the start position and repeating the procedure. In our pilot studies, we observed that movement speed may be another crucial variable when determining DTs (aligned with Hall and McCloskey [141]); therefore, we controlled for movement speed to isolate the effects of force and distance (position and movement). In our protocol, we instructed participants to translate the object with a consistent ‘normal’ speed in the warm-up task. Later, participants were informed when they moved the object too slow/fast, outside of the initially established time frame. After completing the four conditions, we asked them to fill in a SSQ in VR [177]. The total experiment took about one hour.

4.2.3.5 Data Collection

We collected data from five sources: a pre-study questionnaire for demographic information; the subjective responses to the forced-choice staircase question; system logs (including trial times, traveled distance, velocity and acceleration at a sample rate of 5 ms, about 100.000 data points), field notes and observations, and a post-study SSQ in VR using our *VRQuestionnaireToolkit* [94].

4.2.3.6 Hypothesis

In addition to determining the CDTs, we had the following two hypotheses for this experiment:

*We formulated the
following hypothesis*

H1: Manipulation distance has a significant effect on the DTs. Previous work investigating DTs indicated a potential effect of the scale of movement [2, 384]. We hypothesize that this effect should be evoked in our experiment—increased manipulation distance leads to smaller DTs.

H2: We hypothesize that linear stretching has lower DTs since the added relative resistance force provides an additional kinesthetic cue, which may support proprioception. Due to the proportional relationship between travel distance and resistive force, we give users an additional proprioceptive channel, which may result in an earlier semantic violation and potentially lead to quicker detection of manipulation.

4.2.4 Results

We report our estimates for the CDT for linear translation and linear stretching at the two travel distances 7 cm and 14 cm. We then analyze the CDT with respect to our hypothesis. Finally, we take a deeper look at the effects that different C/D ratios have on the overall performance and naturalism of the interaction. This helps to better understand the potential effects that visuo-haptic illusions for linear manipulations might have.

4.2.4.1 Detection Thresholds for Linear Translation and Linear Stretching

Overall, we collected 4080 responses as a result of our interleaved staircase procedure. On average, participants completed 170 trials ($SD = 8.20$). Each participant contributed one CDT per condition which was computed as the average of the last four reversal points out of each staircase sequences [181]. To determine the overall CDTs, we computed the means across the 24 CDTs for all eight conditions separately. The result can be found in Table 2 and Figure 4.16. For further analysis, we plotted participants' responses to see how they converged (for example, see Figure 4.13). All 192 staircase plots are available in the supplementary materials of the paper. Results from the SSQ questionnaire suggest that the haptic illusion did not trigger significant motion sickness. The Total Severity (TS) score was $mean = 18.23$, $SD = 10.54$.

Our estimates for CDTs

4.2.4.2 Effect of Manipulation Techniques and Travel Distance

To further analyze our collected data, we performed a Two-Way ANOVA on the two independent variables, manipulation techniques and travel distance with two levels each. The data was split into two groups, upper or lower threshold for analysis. We used Levene's test to check the homogeneity of variance and Shapiro-Wilk test to verify a normal distribution. Both threshold data sets, upper and lower meet the ANOVA assumptions at $\alpha = .05$. The Two-Way ANOVA

Verifying test assumptions

Table 2: CDTs indicate the C/D ratio disparity that can be used without being detected by a user. SD here indicates the variation across participants.

Condition	Upper CDT		Lower CDT	
	Mean	SD	Mean	SD
Translation_7 cm	1.62	0.18	0.70	0.09
Translation_14 cm	1.50	0.20	0.76	0.11
Stretching_7 cm	1.54	0.16	0.75	0.12
Stretching_14 cm	1.42	0.12	0.80	0.08




Figure 4.16: Violin plots visualizing the probability density of the collected CDT for each condition.

revealed a significant difference for travel distance, 7 cm and 14 cm for the lower ($F_{(1,92)} = 6.245$, $p = .014$) and upper ($F_{(1,92)} = 10.845$, $p = .001$) thresholds. This supports our hypothesis that travel distance ultimately determines how much offset can be introduced while remaining unnoticed. This was also supported by various post-study comments, for instance: “I just had much more time to see if it moved faster” (P8). This provides evidence for our first hypothesis that translation distance significantly affects the DTs. Next, we found the manipulation technique to also have a significant effect on both, lower ($F_{(1,92)} = 4.753$, $p = .031$) and upper ($F_{(1,92)} = 4.548$, $p = .035$) thresholds. This supports our second hypothesis that adding relative proprioceptive cues and, therefore, makes it easier to detect, resulting in lower thresholds. There was no interaction effect between manipulation techniques and travel distance for lower ($F_{(1,92)} = .086$, $p = .076$) and upper ($F_{(1,92)} = .004$, $p = .949$) threshold.

Manipulation technique and travel distance significantly affect CDTs. Our analysis showed that linear stretching has significantly lower thresholds than linear translation. Further, we determined that smaller travel distances allow for higher C/D ratios regardless of the manipulation technique.

4.2.4.3 Proprioceptive Sensitivity

Even though the violin plots in Figure 4.16 support the assumption of a normal distributed data set (verified through Shapiro-Wilk test), it is inevitable that the individual thresholds can differ quite drastically. Most participants fluctuated around the threshold; however, we were also interested in the extremes, and we wanted to understand whether some participants’ proprioceptive senses are more sensitive than others for these kinds of tasks [167]. Therefore, we used an extreme groups approach [273], allowing us to conceptually compare our population as if we had sampled high and low sensitivity groups. To do so, we computed participants’ proprioceptive acuity as an overall performance score by adding up all upper thresholds and the

Distance affects the amount of unnoticeable offset

Stretching results in lower thresholds

Individual differences in CDT




Figure 4.17: Visualizing the correlation between high- (blue) and low-(red) performance groups. 'x' indicates overall CDT.

multiplicative inverse for all lower thresholds across the eight conditions for each individual participant. We then assigned participants to quartiles (groups of six) based on their overall performance score. Thus, we ended up with a high and low-performance group as well as two average groups. We were mostly interested in whether participants in the high-performance group performed consistently better than the main group or underly a random spread, and vice versa. Figure 4.17 shows the low- (red) and the high- (blue) performance group. Connected data points in the high- and low-performance groups represent an individual participant. Here, we see a strong tendency towards individuals in the high-performance group consistently performing better and constantly staying below the overall CDT (marked with 'x'), and vice versa. Following this, we investigate if there is a significant difference between the two performance groups by comparing them to the average group. Levene's test revealed a violation of the homogeneity of variance assumption at $\alpha = .05$. Therefore, we ran a Kruskal-Wallis-Test for unequal variances which did not indicate a significant effect ($H(2) = 1.343$, $p = .510$). Participants in both groups reported mixed experiences with VR and had various backgrounds, countering our initial assumption that with more VR experience thresholds might become lower. Hence, we conclude that low- and high-performance groups belong to a single homogenous group.

Individuals did not significantly differ in their proprioceptive sensitivity in our study. We found a trend indicating individual differences in proprioceptive sensitivity (similar to [144, 167]) leading to an earlier semantic violation regardless of participants' prior VR experience and their professional background. However, this effect was not statistically significant in our experiment.

Extreme groups approach

Participants have consistently high or low thresholds




Figure 4.18: Trial completion times. Comparing two C/D ratios within the individual CDTs and the two extremes (0.5 and 1.8) against the baseline (C/D ratio = 1.0). *** = $p < .001$, ** = $p < .01$, * = $p < .05$, ns = not significant.

4.2.4.4 Interaction Times and Movement

The question we address here is whether C/D ratios within the CDT have acceptable performance and accuracy. We contrast this with performance and accuracy on C/D ratios outside of the CDT. These results help us to understand which factors contribute to semantic violations and, therefore, how different C/D ratios affect proxy-based interactions across the different conditions. First, we look at the trial completion times for all conditions at different C/D ratios by choosing the two initial staircase C/D ratios (0.5 and 1.8) as extremes, two C/D ratios within the threshold in each individual condition and the baseline C/D ratio = 1.0. We analyzed the trial time data at five C/D ratios depending on the conditions. Shapiro-Wilk indicated a violation of the normality assumption at $\alpha = .05$.

Following this, we ran non-parametric Kruskal-Wallis-Test which revealed a significant effect in all four conditions, Translation 7 ($H(4) = 97.24$, $p < .001$), Stretching 7 ($H(4) = 88.45$, $p < .001$), Translation 14 ($H(4) = 40.00$, $p < .001$) and Stretching 14 ($H(4) = 65.17$, $p < .001$). We performed Wilcoxon post-hoc comparisons with Bonferroni corrections for each condition, comparing baseline 1.0 to the four other C/D ratios. As shown in Figure 4.18, it took participants significantly longer to move the slider to the goal position using a 0.5 C/D ratio. In contrast, for each condition's lower threshold, there was no significant difference between the threshold and the baseline in terms of trial completion time. However, in the upper conditions, we did not see such consistent results. In the 7 cm conditions, participants were significantly faster within their CDT. Contrary, there was no significant difference in the 14 cm condition. At the extremes (C/D ratio: 1.8), participants reached significantly faster completion times in the translation condition. However, there was no effect in the stretching condition. Throughout the study, we frequently observed that participants overshoot the goal as a consequence of exposing them to higher C/D ratios. Moreover, some participants reported that they used this effect as an indicator to detect a manipulation. We followed up on this observation by investigating the movement

Potential side effects
of applying illusions

C/D ratio
significantly affects
the interaction time




Figure 4.19: Virtual distance traveled – time relationship graph for all conditions. Comparing two C/D ratios within the individual CDT and the two extremes (0.5 and 1.8) against the baseline (C/D ratio = 1.0).

data. To do so, we plotted the virtual distance–time relationship graph for again five different C/D ratios consisting of the two extremes (0.5 and 1.8), upper and lower CDT for each condition as well as the baseline (1.0). The graph is based on the logged timestamps and the corresponding median distance value across all participants. In Figure 4.19, we can observe that the different C/D ratios show a substantial effect on the movement data. First, the CDT curves (red and yellow) are closest to the baseline and result in a consistent, stable movement, i.e., accelerating at the beginning and slowing down when approaching the goal position. Aligned with our observations, the median curve for C/D ratio = 1.8 (purple) shows that participants frequently overshoot and had to correct for it, especially in the translation condition. However, stretching seems to prevent overshooting, even with higher C/D ratios leading to high accuracy. In contrast, the curve in lower extrema condition (blue) slowly approaches the goal and in the stretching 14 cm condition almost ends-up in a linear motion (see Figure 4.19: bottom right).

Performance and accuracy remain stable within the CDTs. Our analysis showed that higher C/D ratios generally result in quicker completion times but create problems with accuracy. However, staying within the CDT seems to prevent these effects from occurring.

4.2.4.5 Summary

In this section, we reported our estimates for the CDTs for all four conditions. Further, we demonstrated the significant effect of travel distance on the thresholds, smaller distances allow for higher C/D ratios. Moreover, we found that added relative resistance force feedback in the linear stretching condition provides an additional pro-

*Hand movements
affected by offsets*




Figure 4.20: Lisa translates the ‘high-pass filter’ slider (A – s2) through the physical proxy and then wants to operate the ‘track-speed’ slider (B – s1). The proxy slider resets itself to the previously stored position of s1 (see B - bottom Log). For the proxy to work as a stand-in for multiple virtual sliders, an on-the-fly HR technique is used (a and b). Finally, (1) shows a user interacting with the VR-DJ environment.

prioceptive cue, which supports people in detecting manipulations, consequently leading to significantly lower thresholds. Then, we ran an investigation regarding the possible differences in human proprioceptive capabilities, identifying a trend towards some people being more sensitive to visual-proprioceptive conflicts than others; however, this effect was only numerical present and not statistically significant in our experiment. Finally, we show that keeping C/D ratios within the CDT generally preserves performance and accuracy.

4.2.5 Applications & Use Cases

4.2.5.1 Linear Translation: Extending the VR-DJ Experience

To illustrate how *Visuo-Haptic Translation* can be used, we extended our VR-DJ experience. The main goal of this application was to demonstrate the possibilities that visuo-haptic illusions offer by replicating several virtual sliders (Figure 4.20B – s1, s2 and s3) of different lengths on a DJ desk through a single functional proxy slider. Here, Lisa can enter a virtual discotheque in order to practice her performance without the need of expensive mixer equipment². Various effects can be triggered, e.g., a fog machine can be triggered by using a switch. The single physical slider stands in for every slider in the virtual DJ mixer. Here, we use a 10 cm linear potentiometer (RSAoN11M9-LIN10k) offering 10 k Ω resolution which can be moved by using a 10V DC-motor. The DC-motor was controlled using an Arduino Uno Rev3 (AVR ATmega328) and a L298N Dual H Bridge motor driver powered by a 12V (1.0 A) DC power supply. When interacting with the device, the motor needs to be turned off. Therefore, we included capacitive touch sensing capability to the proxy by 3D printing a conductive slider mount (600 Ω resistance) using composite PLA – Electrically Conductive Graphite \varnothing 1.75 mm.

The sensor states and touch events were transmitted via serial communication to the VR machine and adequately mapped to the different slider lengths of 12 cm ($s_1 = C/D$ ratio: 1.2), 10 cm ($s_2 = C/D$ ratio: 1.0) and 8 cm ($s_3 = C/D$ ratio: 0.8) remaining within the CDT. On the VR side, we use a similar implementation as in our user study utilizing an HTC VIVE for VR and Leap motion for hand-tracking. Figure 4.20 illustrates how the system works. A user manipulates the 'high-pass filter' slider (Figure 4.20A – s_2) to a position and releases the slider. The system stores the slider state (analog potentiometer signal) for s_2 in a position log. Then, the user wants to manipulate the 'track-speed' slider (Figure 4.20B – s_1) which is still at the default position. The system fetches the current state of s_1 from the position log and resets the physical proxy slider accordingly (Figure 4.20B - bottom). For the proxy slider to function as a stand-in for multiple virtual sliders, we used Cheng et al. [52]'s *on-the-fly* HR algorithm. Hence, a user touches a different virtual slider but, in fact, has been redirected to the same physical proxy (Figure 4.20a and b). Finally, she starts to manipulate s_1 which updates the position log simultaneously (Figure 4.20b - bottom).

Supported interactions with the sliders on the virtual DJ desk

4.2.5.2 Linear Stretching

There is huge potential in applying our findings in the context of linear stretching. For instance, toolkits such as *TanGi* [91] and *VirtualBricks* [16] allow the creation of proxy objects including stretchable parts. Our findings help to expand their interaction space, enabling multipurpose manipulable proxies. Moreover, several haptic VR controllers utilize resistive forces for example, *ElasticVR* [347], *Haptic Links* [327] and *ElastiLinks* [358]. They support a rich set of interactions which can further be enhanced by using visuo-haptic illusions.

The benefits for stretching interactions

4.2.6 Discussion

Based on our study, we discuss visuo-haptic illusions for linear translation and stretching in VR. Finally, we identify potential future directions enabling the design of more generic proxy objects.

4.2.6.1 Proprioceptive Limits

Through our study, we determined the CDTs for linear translation and stretching at the travel distances 7 and 14 cm. We found that in some cases, relatively high C/D ratios remained undetected; thus, it appears that individuals' proprioceptive capabilities vary. We assume these differences are grounded in more fundamental human experiences. For instance, some people participate in sports

Wide range of undetectable offsets

² Unity assetstore: <https://tinyurl.com/mr26a34j>. Last accessed: Nov 1, 2024

[167], arts and crafts, or playing musical instruments—perhaps these activities select for or enhance people’s proprioception skills since these activities demand moving limbs out of binocular vision—which is not necessarily connected to people’s prior experience with VR. Investigating these differences appears to be an interesting and valuable direction for future VR research.

*Participants seem to
rely more on their
proprioceptive sense*

The difference between translation and stretching shows that not only the absolute limb position and potential differences to the virtual position can lead to a semantic violation. Instead, the entire arm chain, muscle contractions and moreover, the required force to manipulate an object may contribute to a semantic violation. This appears to be an unconscious process since none of our participants reported that it was easier to detect a C/D manipulation in the stretching conditions. When designing generic proxy objects, these findings need to be carefully considered to create a compelling VR experience. Addressing the question to what extent these individual differences might impact the practical feasibility of using visuo-haptic illusion is discussed in the section below.

4.2.6.2 Role of Movement Speed

*Faster hand
movements may be
easier to detect*

In our study, we controlled for movement speed, which appears to be another important variable contributing to a semantic violation. Our work provides supporting evidence for the role of position and force feedback as relevant proprioceptive factors; however, at this point, we cannot disentangle the effects of movement, speed, and force as well as possible correlations between them. Therefore, we propose this as an important direction for future work, allowing us to better understand how closely physical proxy and virtual objects need to match in order to enable the design of truly multipurpose proxies.

4.2.6.3 Practical Feasibility

*Study setup let to
smaller CDTs*

Our reported estimates are the result of investigating the most CDTs by informing participants about the procedure, reducing distractions to a minimum level, and converging at 0.5 probably (CDT) for a correct answer in the staircase procedure. In a more realistic VR experience, users are exposed to other distracting factors such as ambient sounds, incident light, multiple objects and so on. Thus, being immersed in the virtual environment most likely allows for higher manipulation factors while remaining unnoticed or at least are not experienced as disruptive [52, 83]. This is also supported by the SSQ results which did not indicate any significant motion sickness because of the illusion, even though we exposed participants to C/D ratios above their DT. There is also an interesting trade-off regarding the interaction speed and accuracy. For instance, lower C/D ratios result in longer and con-

stantly more stable interactions. In contrast, a designer likely wants to expand the virtual interaction space (e.g., let a virtual slider appear longer than its real counterpart is) by applying a higher C/D ratio. Here, a manipulation factor within the CDT could prevent overshooting, which impacts accuracy. Designers should be aware that using visuo-haptic illusions may affect accuracy and overall performance.

4.2.6.4 Generalizability for Proxies Design

In this first iteration, we abstracted from different grasping types and form factors. However, this raises the question: *to what extent can these thresholds be applied to a larger set of translatable objects and interactions?* As already discussed, in this work, we look at the most conservative case. Therefore, we expect these thresholds to work for different grasping types and form factors. We explore these aspects in the next Chapter 5 as researchers begin continuing to expand the scope of such proxies in future VR systems. Moreover, visuo-haptic illusions offer great potential to enhance proxy-based design by incorporating richer object characteristics such as bendable, twistable, and deformable object parts [91, 148, 239]. By doing so, we allow for re-usable, multipurpose, and realistic proxy objects pushing towards tangible VR. Finally, going beyond proxy design, this research also poses an interesting question to the HR approach, as it shows that force feedback is a proprioceptive factor; to what extent does lifting and holding an object affect HR.

Form factors and grasping gestures

Visuo-haptic illusions for other types of object manipulation

4.2.7 Conclusion & Contributions

In this section, we primarily contribute to **RQ2** by presenting a novel undetectable visuo-haptic technique, *Visuo-Haptic Translation*, which allows a single physical proxy slider to act as a stand-in for multiple virtual sliders of different length (**C2**). We base our technique on the findings from a controlled lab study involving 24 participants. Our work identifies CDTs for both linear translation and linear stretching (**C3**). The reported CDT help VR designers to incorporate haptic feedback into their designs with slider proxies. For instance, a VR designer interested in supporting novice DJs might build a virtual DJ desk without the need for expensive equipment; however, suppose the designer only has a single physical slider available. Because the designer cannot have a proxy for every virtual slider (and their lengths), she would employ visuo-haptic illusions alongside a single physical slider. Our findings would inform her the limits of what she can simulate without users noticing (**C2**).

On the quest to multipurpose proxies

By determining and investigating the CDTs, we further enhance the understanding of multisensory integration [80] during visuo-proprioceptive conflicts. Due to the absence of a corresponding

What happens
during sensory
mismatches

visualization for linear stretching, participants immediately experienced a mismatch between the sensory modalities, which led them to rely more on their proprioceptive sense. As a result, we observed significantly lower CDTs than in the translation condition. Thus, force feedback along with visualization likely impacts detectability (C3). In addition, smaller translations allowed for significantly higher relative C/D gains. Although we did not find statistical evidence for significant differences in participants' sensitivity to visuo-proprioceptive offsets, we found that individuals' thresholds are generally consistently high or low, covering a wide range, similar to Hartfill et al. [144].

Unlocking the
potential of
multipurpose
proxies

To demonstrate our novel *Visuo-Haptic Translation* technique, we developed a proxy slider that is capable of sensing and actuation (C1). The proxy slider can store multiple virtual slider positions and can quickly reset itself depending on the user's virtual target. To achieve this, we combined our visuo-haptic technique with *Haptic Retargeting* [17] to seamlessly redirect users' hands to our proxy. Together with our virtual DJ environment, this creates an immersive experience that haptically supports interactions with sliders through a single multipurpose proxy (C5). Our findings also help to outline the effects of noticeable vs. unnoticeable visuo-haptic illusions on the interaction, which we later build upon in Chapter 6.

4.3 SUMMARY

In this chapter, we presented two novel undetectable hand-based illusion techniques, *Pseudo-Haptic Resistance* and *Visuo-Haptic Translation*, allowing us to expand the scope of proxy-based interactions in VR (RQ2). As a result, our techniques help to overcome the physical limitations of proxy design while not disrupting the immersive nature of the VR experience (C2).

Two novel
hand-based illusion
techniques

Changing perceived
kinesthetic feedback

We established our illusion techniques through two controlled lab studies with a total of 48 participants using a range of psychophysical methods. Through this, we ensured that (1) we can effectively trigger perceptual illusions and (2) users do not notice the manipulations of their interactions within the IVE. Both techniques exploit the visual-dominance phenomenon by offsetting visual from proprioceptive sensory information to create haptic effects. We quantified the undetectable magnitudes of these types of illusions by reporting estimates for the CDT (C3). In Section 4.1, we presented *Pseudo-Haptic Resistance* which allows the system to change the perceived resistance of a knob upon rotational interactions. This can be achieved by visually slowing down Lisa's real-world rotations of the knob. As a result, she needs to cover more distance to reach the target location, which translates

into the sensation of more resistance (**C2**). We presented a first model based on multisensory integration that describes the relationship between C/D gain and perceived resistance, which helps designers to effectively apply our results to new and existing VR experiences. In Section 4.2, we presented *Visuo-Haptic Translation*, enabling a single physical proxy slider to act as a stand-in for multiple virtual sliders of varying lengths. This can be achieved by visually scaling up or down Lisa's real-world translations of the slider. Consequently, the haptic limits of virtual sliders are adequately represented, tricking Lisa's perception into believing that the proxy and virtual slider match (**C2**).

*Model describing
C/D gain and
perceived resistance*

*Multipurpose
proxies through
dynamic re-mapping*

As part of our investigation, we developed two fully functional proxy prototypes (**C1**) and designed the corresponding interactions and IVEs. Our presented applications and use cases demonstrate the rich interaction space of our techniques and outline the possibilities for interaction design with proxies (**C5**). During our investigations, we found many factors that had a measurable effect on the detectability of our hand-based illusions. For instance, travel distance and adding a distinct haptic cue (i.e., linear stretching) (**C3**). Thus, we were questioning the robustness and generalizability of our techniques, allowing them to be applied to practical VR settings rather than controlled lab environments. To better understand these aspects, we used the next Chapter 5 to study how different factors related to the user, the interaction, and the proxy affect the detectability of hand-based illusions that apply visuo-proprioceptive offsets (**RQ3**).

*Effectiveness
demonstrated
through applications*

*What about
robustness and
generalizability?*

DETECTABILITY OF HAND-BASED ILLUSIONS

Videos about the work and prototypes presented in this chapter are available online and can be accessed through the QR codes. Images and parts of the text in this chapter, as well as the presented figures, tables, ideas, concepts, implementations, applications and uses cases, studies and experiments, results, discussions, and conclusions, have been published previously in:

[103] **Martin Feick**, André Zenner, Simon Seibert, Anthony Tang, and Antonio Krüger. *In Proceedings of ACM CHI 2024*. The Impact of Avatar Completeness on Embodiment and the Detectability of Hand Redirection in Virtual Reality.



Video

[99] **Martin Feick**, Kora Regitz, Anthony Tang, and Antonio Krüger. *In Proceedings of ACM CHI 2022*. Designing Visuo-Haptic Illusions with Proxies in Virtual Reality: Exploration of Grasp, Movement Trajectory and Object Mass.



Video

5.1 FACTORS INFLUENCING DETECTABILITY



Figure 5.1: Lisa moves a proxy whisk inside a proxy pot (left). *Visuo-Haptic Rotation* works by offsetting the virtual hand/tool from the real hand/tool. We can simulate virtual pots of different sizes, which provide realistic haptic sensations when stirring. Our studies investigate the extent to which we can use such illusions by exploring the potential effects of grasp, movement trajectory and object mass on the offset which can be introduced while remaining unnoticed by a user.

5.1.1 Introduction

In addition to our two novel hand-based illusion techniques for VR presented in the previous Chapter 4, researchers introduced several other illusion techniques that can enhance interactions with proxies

Illusions need to remain undetectable

Factors influencing the unnoticeable offset

Investigate grasping type, trajectory and object mass

such as *Redirected Touching* [193], *Pseudo-Haptic Weight* [288], *Pseudo-Haptic Stiffness* [360] or *Resized Grasping* [30], that all decouple visual information from proprioception. These approaches present an effective and inexpensive way to achieve a variety of haptic sensations with proxies. Nevertheless, these techniques are only effective as long as they do not noticeably disrupt the VR experience—e.g., if we apply a visual offset that is too large, it disrupts Lisa’s sense of presence [128, 265, 314]. Previously, we noticed factors that influence how much offset can unnoticeably be applied, e.g., movement distance (see Section 4.2). Ongoing work in the field identified factors such as movement direction [25], complexity of the task [83] or the visualization of the virtual hand within an IVE [262]. To contribute to this stream of research and promote the use of hand-based illusions outside lab settings, we designed two experiments investigating: (1) effects of grasping type and the manipulation trajectory, and (2) the effect of different grasping types and object mass (see Figure 5.4). To do so, we varied the offset between the physical proxy and virtual object position by applying different C/D ratios for simple manipulation tasks. Additionally, we propose an extension of our *Visuo-Haptic Translation* technique, called *Visuo-Haptic Rotation*.

5.1.2 Impact of Grasp, Object Mass and Motion

There are several variables which may contribute to a semantic violation. Here, we explore the effect of three such variables: grasping type, object mass, and movement trajectory. In the following, we outline our selection process and discuss why it is important to understand the impact of these variables to develop a generalizable design approach for visuo-haptic illusions.

Does how we hold an object affect how much offset may be introduced?

The role of grasp type

Using grasp types taxonomies

Humans choose the correct grasping type based on the underlying task requirements [45, 62, 105] and objects’ characteristics [106] (particularly, the shape of the object [62]). These variables remain entangled and can therefore only be considered holistically. Cutkosky [62]’s grasping taxonomy broadly distinguishes between power, (intermediate) and precision grasps. As the name already suggests, power grasps are used to manipulate heavier/larger objects or when dexterity is secondary. On the other hand, precision grasps are primarily for fine-grained manipulations. Hence, different muscle groups are involved when changing or adjusting the grasping type [325] which motivates the question of whether the grasping pose itself affects how much offset between the real and virtual world may be introduced. Below, we outline the selection process of the four grasping types we chose for our studies. We analyzed several grasping taxonomies describing between 14 and 33 grasping types by comparing their sim-




Figure 5.2: The four grasping types: Lateral (A), Medium Wrap (B), Tripod (C) and Writing Tripod (D) – that we chose following the selection process described above. The objects were designed and 3D-printed according to the grasp requirements.

ilarities and differences [62, 107, 174, 325]. Our goal was to include one representative grasping type per established category across the different taxonomies, maximizing the likelihood of identifying potential differences. To do so, we prioritized the grasping types in each category based on their usage frequency in four different application areas: housekeeping, machinery, food preparation, and laundry [106, 389]. We selected the four grasping types for our study according to their usage frequency, kinematic differences [325, 370], and distinct object characteristics (i.e., size and mass) to obtain a diverse set of grasping types and to increase variability. This also aligns with the proposed optimal grasp set by Feix et al. [106]. We designed the corresponding proxy objects (see Figure 5.2 and Table 3) based on Feix’s grasp size analysis with real-world use cases in mind.

Procedure for selecting grasping types

Does how heavy the object is affect how much offset may be introduced?

Besides the shape and size of the object, another important variable is the mass of the object since the (predicted) mass of the object strongly correlates with the chosen grasping type [106]. It is unclear whether the properties of an object, such as its weight, contribute to a semantic violation when manipulating it. This question is grounded in the proprioceptive research field, where there is an interesting trade-off between accuracy and force, where movement accuracy is significantly affected by the force required for the manipulation [11, 274]. Following this, we included different object masses, up to 500 g, in our second study to investigate the impact on the DTs. In this work, we

Is lighter more accurate?

Table 3: Grasp classification, correlated object masses and dimensions, and examples for our four grasping types.

	Grasping Type			
	Lateral	Medium Wrap	Tripod	Writing Tripod
Grasp class [62, 325]	intermediate and flat	power and cylindrical	precision and spherical	precision and distal
Mass (avg.) [105, 106]	150 g	400 g	150 g	20 g
Dimensions [105, 106]	up to 2 cm thick	4.5 cm in diameter	3 cm in diameter	1 cm in diameter, tilt angle of 62.4° [74]
Examples [45, 105, 106, 389]	towels, keys, paper, mug handle, cards	bottles, cans, vacuum, mop, handles	doorknobs, salt/pepper shaker, chess pieces	drawing and writing tasks, kitchen and workshop tools




Figure 5.3: Visualization of how C/D ratio manipulations affect linear (*Visuo-Haptic Translation*) and circular movement (*Visuo-Haptic Rotation*).

consider handheld-sized objects, so we used object masses that correspond with our four grasping types and can usually be encountered in our everyday environments [106].

Does how we move the proxy affect how much offset may be introduced?

*Visuo-Haptic
Rotation*

Findings in the HR domain show that the movement direction and distance with respect to the user's body significantly affect the DTs [144]. For instance, distancing one's hand from one's body allows for much greater offset than vice versa (i.e., bringing one's hand closer to oneself) [25]. We included two distinct manipulation trajectories to explore if possible differences between grasping types occur for different movement directions. Rather than opting for the three main axes and limiting the DoF to only one, we used Lissajous-figures, which are used in the motor learning field [24], to systematically include more complex and rich interactions. We chose a one-to-one frequency ratio and 0° phase offset, and a one-to-one frequency ratio and 90° phase offset, resulting in a linear and circular movement (Figure 5.3). The furthest waypoint was set to 30 cm [384] to ensure that participants could physically reach it without fully extending their arm, which would provide a strong proprioceptive cue. To be able to compare the two movement trajectories, we used the distance point D in relation to a user's torso and mapped the C/D ratio intervals for physical movements and their corresponding virtual representation to one another. As a result, the only difference between the two trajectories is the total movement distance covered (i.e., circle perimeter $u = \pi * D > \text{linear distance } D$). We coined the circular extension of *Visuo-Haptic Translation*, *Visuo-Haptic Rotation*.

*Previously limited to
1D*

In line with our **RQ3**, we formulated the following sub-research questions:

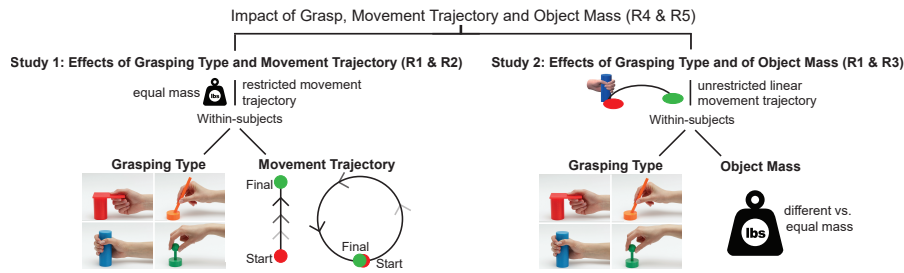


Figure 5.4: Experiment 1 investigates potential effects of grasping type and movement trajectory on the CDTs while neglecting an object's mass (left). Experiment 2 investigates the effects of grasping type and object mass (right).

RQ3.1: Does how a user holds the object affect how much offset may be introduced?

RQ3.2: Does the movement trajectory affect how much offset may be introduced?

RQ3.3: Does how heavy an object is, affect how much offset may be introduced?

RQ3.4: Does performing restricted vs. unrestricted movements affect proprioceptive accuracy?

RQ3.5: Do participants differ in their proprioceptive acuity?

Sub-research questions

To investigate these sub-research questions, we conducted two psychophysical threshold experiments to investigate the effects on the CDT [95, 385] for different independent variables outlined in experiments 1 and 2 below. Both experiments were executed at the same lab facility and used the same simple virtual environment consisting of two tables, the experimental setup, and an instruction screen. Participants remained seated on a chair throughout the experiment and were carefully positioned in front of the physical setup. Participants wore an HMD with their dominant hand being tracked. They were told to manipulate the proxy until it matched a target position displayed in the virtual world. After they successfully established the position, a 1AFC ('yes' or 'no') question appeared, and they were asked whether they noticed a manipulation or not [322]. In the linear movement condition (experiment 1 and 2), they responded to the following statement: "My virtual hand moved faster than my own". In contrast, in the circular condition (experiment 1) they responded to: "My virtual hand moved in a wider circle than my own". In both experiments, participants were informed about the procedure, and we explicitly showed them the effect of C/D ratio manipulations multiple times during the warm-up phase. They were told to report a manipulation as soon as they noticed it, thus targeting the most conservative case.

Conducted two psychophysical threshold experiments

Participants reported offset as soon as they noticed it

5.1.3 Experiment 1: Grasp Type and Trajectory

*Does how we hold
and move objects
influence the
detectability?*

In experiment 1, we compared the four grasping types (medium wrap, lateral, tripod and writing tripod) across two restricted movement trajectories (linear and circular manipulation). We used two different physical setups, enabling us to restrict a user's movement, preventing involuntary path deviations and neglecting an object's weight. This allowed us to isolate the effects that different grasping types and movement trajectories may have on the perception. Through this study, we wanted to understand whether we could use the same thresholds when using different grasping postures and manipulating proxies along different trajectories.

5.1.3.1 Design

We utilized an adaptive psychophysical 1-up-1-down interleaved staircase procedure with a 4×2 within-subjects design. We had two independent variables: GRASPING TYPE (lateral vs. medium wrap vs. tripod vs. writing tripod) \times MOVEMENT TRAJECTORY (linear vs. circular). In total we investigated eight conditions which were counterbalanced using a Latin square ($n = 8$). We used a 1-up-1-down interleaved staircase procedure, exposing participants to different stimuli (C/D ratios) repeatedly. Using a fixed step size, we target the (CDT) (see Section 2.5.2.2). We used the number of reversal points ($r = 5$) as a convergence criterion, and chose 1.0 and 2.0 for our range of manipulation factors with a 0.1 fixed step size [83, 384]. Following our pilot tests, we selected 1.0 (\uparrow asc.) and 1.8 (\downarrow desc.) as the starting values for the procedure to allow for quicker convergence.

5.1.3.2 Participants

*VR experts and
novices participated
in our experiment*

We recruited 24 right-handed participants (11 females, 13 males), aged 20–36 (mean = 26.42; SD = 3.65) from the general public and the local university. Participants had a range of different educational and professional backgrounds, including media informatics, computer science, education, pharmacy, anglistics, neuroengineering, embedded systems, data science and artificial intelligence. All participants reported normal or corrected-to-normal vision and did not report any known health issues which might impair their perception or proprioception. Eight participants had never used VR before, ten had used it a few times (1–5 times a year), no one reported using it often (6–10 times a year), and six others used it on a regular basis (more than 10 times a year). Ten participants reported that they had not played VR games before, nine people responded sometimes or infrequently (1–5 times a year), one often (6–10 times a year), and four people on a regular basis (more than 10 times a




Figure 5.5: Experiment 1 apparatus. Left shows the turntable allowing us to limit users' movements to a circular trajectory, resulting in *Visuo-Haptic Rotation*. The camera slider on the right ensures a smooth linear movement, resulting in *Visuo-Haptic Translation*.

year). Participants not associated with our institution received €10 as remuneration for taking part in the experiment. The study was approved by the Saarland University's Ethics and Hygiene Board.

5.1.3.3 Apparatus

In experiment 1, we used the apparatus shown in Figure 5.5, consisting of an HTC VIVE Pro Eye tracking system and an Optitrack system with five Flex13 cameras. On the software side, we used SteamVR (v.1.17), OpenVR SDK (v.1.1.4) and Motive (v.2.3.0) for motion capturing and running a simple virtual scene, which was developed in Unity3D(v.2020.2.1f1) and was executed on an Acer Predator Orion 5000 PO5-615s offering an Intel® Core i9 10900k CPU, 32 GB RAM and an Nvidia® GeForce RTX 3080. To support the initial grasping phase, we included hand tracking through a Leap Motion controller (core v.4.5.0) using an androgynous hand representation without noticeable characteristics as suggested by Schwind et al. [300] to prevent unwanted effects [262]. We built two different physical setups allowing us to restrict users' movement. For the linear setup, our 80 cm camera slider from Section 4.2 was used, forcing participants to translate the proxy alongside its path, thus, not allowing any path deviations as in our previous chapter. Additionally, this mechanism enables us to ignore object mass. A custom mount was 3D-printed allowing us to quickly swap out the objects for the different study conditions. The circular setup makes use of a lazy Susan turntable (metal bearing) with a laser-cut wooden plate and a custom mount which: (1) could rotate around its center using a second bearing, and (2) hosted the magnetic mount for attaching the different study objects. The two setups were fixed on tables and therefore could not be accidentally moved by our participants. The four objects were 3D-printed using PLA, and included 3D-printed conductive parts (composite PLA – Electrically Conductive Graphite) to enable touch sensing. Following Tinguy et al. [343], we used a combination of optical tracking and capacitive sensors to improve the visuo-haptic synchro-

Used generic hand models to prevent unwanted effects

Included touch sensing to warrant visuo-haptic synchronization




Figure 5.6: Visualizes the effect of C/D ratio manipulations when rotating and translating a proxy along a fixed trajectory. Bottom left shows our four objects augmented with 3D-printed capacitive touch sensors. Colors adjusted for presentation clarity.

nization and immersion in VR. In addition, by snapping the virtual hand to the virtual object when physically touching the proxy, we could avoid hand tracking issues. For touch sensing we used an Arduino Uno running a capacitive touch sensing sketch—transmitting (no-)touch events to the Unity3D program through serial port communication. The experimental logic was implemented using the Unity Experiment Framework (UXF v.2.1.1) [42] and the Unity Staircase Procedure Toolkit [379].

5.1.3.4 Experimental Protocol

After a general introduction to the study, informed consent and explaining the hygiene measures in place, participants filled in the demographics questionnaire. Following this, they were introduced to VR, the system, and the task. Participants were guided through an open-ended practice round to familiarize themselves with the task and the system. In the second step, we exposed them to trials with and without manipulation factors to illustrate the effect and only proceeded once they felt confident in detecting a manipulation. Participants were instructed to grasp the proxy object as indicated and to maintain the pose through each round of the experiment. The experimenter ensured that participants did not change their grasping pose unintentionally. They were told to move the object to the target position with a consistent and comfortable speed. The system monitored that they stayed within a reasonable time limit. Once they reached the goal position, the 1AFC question appeared, and the object needed to stay within a 5 mm distance for the question to remain visible. Participants were instructed to respond to the question as quickly as possible by pointing to either ‘yes’ (there was a manipulation) or ‘no’ (there was no manipulation) using the VIVE controller in their non-dominant hand. In our pilot experiments, we observed that participants carried a bias from the previous staircase round to the next. To address this and to cope with proprioceptive fatigue [280], participants took a longer break after each staircase round (by removing

*Responded with
controller in
non-dominant hand*

their headset). Before starting a new round, participants were given five calibration trials with no manipulation factor, helping them to ‘re-calibrate’ themselves. After completing the eight conditions, participants filled in a SSQ [177]. The total experiment took about 60–70 min per participant.

5.1.3.5 Data Collection

We collected data from five sources: a pre-study questionnaire for demographic information; the subjective responses to the forced-choice staircase question; system logs (including trial times, object position and orientation, and velocity using UXF [42]); field notes and observations; and a poststudy SSQ using our *VRQuestionnaireToolkit* [94].

5.1.3.6 Analysis

We statistically analyzed our data using a Two-Way RM ANOVA on the two independent variables, movement trajectory with two levels and grasping type with four levels. First, we identified two significant outliers using the box plot method, which we removed from the dataset for the analysis step. The dataset met the normality assumptions at $\alpha = .05$, verified through a Shapiro-Wilk test. We checked the assumption of sphericity using Mauchly’s test and applied Greenhouse-Geisser corrections to the within-subject factor grasping type, because sphericity was violated. Additionally, we conducted a Bayesian ANOVA using the BayesFactors R package¹ with default priors (v.0.9.12–4.3). Effects are reported as the Bayes factor for the exclusion of a particular effect (BF_{excl}), calculated as the ratio between the likelihood of the data given the model with the effect vs. the next simpler model without that effect [178].

Applying mixed methods to get a holistic picture

5.1.4 Results

We report our estimates for the CDTs using different grasping types (lateral, writing tripod, medium wrap, and tripod) along two restricted movement trajectories (linear and circular). Then, we analyze the results with respect to our research questions.

5.1.4.1 Detection Thresholds for Grasping Types and Movement Trajectories

We collected 4346 responses through the interleaved-staircase procedure. On average, it took participants 22.6 (SD = 3.6) trials to reach convergence. For each participant, we obtained eight thresholds (i.e., one per condition) by averaging the last four reversal points within

¹ BayesFactor GitHub: <https://tinyurl.com/mr3v3jyp>. Last accessed: Nov 1, 2024




Figure 5.7: Experiment 1 violin plots and data points for the different conditions (left). CDTs reported as mean ± SD (right). Outliers included.

the ascending and descending staircase sequence. The overall thresholds for the eight study conditions were determined by computing the mean across all 24 individual threshold values [385]. The results can be found in Figure 5.7. All 192 staircase plots from experiment 1 are available in the supplementary materials of the paper. Our analysis from the SSQ responses shows an increased Total Severity (TS) score, mean = 21.04, SD = 12.02 (P10 and P21 SSQ data lost). We hypothesize that this was a result of participants wearing a medical mask under the headset (as a COVID-19 hygiene measure) which increased sweating and discomfort, according to participants' post-study comments.

Slight increase in simulation sickness

5.1.4.2 Research Questions RQ3.1 & RQ3.2.

The Two-Way Repeated Measures ANOVA did not reveal a main effect on both variables, manipulation trajectory ($F_{1,21} = 2.292$, $p = .145$) and grasping type ($F_{3,63} = .298$, $p = .827$), within our collected data. There was also no interaction effect ($F_{3,63} = 1.152$, $p = .335$). At this point, it is unclear whether there is (practically seen) no effect, and designers can use the same thresholds regardless of how people grasp and move the proxy object, or there might be an effect that we could not find due to insensitive data. Therefore, we computed Bayes factors, which for manipulation trajectory 0.918 (BF_{excl}) did not favor either hypothesis, and thus, indicates that the data is insensitive [73]. We conclude that more data would be needed to unravel this variable. On the other hand, the BF_{excl} for grasping type is 29.760, suggesting that it is 29.760 times more likely to observe this data under the null hypothesis. Hence, there exists very strong evidence that grasping type did not affect the CDTs [162]. For the interaction effect, we found moderate evidence for the null hypothesis ($BF_{\text{excl}} = 7.590$).

Factors did not affect the CDTs

Strong evidence for the absence of effect

5.1.4.3 Summary

The study showed that we can introduce substantial offsets that are undetectable by humans for all grasping types and across both movement trajectories. Our Two-Way Repeated Measures ANOVA could




Figure 5.8: Experiment 2 linear unrestricted movement task. Our four different objects were augmented with markers for tracking. We fabricated two objects of each kind, with equal weight and different weights. Colors adjusted for presentation clarity.

not reveal a main effect in our data, and the Bayesian analysis suggests that for movement trajectory, this is due to insensitive data. In contrast, for grasping type, the Bayesian analysis provides strong evidence for the null hypothesis, i.e., there exists a high likelihood that grasping type did not affect the CDTs. To this end, we restricted users' motion to isolate the potential effects of grasping type and movement trajectory on the CDTs. However, in our everyday environment, humans regularly perform unrestricted movements with objects, requiring them to lift the objects. Therefore, we conducted a second study to investigate the role of different grasping types and object masses during unrestricted movement.

5.1.5 Experiment 2: Grasp Type and Object Mass

In experiment 2, we compared four grasping types (lateral, medium wrap, tripod and writing tripod) and two mass conditions: (1) all objects had equal mass and (2) all objects had a range of different masses according to the grasping type. We did not restrict user movements in any way (i.e., we apply normal gain-based HR), which introduces some variance. This allows us to isolate the effects that different grasping types and object masses may have on the perception. Through this experiment, we wanted to understand whether these variables significantly contribute to a semantic violation and, therefore, require special consideration when designing hand-based illusions.

Do heavier objects impact with proprioceptive accuracy

5.1.5.1 Design

We utilized an adaptive psychophysical 1-up-1-down interleaved staircase procedure with a 4×2 within-subjects design. We had two independent variables: GRASPING TYPE (lateral vs. medium wrap vs. tripod vs. writing tripod) \times OBJECT MASS (equal vs. unequal mass). In total, we investigated eight conditions, which were counterbalanced using a Latin square ($n = 8$). We applied the same method as in experiment 1.

Experimental 2 design follows experiment 1

5.1.5.2 *Participants*

*Participants from
previous study were
not permitted*

We recruited a new set of 24 right-handed participants (nine females, 15 males), aged 20–37 (mean = 26.70; SD = 5.01) from the general public and the local university. This excludes two participants who were omitted from the analysis due to (1) not reaching convergence in the study and (2) a complete system failure. Participants had a range of different educational and professional backgrounds, including computer science, media informatics, electronics, pharmacy, bioinformatics, HCI, mathematics, and psychology. All participants reported normal or corrected-to-normal vision and did not report any known health issues which might impair their perception or proprioception. Two participants had never used VR before, sixteen had used it a few times (1–5 times a year), one person used VR often (6–10 times a year), and five others on a regular basis (more than 10 times a year). Nine participants reported that they had not played VR games before, eleven people responded sometimes or infrequently (1–5 times a year), one person used it often (6–10 times a year), and three people on a regular basis (more than 10 times a year). Participants not associated with our institution received €10 as remuneration for taking part in the experiment. The study was approved by the Saarland University's Ethics and Hygiene Board.

5.1.5.3 *Apparatus*

*Preparing our
proxies for
experiment 2*

In experiment 2, we used the same apparatus as in experiment 1, but we removed the turntable and the slider. Instead, participants manipulated the proxies directly on the table (see Figure 5.8). For experiment 2, we 3D-printed eight objects, two of each kind using PLA, and included 3D-printed conductive parts to enable touch sensing. The objects were connected via long thin cables to the Arduino Uno, not limiting the interaction space. The design of the lateral object was slightly altered to ensure a more natural manipulation (center of mass). The first set of objects was fabricated to have an equal weight of 40 g (± 1 g tolerance), whereas the second set was weighted (see Figure 5.8, right) using lead shot and secured with super glue (± 2 g tolerance). The lead shot was equally distributed and superglued inside the objects, providing a realistic center of mass to avoid immediate breaks in presence [382]. All objects were augmented with three retroreflective markers, allowing us to precisely track them in 3D space using Optitrack. To improve tracking quality and robustness, we added paper straws to enable an optimal marker setup. The pivots were carefully calibrated for each object.

5.1.5.4 *Experimental Protocol & Data Collection*

We used the same procedure and data collection method as in experiment 1. Participants were instructed to move the object in the most




Figure 5.9: Experiment 2 violin plots and data points for the different conditions (left). CDTs reported as mean ± SD (right).

direct way (linear) to the goal position without dragging it on the table, requiring them to slightly lift the object. This ensured that participants felt the mass of the object.

5.1.5.5 Analysis

We further analyzed our data using a Two-Way Repeated Measures ANOVA on the two independent variables grasping type and object mass. There were no extreme outliers in the dataset. A Shapiro–Wilk test indicated a violation of the normality assumption at $\alpha = .05$ in the lateral/equal mass condition. Hence, we examined the normal QQ plots (see supplementary materials of paper) and computed skewness $\hat{\gamma}_1$ and kurtosis $\hat{\gamma}_2$ values ($|\hat{\gamma}_1|$ and $|\hat{\gamma}_2| < 2.3$) leading to the conclusion that we can run parametric tests [110, 233]. The dataset met the assumption of sphericity verified through Mauchly’s test at $\alpha = .05$. As in experiment 1, we computed Bayes factors to further analyze our collected data.

Justified to conduct parametric statistical hypothesis testing

5.1.6 Results

We report our estimates for the CDTs using different grasping types (lateral, medium wrap, tripod, and writing tripod) with two mass conditions (equal and unequal mass). Then, we analyze the results with respect to our study questions.

5.1.6.1 Detection Thresholds for Grasping Types and Object Mass

Overall, we received 3950 responses in the interleaved-staircase procedure, and it took participants 21.5 (SD = 2.9) trials to reach convergence. As in experiment 1, each participant contributed eight thresholds, i.e., one per condition, which was determined by averaging the last four reversal points in each sequence. The overall thresholds for all eight study conditions were computed by taking the mean across all 24 individual threshold values. The results can be found in Figure 5.9. All 192 staircase plots from experiment 2 are available in the paper appendix. Similar to experiment 1, the analyses of the SSQ re-

Summary of collected results

sponses show an increase in the total severity score (mean = 27.43; SD = 22.56).

5.1.6.2 Research Questions RQ3.1 & RQ3.3.

*Confirms experiment
1 results for
unrestricted
movements*

The Two-Way Repeated Measures ANOVA did not reveal a main effect of either variable, object mass ($F_{1,21} = .371$, $p = .549$) or grasping type ($F_{3,63} = .430$, $p = .732$). In addition, there was also no interaction effect ($F_{3,63} = .757$, $p = .522$). Following this analysis, we computed Bayes factors, and for object mass we found moderate evidence ($BF_{\text{excl}} = 5.527$) in favor of the absence of an effect on the thresholds, i.e., no effect on the CDTs is 5.527 times more likely than that there was an effect. This does not contradict previous findings on the significant impact of force on proprioceptive accuracy. Ansems et al. [11] tested 10%, 25% and 40% of maximum voluntary contraction (MVC) force, which refers to the highest possible load an individual can move using a muscle (group). In fact, these values are substantially greater than the maximum proxy weight of 500 g in our experiment. In line with experiment 1, the Bayes factor for grasping type, 9.623 (BF_{excl}), provides moderate evidence for the null hypothesis (RQ3.1), and for the interaction effect model, strong evidence ($BF_{\text{excl}} = 13.112$) for the null hypothesis given the data.

*Object mass did not
affect the CDTs*

5.1.6.3 Summary

We reported our estimates for the CDTs to quantify the amount of unnoticeable offset. The Two-Way RM ANOVA did not show a main effect. However, the Bayesian analysis provides supporting evidence for accepting the null hypothesis on the variables grasping type and object mass—there exists a high likelihood that neither variable, grasping type or object masses (≤ 500 g), affected the amount of offset which can be introduced. Next, we analyze both studies with respect to our research questions RQ3.4 and RQ3.5.

5.1.7 Movement and Individual Differences

In this section, we analyze both studies, in total 48 participants contributing 384 thresholds, to investigate the differences between linear restricted vs. unrestricted movement type (RQ3.4). We observed a high threshold variance across participants, which led to the question of whether there are consistent differences in humans' proprioceptive acuity (RQ3.5). Finally, we analyze participants' backgrounds with respect to the determined thresholds to better understand where such differences may come from.

5.1.7.1 Restricted vs. Unrestricted Linear Movement

Here, we make the assumption that grasping type did not have an effect on the CDTs following the evidence obtained through our Bayesian analysis in both studies. We analyzed the between-subjects factor linear movement type (restricted vs. unrestricted movement). Since there are two levels in the unrestricted movement condition (equal weight and weighted), we ran two independent samples Welch's t-tests because the dataset did not meet the homogeneity of variance assumption verified through Levene's test. A Shapiro–Wilk test indicated a violation of the normality assumption at $\alpha = .05$. Given our sample size, we examined the normal QQ plots (see supplementary materials in paper) and computed skewness $\hat{\gamma}_1$ and kurtosis $\hat{\gamma}_2$ values (for all conditions $|\hat{\gamma}_1|$ and $|\hat{\gamma}_2| < 3$) leading to the conclusion that there is no severe violation of normality [110, 233]. Further, we applied Bonferroni corrections to account for Type I errors. Additionally, we performed two Bayesian independent samples t-tests using default effect size priors. Results are reported as two-tailed Bayes factors BF_{10} and effect size estimates as median posterior Cohen's δ with a 95% credibility interval (95%CI) [178].

The role of congruent haptic feedback

RESEARCH OBJECTIVE RQ3.4 Our analysis provides strong evidence for an increase in CDTs in the linear restricted (haptic) movement condition (Mdn = 1.45), when comparing to both unrestricted conditions: (1) linear weighted (Mdn = 1.33) ($t_{(180)} = 4.13$, $p < .001$, $d = 0.596$, $BF_{10} = 350.080$, with median posterior $\delta = 0.567$, 95%CI = [0.282, 0.855]), and (2) linear equal weight (Mdn = 1.34) ($t_{(186)} = -4.36$, $p < .001$, $d = -0.629$, $BF_{10} = 826.248$, with median posterior $\delta = -0.600$, 95%CI = [-0.890, -0.313]). These results suggest that providing congruent haptic feedback and limited participants' DoF reduces proprioceptive accuracy and thus allows for greater offsets (see Figure 5.10). This could have potentially two reasons: (1) due to the presence of continuous haptic feedback, the combined visuo-haptic perception was more robust because of higher weighting during mul-

Significant difference between restricted and unrestricted movements




Figure 5.10: Restricted (continuous haptic cue) vs. unrestricted (normal HR gain) linear movement violin plots and data points. *** = $p < .001$ (Bonferroni-adjusted).

*Congruent haptic
feedback may lead to
more robust sensory
integration*

tisensory integration, or (2) based on our observations, we believe that this could also be caused by the ‘somewhat’ artificial movement and the momentum that is generated when smoothly manipulating the object along a fixed trajectory. In contrast, unrestricted linear manipulations resemble a frequently occurring, highly trained and memorized interaction, which may lead to higher accuracy.

5.1.7.2 *Proprioceptive Differences*

Are we all the same?

As illustrated in Figures 5.8, 5.9 and 5.10, participants’ thresholds are widely spread across the entire testing spectrum, which leads to the question: *Are humans equally sensitive to visuo-haptic illusions?* We ran a descriptive analysis on our experiment 1 and experiment 2 data. Then, we computed an overall threshold for each participant by averaging their eight CDTs.

RESEARCH OBJECTIVE RQ3.5 The dataset follows a normal distribution, indicating that it is representative of the general population. Since there are no additional density humps, (which could suggest multiple performance groups) we conclude that all participants belong to the same population (see supplementary materials in paper). However, we were still surprised by how large the threshold spectrum is, reaching from 5% to almost 67% possible C/D gains. Therefore, we analyzed whether these differences are linked to participants’ backgrounds reported in the demographic questionnaires.

Analysis. We conducted multiple Spearman’s ρ rank correlations across all 48 participants, evaluating if there is a relationship between the CDTs and the ratings on the following questionnaire items: participating in physical sports activities, prior experience with VR, 3D interactions or VR gaming, gender, and age.

*Prior VR experience
may impact
tolerance*

Results. There was a positive correlation between the two variables threshold and prior VR experience ($\rho(46) = .15, p = .003$), indicating that with more VR experience, thresholds become smaller. There was no correlation between participants’ thresholds and physical activities ($\rho(46) = .03, p = .567$), experience with 3D interactions ($\rho(46) = .03, p = .502$), VR gaming ($\rho(46) = -.04, p = .385$), gender ($\rho(46) = -.03, p = .592$) or age ($\rho(46) = -.02, p = .682$).

5.1.7.3 *Summary*

We found that restricting a user’s movement results in significantly higher CDT. Additionally, we investigated where differences in proprioceptive accuracy may be linked to. It appears that one of the important factors is previous experience in VR, impacting how much offset can be introduced. In the next section we outline how our results support designers when incorporating visuo-haptic illusions into their workflow.




Figure 5.11: The virtual kitchen space and the questions that Lisa encounters (A, B, and C). She can use the same movement offset (=) regardless of the grasping type (A) and the weight of the different kitchen utensils (B). In addition, Lisa can overcome the mismatch between the physical proxy (document holder) and the virtual oven by applying a greater unnoticeable movement offset, resulting in a realistic haptic sensation (C).

5.1.8 Effects on Applications & Use Cases

Here, we demonstrate the impact of our key findings on the design process of visuo-haptic illusions in VR. For illustration purposes, we use a virtual cooking class scenario in which Lisa wants to experience haptic feedback for her interaction with the various virtual tools. We can achieve this by including our visuo-haptic illusions to enhance the experience (see Figure 5.11). Lisa is thinking about the possible interactions with objects and the environment as well as the physical proxies that are available, and encounters the challenges below:

Grasping Types (A). Lisa is designing the kitchen space consisting of a virtual stove and a single proxy which acts as a stand-in for all the available virtual pots (and maybe pans and bowls) in the kitchen. By using visuo-haptic illusions, she can simulate different-sized pots on the stove. Consider that there are several ways that kitchen utensils such as a whisk could be grasped. Here, she needs to be aware of whether the differences in how the tool is handled affect the extent to which an illusion can be used. Following our results, she would not have to restrict the interaction type in any way. Hence, Lisa can seamlessly transition between different grasping types, resulting in a natural and realistic experience.

Proxy Mass (B). Depending on the desired dish, the students need different kitchen utensils (e.g., a whisk, a hand mixer, or spoons). Clearly, each tool is suited for a specific use case; thus, they differ in their properties and dimensions. Given these uncertainties, Lisa needs to understand which attributes might limit or expand the detectability. Our results suggest that for handheld-sized objects (≤ 500 g), the amount of movement offset which can be introduced is not noticeably affected by the object's mass.

Restricted vs. Unrestricted Movement (C). Finally, we demonstrate how increased CDTs may help to enhance proxy-based interactions. For instance, a student is holding a virtual baking sheet embodied by a physical (document holder) proxy. The physical proxy does not

Illustrate interaction design process of VR illusions with proxies

Applying of Visuo-Haptic Rotation

Considerations of our findings in VR interaction design

perfectly match the depth of the virtual oven. By using a visuo-haptic illusion, Lisa can create a matching depth sensation. She is aware that the oven rails restrict a user's movement; this allows her to include a larger offset, which would otherwise not have been possible.

5.1.9 Discussion

Our work aims to untangle the contributing factors leading to a semantic violation when manipulating a virtual object embodied by a physical proxy. We believe that our results can be seen as a first step towards a generalizable approach for visuo-haptic illusion design. Here, we discuss our findings in a broader context, outline current limitations, and give recommendations for future work.

5.1.9.1 Role of Movement Speed

Future work should consider differences in hand movements velocity

Aligned with our previous studies in Chapter 4, we observed that movement speed seems to be a critical variable. To the best of our ability, we tried to control for it by (1) instructing participants to move the object with a consistent 'normal' speed, (2) giving them a warm-up round to establish a comfortable pacing and (3) monitoring their speed through our study program. As soon as participants moved the object faster or slower than the previously determined threshold boundaries, they were instructed (by using audio feedback) to adjust their speed. Future work should aim to investigate the role of movement speed in visuo-proprioceptive conflicts.

5.1.9.2 Restricted vs. Unrestricted Movement

Congruent haptic feedback impacts CDTs

The potential difference between restricting and not restricting the DoF of a user's motion has powerful implications for the design process of visuo-haptic illusions. In fact, many real-world objects, UI elements and mechanics limit or guide users' movements and provide continuous congruent haptic feedback, such as steering wheels, levers, switches, shifters, door handles, keyholes, sliders, knobs and many others. Capitalizing on this potential when designing VR illusions can help to create more engaging, realistic, and powerful applications. Moreover, there are several haptic devices that steer or guide a user's movement, such as *ElastiLinks* [358] and *Haptic Links* [327], which could greatly benefit from illusions. To this end, our results build on the assumption that different grasping types do not have an effect on the thresholds. Therefore, future work should focus on direct comparison between restricted vs. unrestricted movements and additionally, incorporate other trajectories.

5.1.9.3 Generalizability

It is important to note that we only investigated HR gains ≥ 1.0 , scaling up a user's real-world movement which is in practical terms more prevalent. Nevertheless, at this point, it is unclear whether scaling down (gain < 1.0) may reveal other results. Further, we chose our set of diverse grasping types and object masses based on the selection criteria described above, covering a wide spectrum to detect potential differences; however, there is still the possibility that other grasping types/object masses might influence the amount of offset that can be introduced. In addition, the role of movement trajectory needs further investigation because our experiment could not unpack it. Hence, future studies are needed to gather more evidence leading to a better understanding of the relevant variables in visuo-proprioceptive conflicts. Finally, an interesting question for future work is how the visualization (i.e., only showing a user's hand and not the entire arm chain, thus offering very limited visual cues to detect a manipulation) generally affects the CDTs.

Only investigated practically relevant effects

5.1.9.4 Personalized VR Experiences

We found that individual thresholds differ quite drastically. Our analysis suggests that this is linked to participants' previous experience in VR. Likewise, there are many more variables contributing to an individual's proprioceptive acuity, which we could not assess in our questionnaire [280]. Since the perceptual differences appear to be widely spread, we propose to investigate whether we can establish a method for proprioceptive calibration in Chapter 6. In fact, this could also be expanded to other illusion techniques such as redirected walking [322], haptic retargeting [17] or redirected touch [193]. This approach could first calibrate a conservative base threshold, and depending on other parameters such as complexity of the experience (distracting factors) [52, 83, 384] or time spent in the application [280], we could dynamically adjust the magnitude of the hand-base illusion, delivering a VR experience tailored to the individual.

Large individual differences in CDT

5.1.10 Conclusion & Contributions

In this section, we primarily contribute to **RQ3** by investigating factors related to the user, interaction, and the proxy itself, studying their potential impact on the detectability during HR-based redirection. The results provide a better understanding of the contribution variables leading to a semantic violation (**C3**).

To achieve this, we conducted two experiments with 48 participants, unravelling the extent to which three variables (grasping type, movement trajectory and object mass) impact the amount of

*Movement
trajectory, proxy
mass and grasp type
do not seem to have
practically relevant
effects*

*Continuous
congruent haptics
enlarges CDTs*

*Sensitivity to offsets
seem to depend on
prior experience*

*Introduced
Visuo-Haptic
Rotation*

*How to apply
visuo-haptic
illusions to VR*

offset which can be introduced while remaining unnoticeable for users. Our reported CDTs for the different conditions support the visual-dominance phenomenon for new types of interactions with proxies. Our frequentist analysis did not reveal a significant effect of the study variables grasping types, object masses (≤ 500 g) and the different movement conditions (linear and circular) on the amount of offset which can be introduced. The computed Bayes factors suggest that for movement trajectory, this was due to insensitive data. For grasping type and object mass, there was evidence for the absence of an effect on the CDTs, suggesting that existing hand-based illusions and their perceptual boundaries can be applied (C3). However, we found a significant effect between linear restricted and unrestricted movement. Restricted movement led to smaller CDTs, indicating that proxy manipulations that limit a user's motion along a fixed path, providing continuous congruent haptic feedback, may allow for greater offsets. We identified a wide range of thresholds linked to participants' prior experience in VR, suggesting that we need some sort of proprioceptive calibration process—pushing towards personalized VR experiences (C3 and RQ4). Together, these results strengthen the understanding of the factors that need to be considered when moving from lab to practical VR settings, contributing to the continuously growing number of studies and aiming to establish design guidelines for including undetectable illusions in VR.

Although not the main focus of this section, our investigation also produced a novel illusion technique, *Visuo-Haptic Rotation*, which is an extension of *Visuo-Haptic Translation*. It allows Lisa to, e.g., steer inside virtual pots of different sizes embodied by a single physical proxy. Thus, it contributes to illusion techniques that expand the haptic rendering capabilities of proxies (C2). Finally, we outline how to incorporate *Visuo-Haptic Translation* and *Visuo-Haptic Rotation*, interaction techniques and proxies studied in this section into a practical VR experiences (C5). Through this, we also provide considerations for VR designers on how to apply undetectable hand-based illusion techniques (C3).

5.2 EFFECTS OF AVATAR REPRESENTATION ON DETECTABILITY




Figure 5.12: Left: A user is performing a pointing task with HR in VR. Middle: shows the four avatar completeness levels Hand, Hand + Arm, FullBody and a FullBodyTexture. Right: Depicts the experimental setup, where the user (right) sees himself in the virtual mirror (left).

5.2.1 Introduction

Yet, previous studies (including ours) investigated HR thresholds using mid-air ‘floating’ hands that are disconnected from the user’s body. At the same time, however, there is a growing interest in full-body avatars for social and collaborative VR, games, simulations, etc. [115, 359] supported by more advanced body tracking in consumer hardware²³. This leads to the question: “Does HR still work if users are embodied in full-body avatars?”. There is good reason to expect HR to not function as properly given full-body avatars as when only visualizing virtual hands. When HR is working, we see our virtual hands as moving slower/faster away from our bodies compared to what our proprioceptive sense is telling us. Yet a more prominent and complete avatar could make it even more obvious that an unusual manipulation is going on. Here, we would visually perceive more of an arm instead of just a hand, and we might also see more clearly that our arm is not moving proportionately in the way we would expect. Further, our motivation to study this is grounded in previous work that found that users embodied by a full-body avatar perform better in estimating targets [77, 126], and that higher avatar embodiment generally improves spatial [281] and depth [270] perception in VR. Given that HR exploits users’ perceptual inaccuracies [47], it is possible that previous HR DTs, determined by only visualizing hands, are inflated because users experienced very low embodiment. As a consequence, many of the established hand-based illusion techniques and their perceptual thresholds could unexpectedly disrupt the VR experience if systems begin to use full-body avatars for embodiments. Thus, in accordance to **RQ3**, we formulate the following two sub-research questions:

Does HR rely on mid-air floating virtual hands?

² Meta webpage: <https://tinyurl.com/4te9nknv>. Last accessed: Nov 1, 2024

³ Steam webpage: <https://tinyurl.com/39wxtxhj>. Last accessed: Nov 1, 2024

*We ask the following
research questions*

RQ3.6: Does increased avatar completeness lead to a greater sense of embodiment?

RQ3.7: Does a greater sense of embodiment lead to smaller HR thresholds?

5.2.2 Experiment

We conducted an experiment in a quiet room, using a non-distracting virtual environment consisting of a chair, the experimental setup, and an instruction screen. Additionally, we used a virtual mirror to enable participants to see their virtual avatar in order to enhance their behavior responses [115]. Participants remained seated on a physical chair throughout the experiment.

5.2.2.1 Avatar Completeness Levels

*We defined four
avatar completeness
levels based on RW
and current
hardware tracking
capabilities*

Inspired by Eubanks et al. [84], we defined four levels of avatar completeness: (1) Hand, (2) Hand + Arm, (3) FullBody and (4) a FullBody Texture avatars. The avatar levels are depicted in Figure 5.12. We used an abstract and generic 3D avatar model for completeness level (1)–(3) to prevent unwanted effects caused by the visualization [13, 262] as, in this iteration, we were mostly interested in the effects of avatar completeness. Therefore, we aimed to reduce other influences by uniformly coloring these avatars white and designing them in a gender-neutral way [300]. To facilitate the feeling of embodiment, we developed a calibration tool for scaling the size of an avatar to match the user’s body, specifically, the hand, forearm and upper arm length as well as the height of the user. In addition, we allowed the calibration of hand width, forearm and upper arm diameter, leading to higher anthropomorphic fidelity [77]. By customizing the size of an avatar, we ensured that the virtual body parts correspond to the physical dimensions of the user’s body, promoting better distance and spatial perception [77, 217, 281]. Additionally, we investigate how a textured full-body avatar would affect the DTs compared to an abstract avatar. This decision was motivated by Ogawa et al. [262]’s results, showing that visually more realistic hands make it more difficult to notice HR. To investigate this, we used two fully rigged FullBodyTexture avatars from Microsoft’s RocketBox library⁴ [129] shown in Figure 5.12.

5.2.2.2 Method

DOES INCREASED AVATAR COMPLETENESS LEAD TO A GREATER SENSE OF EMBODIMENT? To answer **RQ3.6**, we used a relevant

⁴ Microsoft GitHub: <https://tinyurl.com/dsaaya6c>. Last accessed: Nov 1, 2024




Figure 5.13: Shows the ten body poses participants had to replicate for a total of 2 min. This was done to give participants sufficient time to explore and familiarize themselves with each avatar completeness level.

subset of Gonzalez-Franco and Peck [130]’s embodiment questionnaire for Body Ownership, Agency, Location and Appearance. Questions regarding Tactile Sensations and the Response to External Stimuli did not apply to our experiment. Participants were asked to replicate various body postures with the help of the mirror shown to them in the virtual environment (see Figure 5.13). This was done to give participants sufficient time to familiarize themselves with each avatar completeness level, which is crucial to develop the sense of embodiment towards a virtual avatar [172]. Participants adopted a total of ten postures and were required to hold each position for two seconds before the system displayed the next one. There was also a free exploration where participants could explore their avatar. Overall, the familiarization phase took 2 min.

Our investigation is twofold

DOES A GREATER SENSE OF EMBODIMENT LEAD TO SMALLER HR THRESHOLDS? To investigate the potential effects on the detectability of HR (RQ3.7), we conducted a psychophysical DT experiment using a 2IFC procedure, investigating the effects of different avatar completeness levels on the CDT. Participants were repeatedly required to move their hand from a start to a target location displayed in the virtual world (30 cm in front of them), indicated by a sphere. In each trial, participants performed two successive movements and were consequently asked to respond to the two-alternative forced choice statement “*The two movements felt the same*” with either ‘yes’ or ‘no’ using a Vive Controller in their non-dominant hand. During one of the movements, no HR was applied (= baseline), whereas during the other, we applied a certain magnitude of HR (= the stimulus of this trial). The order of baseline and stimuli was randomized. We used a 1-up/1-down interleaved staircase method, using a fixed step-size of $\Delta_{\text{step}} = 0.7$ cm. We chose the number of reversals = 5 as the termination criterion, and 0 cm for the starting stimulus of the ascending and 8.4 cm for the descending sequence, respectively. These values were determined based on the existing literature and pilot testing. The final CDT was computed by averaging the last four reversals in each sequence. Participants were informed about the procedure, and we explicitly showed them the effect of HR multiple times during the warm-up phase to ensure that they understood the effect. They were




Figure 5.14: Example of a 1-up/1-down interleaved staircase procedure, consisting of an ascending (red) and a descending (blue) sequence. Depending on a participant's previous response the procedure either increases or decreases the next stimulus within a sequence. A directional change within a sequence is called a reversal point.

told to report HR as soon as they noticed it, thus targeting the most conservative case.

5.2.2.3 Design

We used a within-subjects design and included four avatar completeness levels: Hand, Hand + Arm, FullBody, and the FullBodyTexture avatar. We fully counterbalanced the conditions using a Latin square ($n = 4$).

5.2.2.4 Participants

We recruited 24 right-handed participants (nine female, 15 male), aged 20–38 (mean = 26.42; SD = 3.65) from the general public and the local university. This excludes one participant (P2), in whose case we had to stop the experiment due to system failure. Participants had a range of different educational and professional backgrounds, including media informatics, computer science, education, pharmacy, anglistics, neuroengineering, embedded systems, data science and artificial intelligence. All participants reported normal or corrected-to-normal vision and did not report any known health issues which might impair their perception. Twelve participants had never used VR before, seven had used it a few times (1–5 times a year), no one reported using it often (6–10 times a year), and 5 others used it on a regular basis (more than 10 times a year). Participants not associated with our institution received €10 as remuneration for taking part in the experiment. The study was approved by the Saarland University's Ethics Board.

5.2.2.5 Apparatus

Our apparatus consisted of a HTC VIVE Pro Eye system, three Vive Trackers (v.3), a Vive Controller, two base stations, a chair and a splint to fix the position of the index finger. For tracking users' movements, we attached two Vive Trackers to users' feet and a Vive Tracker to their

dominant hand. In addition, users were holding a Vive Controller in their non-dominant hand that also acted as a tracker. In order to realistically animate the virtual avatar, we used FinalIK⁵; specifically, the *VR IK* algorithm that is optimized for human locomotion. Otherwise, inaccuracies may affect the user's motor behavior without their noticing [204], which Gonzalez-Franco et al. [127] describe as the *self-avatar follower effect*. The virtual environment was developed in Unity3D (v.2022.3.10f1) and was rendered on an Nvidia® GeForce RTX 3080. The experimental logic was implemented using the Unity Experiment Framework (UXF v.2.4.3) [42], the Unity Staircase Procedure Toolkit [379] and our *VRQuestionnaireToolkit* [94].

REDIRECTED MOVEMENTS Redirecting users' hand movements is most often done by gradually offsetting the virtual hand from its physical counterpart as users reach towards a target. Cheng et al. [52], Azmandian et al. [17] and Kohli [193], for example, presented algorithms to achieve this, where the virtual hand can be offset horizontally, vertically and in the depth axis. We only opted for gradual horizontal displacement [52] in this experiment, because it is a well-studied direction [384] and common in scenarios employing HR for haptic retargeting. The previous chapters of this dissertation did not investigate this type of mid-air HR with the absence of a haptic stimuli. Following the experimental design of Zenner et al. [385], we offset the virtual hand only to the right in our study. As a result, users need to compensate with their real hand moving towards the left in order to reach the virtual target (conceptually illustrated in Figure 5.12 left).

Horizontal HR, a simple type of hand-based illusion

5.2.2.6 Experimental Protocol

Participants were given a general introduction to the study. Next, we gathered their informed consent, and asked them to fill in a demographics questionnaire. Then, we guided them through our calibration routine to adjust the avatar's size to the participant's, and we showed them the tracking setup, before equipping them with the different sensors. We fixed participants' index fingers using a splint and carefully calibrated limb positions [384] to ensure high visuo-proprioceptive alignment. Next, the participants were introduced to VR and the task through an open-ended practice round. Participants chose the FullBodyTexture avatar they could best identify with for the study. Once everything was set up, we started with the familiarization phase, asking participants to replicate a series of body poses shown to them. Afterwards, they filled in the embodiment questionnaire inside the VR environment (RQ3.6). After completing this part for a specific avatar completeness level, we continued with the second phase, where participants ran through the DT experiment (RQ3.7). Af-

Participants followed this study procedure

⁵ Unity assetstore: <https://tinyurl.com/5n6v8jbf>. Last accessed: Nov 1, 2024

ter completing both phases for all four avatars, participants filled in a SSQ [177]. The total experiment took about 70 min per participant.

5.2.2.7 Data Collection

We collected data from six sources: a pre-study questionnaire for demographic information, the responses in the avatar embodiment questionnaire, the subjective responses to the forced choice questions, field notes and observations, a short post-study interview and a SSQ inside the IVE.

5.2.2.8 Analysis

First, we removed significant outliers using the box plot method. We statistically analyzed our data using Friedman tests, because Shapiro-Wilk and QQ plots indicated a violation of the normality assumption at $\alpha = .05$. Post-hoc pairwise tests were corrected using Bonferroni-Holm adjustments. To further investigate our data, we conducted a Bayesian analysis using the BayesFactors R package¹ with default priors (v.0.9.12+–4.4). Effects are reported as the Bayes factor for the inclusion of a particular effect (BF_{incl}), calculated as the ratio between the likelihood of the data given the model with the effect vs. without that effect [178]. Additionally, we performed paired Bayesian t-tests using default effect size priors. Results are reported as two-tailed Bayes factors BF_{10} [178].

*We apply
Frequentist and
Bayesian analysis
methods*

5.2.3 Results

First, we examine whether greater avatar completeness leads to an increase in embodiment (**RQ3.6**).

INCREASED AVATAR COMPLETENESS LEADS TO A GREATER SENSE OF EMBODIMENT. Analogously to Gonzalez-Franco and Peck [130] we computed an embodiment score (range: –3 to 3) based on our subset of questions for Body Ownership (Q1–Q5), Agency (Q6–Q9), Location (Q14) and Appearance (Q17–Q18) (we abbreviate the score here as OALA). We used the following Equation 4 to compute the OALA score:

*Determine users
perceived
embodiment*

$$OALA = \frac{\frac{Ownership}{5} \cdot 2 + \frac{Agency}{4} \cdot 2 + Location \cdot 2 + \frac{Appearance}{2}}{7} \quad (4)$$

The results are depicted in Figure 5.15. Friedman tests revealed a main effect ($\chi(3) = 13.946$, $p = .003$, $W = .194$, $BF_{incl} = 1148$). Post-comparisons showed a significantly lower embodiment for Hand (mean = 0.418, SD = 1.020) and Hand + Arm (mean = 0.489,

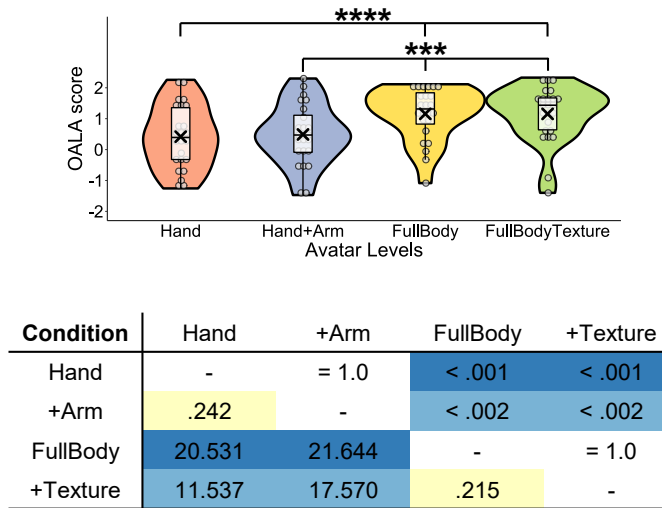


Figure 5.15: Top: OALA score across conditions, showing that participants perceived significantly higher levels of embodiment towards FullBody and FullBodyTexture than Hand, as well as between FullBody and FullBodyTexture vs. Hand + Arm avatars. Perceived embodiment appears to be unaffected when visualizing arms or textures. Bottom: Confusion matrix shows the supporting Bayesian factors (BF_{10}) and p -values.

SD = .983) in comparison to FullBody (mean = 1.157, SD = .849) and FullBodyTexture (mean = 1.155, SD = .917). There was no significant differences found between Hand and Hand + Arm ($p = 1.0$, $BF_{10} = .242$) as well as FullBody and FullBodyTexture ($p = 1.0$, $BF_{10} = .215$). The Bayesian analysis provides supporting evidence for the absence of an effect between the two conditions, suggesting that it is about four times more likely to observe this data under the null hypothesis. Thus, the FullBodyTexture avatar did not notably increase the level of embodiment felt towards the avatar. Participants had mixed feelings about this avatar, given the limited choice. For example, P8 stated “[I] could really identify myself” with the avatar, whereas P12 stated “The avatar does not look like me at all”. Finally, we identified a weak positive correlation between avatar completeness levels and the OALA score ($\rho(94) = .33$, $p = .001$, $BF_{10} = 1.470$), also visible in the distribution across conditions in Figure 5.17. As a result, we can confirm that the avatar completeness level generally increases the sense of embodiment (RQ3.6) which supports previous findings on avatar embodiment research [84, 124].

Significant lower perceived embodiment for hand and hand + arm

Participants subjective comments on the avatar completeness levels

GREATER SENSE OF EMBODIMENT DOES NOT SEEM TO IMPACT HR THRESHOLDS. All participants reached convergence in the staircase procedure (see Figure 5.14 for an example). All staircase plots are available in the paper’s supplementary materials. The CDT for the different conditions (Hand = 3.7 cm, Hand + Arm = 3.8 cm, FullBody = 4.0 cm, FullBodyTexture = 4.1 cm) are comparable to the thresholds reported in the literature for this type of HR [384].




Figure 5.16: Results from our psychophysical experiment suggest that our four avatar completeness levels do not affect the CDTs for HR.




Figure 5.17: Relationship between CDT and OALA, showing that higher embodiment does not seem to affect the CDTs. It can also be observed that data points for Hand and Hand + Arm tend to have lower embodiment, whereas FullBody and FullBodyTexture show higher embodiment.

*DTs remain robust
regardless of avatar
completeness*

There seems to be an increase in DTs with increasing avatar completeness level. To further investigate this, we statistically analyzed our data by running a non-parametric Friedman test, which did not reveal a main effect ($\chi(3) = .445$, $p = .931$, $W = .006$, $BF_{incl} = .134$). In fact, the Bayesian analysis provides evidence for the absence of an effect, suggesting that it is 7.46 times more likely that there is no effect that can be attributed to our four study conditions within our collected sample. This is supported by the violin plots depicted in Figure 5.16, revealing no substantial differences in the CDT between conditions. Contrasting the CDT with the OALA embodiment score did not suggest a tangible relationship between them ($\rho(94) = .042$, $p = .685$, $BF_{10} = .138$). The Bayesian analysis provided a 7.23 higher likelihood for the absence of an effect (Figure 5.17). Thus, we conclude that embodiment most likely did not influence the detectability of HR in our study (**RQ3.7**).

Finally, the SSQ results did not suggest simulator sickness caused by exposing participants to sometimes noticeable offsets above their DTs (Total Severity (TS) score: mean = 17.04, SD = 5.12) compared to our previous work and [2].

5.2.3.1 *Summary of Results*

Our study demonstrates that both full-body avatars lead to significantly higher levels of embodiment compared to Hand and Hand + Arm (**RQ3.6**). However, we found no difference between Hand and Hand + Arm as well as FullBody and FullBodyTexture avatars. The HR thresholds determined in our study confirm the results from previous studies. But our statistical analysis suggests that neither embodiment nor avatar completeness had measurable effects on the detectability of HR. Hence, we conclude that HR can be effectively used for a large set of illusion techniques, regardless of the avatar completeness level (**RQ3.7**). Our results contribute towards the trend of visualizing more complete avatars (e.g., in consumer applications), and suggest that designers and practitioners can apply existing illusion techniques and their perceptual thresholds even when using full-body representations of the user in VR

5.2.4 *Discussion*

5.2.4.1 *Impact on Avatar Embodiment*

Our results showed an increasing sense of embodiment felt towards full-body avatars. However, the absence of an effect between the FullBodyTexture and the FullBody avatar contradicts Ogawa et al. [262]’s results on more realistic hand representations. We believe that this could have been a result of participants having to choose between two very basic textured avatars with no room for customization in terms of appearance in our study. Thus, they had very different experiences under this condition, some participants identifying with the avatar, while others did not [29, 300]. As a result, one participant stated: “I think the robot avatar works better for me than the realistic one because...it feels more comfortable.” (P6). Despite this, the completeness level still appears to be the dominating factor, because both full-body avatars created a significantly greater sense of embodiment than hands or arms [84]. We recommend that future studies should allow for avatar customization or consider using a realistic full-body 3D scan of participants, which was out of the scope of this work.

Appearance of avatars did not match participants

Despite a small numerical increase in the CDT from Hand to Hand + Arm, our statistical analysis did not reveal a significant increase in embodiment between the two conditions. In fact, we provide evidence for the absence of an effect which supports the findings of Tran et al. [345]. We included Hand + Arm based on previous work [84, 109] and since it is commonly used by consumer systems such as *Leap Motion*⁶. However, our qualitative findings suggest

Extending the avatar completeness levels

⁶ Ultraleap webpage: <https://tinyurl.com/3ynfpn6m>. Last accessed: Nov 1, 2024

that participants still felt a stronger connection to the Hand + Arm, supported by participants' comments (such as those of P4: "*Arms and hands feel way more realistic and comfortable than only hands.*"), and the general correlation between higher avatar completeness levels and embodiment. Finally, future research should extend the avatar completeness levels, e.g., by including facial tracking, which might enhance the sense of embodiment even further.

5.2.4.2 Practical Implications for Hand Redirection

Ongoing research, including our in Chapter 4, identified many factors that can affect the detectability of HR [25, 83, 262]. Our estimates for the CDT are in line with previous work [384], confirming the validity of our experimental design. Moreover, they suggest that the most recent trend towards more complete avatars does not limit the practical use of hand-based VR illusions. Thus, VR designers could apply existing hand-based illusion techniques [17, 193, 288, 360] and their perceptual thresholds, regardless of the avatar representation, because we provide supporting evidence for the absence of an effect on the CDT through our Bayesian analysis. However, we would like to note that additional studies are required to confirm this. Also, the configuration of our psychophysical experiment (i.e., step size, range of tested offsets, number of steps, etc.) and the number of participants limit the kind of effects that we can detect reliably. For instance, there could be much smaller effects than we anticipated. However, we would like to note that such minimal effects—if existent—would then be merely of a theoretical nature and would have no practical implications for the use of HR. To investigate the effects of such magnitudes, we would also approach thresholds for which the Vive tracking system, as well as the IK heuristic, might not be accurate enough to arrive at reliable results. To this end, more advanced tracking systems, such as *Optitrack*⁷, could improve tracking accuracy and thus also reduce potential *Self-Follower Effects* [127]. Finally, it is again important to note that we investigated only a specific type of redirection, i.e., gradual, horizontal hand offsets towards the right. This type of redirection is commonly used in haptic retargeting applications [17, 132] and remapped interfaces [237]. Nevertheless, it remains to be explored in future studies if our results might be affected by the direction [25, 144] or the trajectory [385] of HR. This is important because other types of hand-based illusions [288, 360] make use of them.

5.2.4.3 Hand Redirection Beyond the Reachable

We only investigated CDT for HR in a '*reasonable*' area, i.e., the virtual hand position was always within reaching distance given the human body's biomechanics. However, HR can also be used for interacting

HR is not affected
my avatar
representation

Experiment was
designed to identify
large effects

Only one type of HR
was tested

⁷ Optitrack webpage: <https://www.optitrack.com/>. Last accessed: Nov 1, 2024

with distant objects in order to improve usability and ergonomics [272]. Here, out-of-reach interactions are often necessary, especially in the depth axis, which allows for the greatest offset between the real and virtual hand position to go unnoticed [25]. This can be studied easily with mid-air floating hands but becomes more challenging when using arms or full-body avatars. Feuchtner and Müller [109] presented an interaction technique for out-of-reach interactions by elongating a user's virtual arm. While preserving high body ownership over the virtual arm to improve usability, the visual manipulation can be clearly noticed. With a similar goal in mind, Dewez et al. [71] introduced dual body representations, a technique visualizing the real world as well as the visually offset position of users' arm movements to preserve high embodiment—while neglecting noticeability. *MultiSoma* [244] extends this by allowing a user to control up to four bodies at the same while maintaining a high sense of embodiment and agency. This raises the question of how unnoticeable redirected movements can be achieved with higher levels of avatar embodiment or even multiple synchronized avatars. For example, is it more reasonable to 'grow' the virtual hand, forearm, and upper arm uniformly, or should designers manipulate the dimensions of one limb to hide the offset, and if so, which one? Here, we still see the greatest risk for potential disruptions in the VR experience caused by the avatar completeness levels because there still exist no visualization techniques or algorithms for redirected movements that should remain unnoticeable for users.

*Beyond reach
interactions*

*Controlling avatars
with multiple limbs*

5.2.5 Conclusion & Contributions

In this section, we contribute to (RQ3) by studying the potential effects of more complete avatars on the sense of embodiment and the detectability of horizontal HR. To do so, we conducted a psychophysical experiment with 24 participants, investigating the four levels of avatar completeness: (1) Hand, (2) Hand + Arm, (3) FullBody and (4) FullBodyTexture. Our results help to understand the potential factors that designers need to consider when including hand-based illusions in VR to improve haptics or ergonomics (C3).

Participants experienced four avatar completeness levels by replicating a series of body postures to promote the onset of embodiment. Next, we assessed participants' sense of embodiment felt towards the avatars and found that they experienced significantly higher levels of embodiment towards more complete avatars. There were significant differences between Hand vs. FullBody and FullBodyTexture as well as Hand + Arm vs. FullBody and FullBodyTexture. As a result, we contribute supporting evidence for previous findings on avatar embodiment research [84, 124]. Our participants' ability to de-

*Sense of embodiment
increases with
higher completeness*

*CDTs for horizontal
HR remains
unaffected*

tect HR was comparable to that reported in existing literature [385], regardless of the avatar completeness level. Through this, we ensure the validity of our result and confirm the established perceptual boundaries for this type of HR. Moreover, we are the first to report CDTs for HR that does not use mid-air floating virtual hands. Our analysis provides evidence for the absence of practically relevant effects on the detectability of HR, suggesting that designers and practitioners can continue to apply existing hand-based illusion techniques with varying levels of avatar completeness. Our findings contribute towards a better understanding of the factors that limit unnoticeable HR-based illusion techniques that use visuo-proprioceptive offsets (C₃).

5.3 SUMMARY

In this chapter, we investigated how factors related to the user, interaction, and proxy affect the detectability of hand-based illusions (RQ₃). Our findings raise awareness of the factors that should be considered when designing VR experiences that utilize hand-based illusion techniques and their established perceptual thresholds (C₃).

To achieve this, we conducted three controlled lab experiments with a total of 72 participants using psychophysical methods to determine CDTs for new and existing types of interactions, with and without proxies, contributing to the landscape of perceptual studies concerning illusions that exploit visuo-proprioceptive offsets. Next, we analyzed the CDTs through a mix of frequentist and Bayesian statistics to infer a more holistic picture of the potential effects during multisensory integration.

*Congruent haptic
cues enlarge CDT*

*Set of factors that
may influence
detectability is long*

For example, we found that *Visuo-Haptic Translation* leads to an increase in CDT compared to free-form translation of a proxy, which may be attributed to stronger sensory weighting during integration, because of the presence of congruent haptic cues. Moreover, our results suggest that hand-based illusions seem to work regardless of the avatar representation, grasp type, and proxy mass, which is crucial for many VR applications. Overall, we saw that individual thresholds differ quite drastically, which could be linked to their previous experience in VR. Nevertheless, there most likely exist many more factors that contribute to an individual's proprioceptive sensitivity that we did not assess through our questionnaires, raising the need for an alternative approach. This chapter contributes to the ongoing string of research that attempts to untangle the dominating factors during visuo-proprioceptive conflicts with and without proxies [26, 133, 384].

We developed a design scenario using a cooking application that illustrates the process of how illusion techniques can be applied in

practical VR settings (C5). Through this application, we also highlight the relevance of the techniques studied. While this was not the primary focus of this chapter, as a part of this investigation, we also presented a novel hand-based illusion for VR, *Visuo-Haptic Rotation* (C2), enabling users to interact with virtual containers of varying sizes represented by a single physical proxy. This method employs a similar concept to *Visuo-Haptic Translation*, which was previously introduced in Section 4.2. Consequently, it extends the landscape of perceptually validated hand-based illusion techniques, and we quantify the limits of the techniques in the form of CDTs.

*Visuo-Haptic
Rotation in action*

During our investigations, we found many factors that had a measurable effect on the detectability of our illusions and others that did not seem to be practically relevant (C3). Together with ongoing work in the field, our results paint a complex picture between the type of interaction, individual's sensitivity, and the IVE on the amount of undetectable offset. Therefore, in the next chapter, we investigate an alternative method with the goal of allowing dynamic adaption of the undetectable offset independent of these factors (RQ4).

*Where do we go from
here?*

INDIVIDUALIZING HAND-BASED ILLUSIONS

Images and parts of the text in this chapter, as well as the presented figures, tables, ideas, concepts, implementations, applications and uses cases, studies and experiments, results, discussions, and conclusions, have been published previously in:

[98] **Martin Feick**, Kora Regitz, Lukas Gehrke, André Zenner, Anthony Tang, Tobias Jungbluth, Maurice Rekrut, and Antonio Krüger. *In Proceedings of ACM UIST 2024*. Predicting the Limits: Tailoring Unnoticeable HR Offsets in Virtual Reality to Individuals' Perceptual Boundaries.

[97] **Martin Feick**, Kora Regitz, Anthony Tang, Tobias Jungbluth, Maurice Rekrut, and Antonio Krüger. *In Proceedings of IEEE VR 2023*. Investigating Noticeable HR in Virtual Reality using Physiological and Interaction Data.

6.1 PHYSIOLOGICAL AND INTERACTION DATA DURING ILLUSIONS

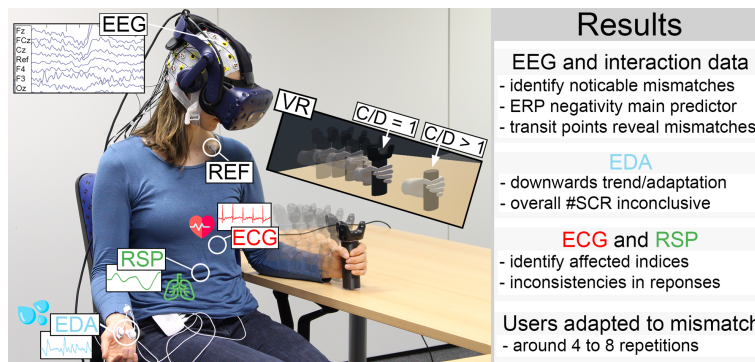


Figure 6.1: Study setup: A participant is moving a virtual object embodied by a physical proxy. The VR view shows the one-to-one mapping and the redirected position of the virtual hand/object.

6.1.1 Introduction

The extent to which hand-based illusions can be used without detection is a constantly growing field with the current state of the art of conducting psychophysical threshold experiments to determine humans' tolerance to offsets, reporting DTs that future work can

Psychophysical experiments are time-consuming and lack generalizability

build upon [2, 30, 133, 342, 372, 384]. However, this approach has three major limitations: (1) the psychophysical experiment needs to be conducted for each new device and illusion technique; (2) contextual factors limit the generalizability of the technique; and (3) the experiments do not account for individual differences, which is more meaningful compared to a group mean accordingly to Hartfill et al. [144] and other previous findings in Chapter 5. The general question remains, *how far can we push an illusion without an individual noticing it?*

Tailor illusions to users' current sensitivity to offsets

In this section, we consider individuals' perceptual boundaries of these illusions as a form of personalization. We explore the extent to which a system can automatically detect when these thresholds have been crossed for an individual—i.e., *how can we assess whether a semantic violation has occurred without a user telling us?* Our long-term vision is to use this information to automatically tailor illusions based on an individual's perceptual boundary in an *on-the-fly* fashion. Therefore, the goal of this section is to understand which sensor modalities can be reliably used to detect noticeable visuo-proprioceptive offsets by investigating EEG, ECG, EDA, RSP and interaction data. As part of this section, we also investigate how quickly participants adapt to such offsets and study how long this adaptation process takes.

6.1.2 Physiological & Interaction Data

In this section, we investigate if noninvasive physiological measures may show a unique signature that can be used to reliably detect gain-based HR that goes beyond individuals' perceptual boundaries. Additionally, we also incorporate measures that are linked to the interaction itself. Each of these measures has been used in prior related work, and we order them based on whether the measure was likely to provide a strong signal.

Brain shows response to noticeable offset

EEG: As discussed in Section 2.5.2.2, there exists a body of work which uses EEG, specifically ERPs, to detect errors in VR. Commonly, these studies report an effect in the frontal cortex area (FCz) between 100 to 360 ms after the stimulus onset [118, 265, 303, 306]. We hypothesize that noticeable gain-based HR shows a similar ERP pattern to errors in VR in the frontal cortex area, but in contrast to prior work, it is unclear when the effect occurs (e.g., beginning vs. middle vs. end of the movement phase).

Movement phases: Aimed movements can be separated into two distinct movement phases, ballistic and (an optional) correction [221]. It has been shown that introducing offsets between the real and the virtual world influences the execution of targeted movements because users need to compensate for the offset [17, 131, 207].

Therefore, we expect that noticeable gain-based HR also results in consistent shifts of movement phases, leading to a much longer correction phase when experiencing a noticeable hand offset because users would need to correct for the unexpected offset.

EDA: Facing noticeable gain-based HR may trigger physiological arousal and thus increases in the number of SCR peaks. We formulated this hypothesis based on findings reporting a correlation between greater self-reported VR immersion scores and lower physiological (EDA) responses [78].

Modalities we consider in our investigation

RSP: Performing precise interactions while coping with changing visuo-proprioceptive offsets is a challenging task. Respiration appears to be affected by various environmental stimuli in VR [86, 140]. Therefore, we hypothesize that changes in respiratory rate appear when encountering noticeable gain-based HR as a direct response to the unexpected mismatch.

ECG: Changes in heart rate or HRV measures can be observed when studying long-term immersion [231] and therefore may also be observed when exposing participants to a noticeable gain-based HR, which may disrupt presence, because of violated SCs [128, 310]. However, our work involves a comparatively short interaction, and it is unclear whether this would affect the rather slow-adapting HRV.

HRV measures typically respond slower

6.1.3 Experiment

We conducted an experiment in a quiet room with air conditioning to ensure a room temperature of 22–24°C, which is ideal for high-quality physiological data acquisition. We used a non-distracting virtual environment consisting of two tables, the experimental setup, and an instruction screen. Participants remained seated on a chair throughout the experiment. They wore an HMD, an EEG headset and various sensors on their body. They were told to move the virtual object (embodied by a physical proxy) forward until it matched a target location displayed in the virtual world. After they successfully established the position, they were required to maintain the position for two seconds before moving the object back to the start location. We explicitly showed participants the effect of C/D gain manipulations multiple times during the warm-up phase to ensure that they understood the effect.

Creating a controlled environment

6.1.3.1 Research Questions

Based on our (RQ4), we formulated the following five sub-research research:

RQ4.1: Are EDA, RSP and ECG responses triggered when experiencing a noticeable gain-based HR illusion?

RQ4.2: Can interaction data reveal a noticeable gain-based HR?

RQ4.3: Does a noticeable gain-based HR illusion show a distinguishable ERP?

RQ4.4: When in the interaction does a semantic violation happen?

RQ4.5: Do humans adapt to noticeable gain-based HR over time, and if so, how long does it take for them to adapt to these offsets?

6.1.3.2 Design

In this experiment we use a within-subjects design. We had four study conditions: two Baselines, a Steady and a Mixed, each consisting of 16 trials (see Figure 6.2).

Collecting baseline
data

Baseline: Uses a one-to-one mapping between real and virtual movement corresponding to a C/D gain of 1.0. These conditions are used to collect ground-truth data [118] about what participants experience to be ‘normal’, which needs to be captured in VR [78]. By including two baseline conditions, we are able to perform consistency checks on our collected sample.

Investigating
adaption

Steady: Applies a fixed C/D gain that lies above an individual’s perceptual boundary, i.e., they can detect the offset. To establish this per-participant gain, we use a pre-calibration described below. This condition enables us to address (RQ4.5) because participants may adapt to the pre-calibrated threshold, i.e. they do not notice the manipulation anymore, even though they did initially.

Exposing
participants to
sudden noticeable
changes

Mixed: Randomly jumps between two C/D gains (1-to-1 mapping from Baseline, and the calibrated offset used in Steady condition). The system ensures an equal occurrence of both C/D gains and only allows for a maximum of three consecutive trials with the same C/D gain. With this, we want to provoke a situation where a participant fails to adapt because of the repeatedly changing C/D gains.




Figure 6.2: Study design. Even participant#: AB, odd participant#: BA.

Each of the Steady and Mixed conditions was paired with a preceding Baseline condition, respectively. This was done to ensure that participants could establish a stable model of the environment, i.e. re-calibrate themselves, before experiencing a condition with visuo-proprioceptive offsets. To counterbalance the study, we alternated the order of the two blocks A and B for each participant, resulting in a factor two design (see Figure 6.2).

6.1.3.3 Participants

We recruited 22 right-handed participants (five females, 17 males), aged 18–31 (mean = 25.05; SD = 3.05) from the general public and the local university. Participants had a range of different educational and professional backgrounds, including media informatics, computer science, education, pharmacy, cybersecurity, entrepreneurship, biomedical engineering, data science and artificial intelligence. All participants reported normal or corrected-to-normal vision and did not report any known health issues which might impair their perception or proprioception. Nine participants had never used VR before, ten had used it a few times (1–5 times a year), no one reported using it often (6–10 times a year), and three others used it on a regular basis (more than 10 times a year). Participants not associated with our institution received €15 as remuneration for taking part in the experiment. The study was approved by the Saarland University's Ethics and Hygiene Board.

No one reported health issues that impair perception or proprioception

6.1.3.4 Apparatus

We used the apparatus consisting of an HTC VIVE Pro tracking system with SteamVR (v.1.22) and OpenVR SDK (v.1.16.8). The simple virtual scene was developed in Unity3D (v.2022.1.0). We used an Acer Predator Orion 5000 PO5-615s offering an Intel® Core i9 10900k CPU, 32 GB RAM and an Nvidia® GeForce RTX 3080 for running the experiment. We included an androgynous hand representation without noticeable characteristics as suggested by Schwind et al. [300] to prevent unwanted effects [262]. To avoid mismatches due to error-prone hand tracking [303], the virtual hand was affixed to the virtual object, and initially aligned with a participant's real hand. The experimental logic was implemented using the Unity Experiment Framework (UXF v.2.1.1) [42] and the Unity Staircase Procedure Toolkit [379].

We extended the basic system from our previous DT experiments to ensure comparability

EEG SETUP EEG data was recorded from 32 actively amplified electrodes using BrainAmp DC amplifiers from BrainProducts¹. Electrodes were placed according to the international 10–20 system with

¹ Brainproducts webpage: <https://www.brainproducts.com/>. Last accessed: Nov 1, 2024

nasion/inion as reference points (see Figure 2.31 in RW). We applied conductive gel to establish a proper connection between electrodes and scalp, and brought impedance of all active electrodes below 20 kOhm, before continuing in the experiment. Impedance was verified before and after the study. EEG data was recorded with a sampling rate of 500 Hz. The validity of EEG recordings with an HMD has been verified previously [149].

*We used the
Biosignalsplux
physiological sensors*

BIOSIGNAL SETUP To collect ECG, RSP, and EDA data, we used Biosignalsplux's² multi-sensor research platform together with the OpenSignal software, allowing for medical-grade data acquisitions. The corresponding sensors were attached to the human body using six pre-gelled and disposable electrodes. Signals were streamed at a 500 Hz sampling rate and 16-bit resolution per channel. Overall, participants wore four sensors consisting of (1) an ECG sensor with an Einthoven Lead I setup, (2) a piezo-electric respiration (PZT) sensor between the 8th and 10th rib, (3) a two electrode EDA sensor attached to participants' palms and (4) a reference on their right collarbone. We synchronized tracking, EEG and biosignal data, and events of the study procedure using labstreaminglayer (LSL) [200].

6.1.3.5 Experimental Protocol

*Initial setup phase
for participants*

Participants were given a general introduction to the study, i.e., we showed them the setup and explained each sensor to make sure that they were comfortable having their physiological data tracked. Next, we gathered participants' consent and asked them to fill in a demographics questionnaire. We then started with the procedure of attaching the physiological sensors. There were always two experimenters available, one identifying as male and another one as female. Participants could choose who should assist during the procedure to ensure proper placement of the sensors. Next, both experimenters fitted the EEG cap on participants' heads. Overall, the preparation time was about 40 min.

Following this, participants were placed in the VR environment and guided through an open-ended practice round, showing them the effect of C/D gain manipulations (i.e., the virtual hand moves faster than their real hand). By doing so, we allowed them to familiarize themselves with the task and the system. Once they felt comfortable, we moved to the second phase, where we calibrated their personal DT as described in the section below. Participants were instructed to grasp the proxy object as visually indicated and to maintain this pose throughout the experiment. They were told to sit comfortably and

² Biosignalsplux webpage: <https://tinyurl.com/sv7r43ju>. Last accessed: Nov 1, 2024

only move the object to the target position with a consistent and comfortable speed. The system monitored that they stayed within a reasonable time range. Once they reached the goal position, the object needed to remain in that position for 2 sec, before the visual start target appeared again. Participants were required to stay within a 5 mm distance (C/D gain = 1.0) for the countdown to remain active. To account for increased task difficulty caused by greater C/D gains, the threshold distance was adjusted using the Fitts's law method [228]. Participant and experimenters were not allowed to talk during each trial, to avoid interrupting the continuous docking task or introducing artifacts in the data. To minimize carry-over effects and cope with proprioceptive fatigue [280], participants took a break after each study conditions, and filled in a questionnaire in VR. On average, the data-collecting process was 25 min long, during which participants were not allowed to remove the VR headset. The total experiment took about 90 min per participant.

Participants performed the same movement with and without HR applied

6.1.3.6 Determining Individuals' Perceptual Boundaries

To tailor our study to each participant's individual perceptual boundaries, we calibrated their personal DTs, which were then used in the mixed and steady conditions. This was done because we found large differences between DTs, ranging from 5% to 67%, may be undetected in Section 5.1: what is an 'obvious' C/D manipulation for one person may not be perceivable by another. To achieve this, we used a 3-up-1-down interleaved staircase procedure, exposing participants to different stimuli (C/D gains) repeatedly. Using an unequal step size, we target the DT, meaning that a participant can detect the manipulation 75% of the time (see Section 2.5.2.2). We can compute the required step-size (ψ_{target}) for the step Up(Δ^+)/Down(Δ^-) method and DT = .75 as follows:

Compute individual C/D gain for HR

$$\psi_{\text{target}} = \frac{\Delta^+}{\Delta^+ + \Delta^-} \Rightarrow \frac{\Delta^-}{\Delta^+} = \frac{1 - \psi_{\text{target}}}{\psi_{\text{target}}} = \frac{(1 - 0.75)}{0.75} = \frac{1}{3} \quad (5)$$

We used the number of reversal points ($r = 3$) in each sequence as a convergence criterion. Based on our previous studies, we chose 1.0 and 2.0 for our range of manipulation factors with a 0.1 step size. Following our pilot tests, we selected 1.0 (\uparrow asc.) and 1.8 (\downarrow desc.) as the starting values for the procedure to allow for quicker convergence. Finally, we added a relative 25% to the personalized DT and since perception is known to behave non-linear [90], this ensures that the C/D gain is always noticeable.

Configuration of the threshold experiment

6.1.3.7 Data Collection

We collected data from nine sources: a pre-study questionnaire for demographic information; EEG, ECG, RSP, EDA and interaction data;

system logs (including trial times, object position and orientation, and velocity using UXF [42]); field notes and observations; and a subset of the avatar embodiment questionnaire [130] after each condition in VR using our *VRQuestionnaireToolkit* [94].

*Used a subset of
embodiment
questionnaire*

AVATAR EMBODIMENT QUESTIONNAIRE To establish a basis for the overall analysis, we first need to ensure that participants feel that the avatar is easily controlled and that the HR is noticeable. Therefore, we use a subset of the complete avatar embodiment questionnaire [130] concerning only body ownership, agency and location, resulting in ten questionnaire items. Responses were collected on a 7-point Likert scale (-3 = strongly disagree, 0 = neutral, 3 = strongly agree).

6.1.3.8 Analysis

For the analysis, our data was split into epochs corresponding to the conditions and the trials within them. Then, we pre-processed, filtered and analyzed the data using the method described below.

*Pre-processing of
EEG data*

EEG We utilized the EEGLAB³ and MoBILAB⁴ toolboxes inside the MATLAB environment for our analysis. We followed Gehrke et al. [118]'s methodology to pre-process, filter and analyze the EEG data and extract ERPs. To summarize, raw EEG data was re-sampled to 250 Hz, high-pass filtered at 1 Hz and low-pass filtered at 125 Hz. Then, the data was re-referenced to the average of all channels, followed by applying ICA to reject eye and line noise activity.

Extracting ERPs. To obtain the ERPs (shown in Figure 6.5 right), we filtered the EEG data with a 0.2 Hz high-pass and 35 Hz low-pass filter. Then we created epochs from -0.2 sec to 0.9 sec around the *MovementStartEvent*, i.e., past the jitter threshold. To guarantee robust data, we rejected 10% of the noisiest epochs [137]. We focused our analysis on one electrode, FCz, located on the forehead, which previous studies found to be a strong predictor for visual and haptic mismatches [118, 303, 306]. Furthermore, we automatically extracted the ERP negativity peaks and their latencies by locating the minimum peak in a 400–700 ms time window after the *MovementStartEvent*, using a 10 Hz low-pass filter. The time window was derived from visual inspection of the mean difference ERP wave in Figure 6.5 left.

6.1.3.9 Biosignals

We used NeuroKit2 [230] with its automated pipeline for pre-processing ECG, RSP and EDA signals.

³ EEGLAB webpage: <https://tinysurl.com/5n99v6a2>. Last accessed: Nov 1, 2024

⁴ MoBILAB webpage: <https://tinysurl.com/n47zr33x>. Last accessed: Nov 1, 2024

ECG: Raw data was cleaned using a fifth order 0.5 Hz high-pass Butterworth filter and powerline filtering at 50 Hz. R peaks were extracted using NeuroKit2's default method.

RSP: Raw data was cleaned using linear detrending followed by a fifth order 2 Hz low-pass IIR Butterworth filter. To extract RSP peaks, we applied Noto et al. [260] analysis method.

EDA: Raw data was cleaned using a fourth order 3 Hz high-pass Butterworth filter. We accepted SCR peaks below the rejection threshold set relative to the largest amplitude in the signal (min 0.1).

Finally, the results were then processed for descriptive and/or statistical analysis. After cleaning the signal, we extracted features such as number of skin conductance peaks (EDA) and we computed common event and interval-related, e.g., HRV indices [269].

Pre-processing and analysis of biosignals

6.1.4 Results

In this section, we investigate whether all participants show consistent and systematic responses across the different modalities, seeking to understand whether the monitored measures show a consistent effect across all participants when knowingly exposing them to noticeable virtual/real hand mismatches. Here, the goal is to distinguish between VR experiences that do vs. do not use noticeable HR. Our data analysis is split into five parts based on the data source; for each, we discuss which research objective the data source addresses. First, we report the results from the personalized DTs procedure. Next, we analyzed the questionnaire responses to ensure that the avatar did not consciously create dissonance for our participants. Then, we look at the EEG data followed by the interaction data, and finally, we report our results from the EDA, RSP, and ECG data analysis. The biosignals analysis excludes P9 due to data loss, and the EEG analysis excludes participants 1–4 because of an experimenter error with the electrode setup. Statistical tests were chosen based on parametric test assumptions at $\alpha = .05$, and we use outlier removal with the box plot method.

Analysis split into five parts corresponding to data source

6.1.4.1 Personalized Detection Thresholds

We collected 277 responses through the interleaved-staircase procedure. On average, it took participants 12.6 (SD = 2.6) trials to reach convergence. For each participant, we obtained one personalized DT by averaging the last three reversal points within the ascending and descending staircase sequence [385]. The results can be found in Figure 6.3. The average DT was 1.7 (SD = 0.12). All 22 staircase plots are available in the supplementary materials of the paper.

All participants converged in our staircase




Figure 6.3: DT distribution and 3-up/1-down staircase.

6.1.4.2 Avatar Embodiment Questionnaire

Here, we analyze the responses from the virtual embodiment questionnaire, namely the items concerning body ownership, agency/motor control and location, by running Friedmann tests on the aforementioned three categories. The results are depicted in Figure 6.4; all questionnaire items showed a main effect. Consequently, we ran multiple pairwise comparisons using paired Wilcoxon signed-rank tests and Holm adjustments.

Body ownership. One can have high body ownership over, e.g., a virtual hand that is dislocated from the position of the real hand when visuo-tactile stimuli are provided synchronously [130]. In the case of noticeable HR gains, the initial position of the hand is correctly rendered in place, but once a user starts the movement, a gradually changing offset is applied, leading to a break in body ownership. Q1 - “I felt if as my virtual hand was my real hand” supports this by showing significantly greater scores in both baseline conditions compared to the mixed condition.

Agency/motor control. These questions target whether participants can move their virtual body parts in synch with their real movements. Interestingly, the baseline conditions showed significantly greater scores than the mixed condition on question Q4 - “It felt like I could control the virtual hand as if it was my own hand”, but this was not the case for the steady condition. This shows that the random jumps between gain factors made it difficult for participants to predict and control their movements accurately. The steady condition, however, lies between baseline and mixed, suggesting that participants




Figure 6.4: Responses from the avatar embodiment questionnaire items. *** = $p < .001$; ** = $p < .01$; * = $p < .05$ (Holm-adjusted).

Effect on the
different dimensions
of avatar
embodiment




Figure 6.5: Mean ERP amplitudes for all study conditions across participants (left) and P8 (right). Time window for peak negativity.




Figure 6.6: Negativity peak amplitudes trend, severity and latency across conditions. *** = $p < .001$.

adapted to the (large, but consistent) offset (RQ4.5).

Location. Here, the differences between the conditions became most visible in participants' ratings. For example, the significant difference between baseline and mixed as well as steady and mixed on Q8 - "I felt as if my hand was located where I saw the virtual hand" suggests that participants clearly noticed the positional offset in the mixed condition. Interestingly, participants frequently asked: "I am not sure how to respond, because at the beginning it was faster (the effect of a high C/D gain), but at the end, it felt normal" (P11) after they finished the steady condition, showing that they adapted to even higher C/D gains (RQ4.5).

Participants noticed the offset early but adapted to it

Summary. The results demonstrate that: (1) we successfully established a solid baseline in our experiment with high body ownership, agency and location scores and (2) the induced HR illusion was noticeable for participants, significantly affecting their questionnaire responses. This ensures that the data analysis below can be linked to the effects we aim to investigate. Next, we report the result from our EEG analysis.

6.1.4.3 EEG Analysis

Since all conditions showed error negativity between 420–630 ms, we investigated the severity of the visuo-proprioceptive conflict by computing and analyzing the global minimum prediction error

*Mismatch is
reflected in the EEG
signal*

amplitudes and their latencies [118]. We found a main effect for the amplitude of a global minimum ($\chi^2(3) = 22.482$, $p < .0001$) between our four study conditions. The mixed condition showed significantly stronger prediction negativity than all other conditions as illustrated in Figure 6.6, right. The strength of global minimum prediction error amplitudes in the mixed condition is comparable to the visual mismatch negative amplitude reported by Gehrke et al. [118], demonstrating the validity of the results. Second, as shown in Figure 6.6, right, we observed no significant differences ($\chi^2(3) = 3.988$, $p = 0.262$) for the peak latencies across the four conditions. **We conclude that it is possible to distinguish the mixed from the other conditions based on error negativity peak amplitude, i.e., the severity of the mismatch is reflected in the negativity (RQ4.3).**

*Participants seem to
quickly re-calibrate
their SCs based on
their most recent
experience*

Next, we explore ERP negativity in the steady condition to consider research objective (RQ4.5), whether humans adapt to noticeable virtual/real hand offsets. We investigate the distribution of negativity peaks within our conditions across the 16 consecutive trials per participant. We observed much greater negativity in earlier trials, which decreases throughout the steady condition, quickly matching the baseline after trial 4 and reaching a stable plateau between trial 8 and 12 (see Figure 6.6). We ran a Spearman correlation analysis on the negativity peaks along the 16 trials for each condition, and found a small positive correlation in the steady ($\text{peaks}_s(2263338) = .17$, $p = .006$) and the baseline ($\text{peaks}_b(2336333) = .15$, $p = .013$) conditions, but not in the mixed ($p = .778$) and baseline_2 conditions ($p = .606$). By trial 4, there was almost no difference between steady and baseline conditions. **Our results suggest that there is a possibility of almost immediate adaptation to HR i.e., participants quickly re-calibrate themselves to the given offset.**

*Moving away from
the real-world
mapping*

Finally, we were specifically interested in the mixed condition because participants reported that this condition was “...all over the place” (P18) and was often perceived as “...movement was either too fast [C/D gain > 1] or too slow [C/D gain = 1].” (P1). We extracted the baseline trials from the mixed condition and re-ran the analysis on the global minimum prediction error amplitudes $\chi^2(3) = 213$, $p < .001$. The mixed_baseline showed significantly greater error negativity than baseline ($p < .001$) and baseline_2 ($p < .001$). **Thus, we conclude that even the one-to-one mapping can be experienced as disruptive if participants had previously adapted to a different C/D gain.**

Summary of EEG Results. We found that all conditions (incl. baseline) provoked error negativity consistently between 420–630 ms and thus, we

conclude that participants notice errors early in their movements. Additionally, the minimum prediction error amplitudes allowed us to distinguish the mixed from all other conditions. Further, participants seem to adapt quickly to the offset, reflected in decreasing error negativity. In the following we look at the interaction data.

6.1.4.4 Interaction Data

We analyze how noticeable HR affects two main movement phases: ballistic and correction. We hypothesized that C/D gain manipulations increase the relative duration of the correction phase and shorten the ballistic phase because participants start to compensate for the offset. There exist variations on how these phases are defined, but we follow Nieuwenhuizen et al. [254] to obtain ballistic and correction phases. According to Liu et al. [221]’s recommendation, we analyze movement phases across two dimensions, *distance* and *time*, to achieve a more encompassing view of the interaction.

Holistic view on movement profiles

Our central interest lies at the transition point between the ballistic and correction phases, which we normalized (by travel distance and time) to account for differences in task completion time and total distance traveled. For analysis, we computed movement profiles for all 1408 interactions (see paper supplementary materials). We ran a Friedman test on the normalized (time and distance) transition points. The Friedman test revealed a main effect for both independent variables, distance ($\chi^2(3) = 53$, $p < .001$) and time ($\chi^2(3) = 53$, $p < .001$). Hence, we conducted multiple post-hoc pairwise Wilcoxon rank tests with Holm adjustments (see Figure 6.7).

Time. The results show a significantly shorter ballistic phase in the mixed, compared to both baselines ($p < .001$) and the steady condition ($p < .001$). However, we could not find evidence for an effect between the steady and baseline conditions on the time dimension. The total task completion time ($\chi^2(3) = 176$, $p < .001$) was significantly faster in the steady condition than in all other conditions, which is unsurprising because we compensated for the increased task difficulty using Fitts’s law.

Considering time and distance

Distance. On the other hand, distance results reveal that both mixed and steady conditions have significantly later transition points ($p < .001$). As illustrated in Figure 6.7, we observed high levels of participants over- and under-shooting the target due to unexpected slow or fast movements. In contrast, baseline conditions showed consistent transition points, before or right at the required distance, indicating high accuracy. Looking at the total distance traveled ($\chi^2(3) = 213$, $p < .001$) reveals that significantly more distance was covered in the steady and mixed condition than in both

Erroneous movements result in greater total distance covered




Figure 6.7: Time and real-world distance needed to reach the virtual target. **** = $p < .0001$ (Holm-adjusted).

baselines. Again, this demonstrates the under- and overshooting in combination with many corrections due to erroneous movements, in contrast to the baseline conditions in which overshooting was nonexistent. With this in mind, we investigate whether participants adapt to virtual/real hand offsets (RQ4.5) by running a Spearman correlation analysis on the normalized transition points alongside the 16 trials for each condition. We found a medium negative correlation in the steady condition on ($\text{distance}_r(20) = -.37$, $p < .001$), and a positive correlation for ($\text{time}_r(20) = .25$, $p < .001$), both indicating that transition points move closer to normal, i.e., to both baseline conditions, with each trial. In contrast, this correlation was not present in the mixed ($\text{time}_r(20) = .08$, $p = .148$; $\text{distance}_r(20) = -.09$, $p = .106$) and baseline_2 conditions ($\text{time}_r(20) = .07$, $p = .219$; $\text{distance}_r(20) = -.02$, $p = .713$). However, baseline showed a weak correlation ($\text{time}_r(20) = .12$, $p = .034$; $\text{distance}_r(20) = .16$, $p = .006$) which may be attributed to learning effects, because it was the first study condition. Figure 6.8 illustrates how P4 adapts to a constant C/D gain in the steady condition with a large correction phase at the beginning (trial 1), and decreasing correction phases in trial 4, 8 and 12 until the end of the study condition. This is in line with our results from the questionnaire and EEG analysis (RQ4.5).

Adaptation is also visible in the movement profiles

Analogously to the ERP analysis, we extracted baseline trials from the mixed condition and re-ran the analysis. Mixed_baseline had shorter and earlier transition points than baseline ($\text{time}_p = .004$; $\text{distance}_p = .003$) and baseline_2 ($\text{time}_p < .011$; $\text{distance}_p < .001$). Again, this provides supporting evidence that there is no ‘absolute’ baseline, but instead, participants quickly adapt to the offsets, and thus, how an illusion is perceived is relative to a user’s most recent experience with this kind of interaction (RQ4.5).

Summary. Movement profiles, time and distance appear to be reliable metrics to distinguish unnoticeable vs. noticeable HR offsets because these are directly affected by the performed interaction (RQ4.2). Further, we found that participants adapt to consistent offsets in the steady condition (RQ4.5), which was not the case in the mixed condition.




Figure 6.8: Movement normalizes and becomes more accurate in the steady condition over time. Correction phase becomes shorter.

6.1.4.5 EDA, RSP and ECG Data Exploration

EDA Participant 15, 17, and 20 were omitted from the analysis because it was not possible to compute SCR peaks, possibly due to noisy data. The corresponding raw/cleaned data plots of all participants can be found in the paper's supplementary materials. To investigate our hypothesis that the total number of SCR peaks increases in the mixed condition because it creates arousal, we ran an RM ANOVA. However, we did not find a main effect on the total number of SCR peaks across our study conditions ($F_{(3,48)} = 2.025$, $p = .123$). We further analyzed the data by examining the SCR plots. Overall, we observed that the majority of SCR peaks seem to appear at the beginning of a study block, which might suggest that this phase is used for initial self-calibration—even in the baseline condition, although there was a break and a questionnaire between study conditions. To further investigate this, we split each study condition into quartiles, analogous to an extreme groups approach [273], and then computed the number of SCR peaks for each quartile/condition. The results are depicted in Figure 6.9, showing a downward trend in the steady condition similar to the baseline, especially in the 3rd quartile (trial 9 to 12), but an equal distribution of SCR peaks in the mixed condition.

Participants used the initial movements for self-calibration

Summary of EDA Results. *Even though we could not identify a significant effect on the overall number of SCR peaks, our results may be interpreted as evidence that participants adapt to even larger offsets (RQ4.5) but fail to adapt to the constantly changing offsets in the mixed condition, reflected in the distribution of SCR peaks.*

RSP AND ECG To the best of our knowledge, we are the first to explore the relationship between ECG and RSP in the context of hand-based illusions. Therefore, hypothesis testing using statistical methods is inappropriate; instead, we use a descriptive and exploratory data analysis approach, splitting the data into four epochs corresponding to our four study conditions. We performed an interval analysis, computing the most-common RSP, time, frequency as well as non-linear ECG indices accordingly to Makowski et al. [230] and Pham et al. [269]. Overall, we obtained 86 indices that

Exploratory nature of our analysis

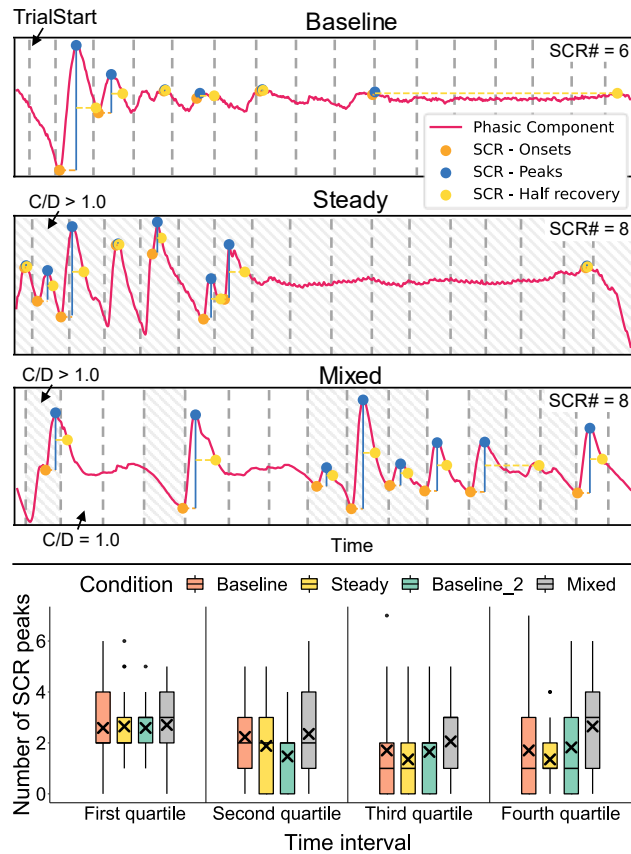


Figure 6.9: Top: Phasic component during baseline, steady, and mixed condition. Bottom: Number of SCR peaks separated in quartiles.

may be selected for statistical analysis or ML approaches. Here, we only report a subset of our results. However, we make the data set publicly available to inform future research hypotheses and studies that can then be evaluated in dedicated studies.

First, we descriptively analyze respiration profiles because, in our study, we observed that participants' responses differed dramatically. For instance, in the mixed condition (see Figure 6.10: bottom), P11 held his breath for several seconds followed by heavily in-/exhaling, because he was focused on matching the target and start position while coping with the effects of frequently changing C/D gains. On the other hand, P5 reacted 'surprised' when the C/D gain changed, triggering smaller exhales corresponding to the trials, leading to a completely different breathing pattern, demonstrating how different individual participants' reactions can be. However, most participants seem to fall in the latter category, according to RSP rate and amplitude measures, when comparing mixed with baseline conditions. Overall, there seems to be a trend towards an increase in RSP rate and more shallow breathing in the mixed condition.

*Respiration was
visibly affected when
experiencing strong
HR*




Figure 6.10: Different RSP responses when exposing P5 (top) and P11 (bottom) to frequently changing C/D gains in the mixed condition.

	Rate Mean	Ampli Mean	RRV CVSD	Inspira tion	Expira tion	Median BB	Mean BB	Cohen's d
Baseline/ Mixed	.333	.743	.502	.493	.162	.507	.400	1.0
Baseline_2/ Mixed	.437	.459	.477	.384	.547	.527	.524	0.8
Baseline/ Steady	.161	.454	.014	.023	.023	.264	.229	0.6
Baseline_2/ Steady	.307	.307	.014	.218	.323	.264	.293	0.4

Figure 6.11: RSP Cohen's d in combination with dimensional reduction.

Principal components analysis. Next, we computed 22 common RSP indices, and we performed a dimensional reduction using principal components analysis (PCA) in order to determine the number of components that are needed to describe the variance in the data. In the next step, we check whether the PCs reveal clusters within our collected sample, allowing us to distinguish study conditions. The PCA showed that 65.2% of the variance can be explained by PC1 and PC2. Five additional PCs are needed to reach the often-propagated total variance of 80% for exploratory data analysis (see Figure 6.12). However, when inspecting the biplots, we see that even the most powerful principal components PC1 and PC2 do not indicate a distinct separation in clusters that correspond to our study conditions (see Figure 6.12). Although this is desirable for both baseline conditions, it shows that there seems to be no consistent pattern within our collected data which can be linked to our study conditions.

Exploratory data analysis suggests that the RSP features lack predictive power

Standardized mean difference. Next, we compared the standardized mean difference between conditions using paired Cohen's d. A d of .5 is generally considered a medium effect size, meaning that 69.1% of one group will be above the mean of the other group (Cohen's U3). Six indices between baseline and mixed conditions




Figure 6.13: ECG Cohen's d in combination with dimensional reduction.

with constantly changing C/D gains in the mixed condition could have an effect on them. However, other well-studied stress indices, such as RMSSD, remained consistent across all study conditions and did not suggest any differences. Figure 6.13 shows the resulting d scores of each variable that contributed to the first seven principal components with at least a medium effect size of $d > .5$.

Summary of ECG Results. We descriptively analyzed ECG data, laying the foundations for future research that aims to detect noticeable illusions using heart rate data. We provide an overview of common ECG indices, helping the community to make informed decisions about research hypotheses by outlining the most promising ECG features. Specifically, the stress-related index SDNN, as well as the SD1SD2 and CSI_Modified indices, appear to be promising candidates for future investigations. However, it must be noted that based on our PCA analysis, there exist inconsistencies in participants' ECG responses within our collected sample (RQ4.1).

6.1.4.6 Summary of Results Towards Research Questions

RQ4.1: We investigated EDA, RSP and ECG, providing an overview of the most promising indices that may be selected for future hypothesis-driven studies. Our analysis revealed inconsistencies for RSP and ECG, while simultaneously demonstrating substantial individual differences in participants' physiological responses. Hence, at this point, we are unable to conclude if these measures can or cannot be used to distinguish study conditions. Our analysis on the SCR peaks was inconclusive, but we identified a downward trend in SCR peaks except in the mixed condition, which seems to indicate that participants adapted to the offset.

RQ4.2: Interaction data, specifically the analysis of movement phases and transition points, showed a significant difference that could be systematically linked to our study conditions.

RQ4.3: ERPs were found in all conditions, but differed in their amplitudes and negativity, which allowed us to distinguish the mixed

What do our results mean in the context of our research questions?

from the other conditions. The steady condition seemed (except for the initial strong negativity during the first four repetitions) mostly indistinguishable from the baseline conditions.

RQ4.4: The ERP analysis suggests that noticeable visuo-proprioceptive offsets were detected early after movement start, which is supported by the movement analysis, showing very early transition points from ballistic to correction phase in the mixed and steady conditions, caused by participants compensating for the offset.

RQ4.5: We found overwhelming evidence that participants adapted to the offsets in the steady condition through analyzing participants' comments, questionnaire, EDA, movement and EEG data. The exact number of repetitions needed most likely depends on the individual. However, our data suggest that somewhere around repetition 4–8, participants were adapted to the offset. Interestingly, steady and baseline conditions showed very similar patterns across all measures, suggesting that participants used the initial phase for self-calibration and to establish a robust model of their environment.

6.1.5 Discussion

6.1.5.1 Combining Physiological Modalities

In this work, we investigated a set of physiological measures that we intend to expand, including eye blinks, pupil dilation and gaze. The next logical step is to combine multiple modalities to classify whether a visuo-proprioceptive offset was noticed or not. Just as we found in our prior research that DTs differed between individuals, we observed that our participants' individual physiological responses also differed. Similarly, we expect that each user would have their own personalized classifier. Based on our data, it is unclear whether these responses would remain stable—an individual may, for instance, differ from one day to another. Running an analysis on individuals across time requires substantially more data. This is difficult because of the trade-off between collecting sufficient samples versus proprioceptive fatigue [280]. It is thus possible that the lack of consistency in our data may be attributed to the small sample size; however, our effect-size analysis provides a foundation and a solid starting point for future hypothesis-driven studies.

*Physiological
responses underlie
large variance*

6.1.5.2 Validity & Applicability of EEG Results

Our results confirm that the frontal cortex area, especially the FCz electrode, and the concept of ERPs [118, 303, 306] can be used to detect noticeable visuo-proprioceptive offsets. However, previous results show a more consistent ERP mean curve, which can be explained

by the distinct events used in their studies. Instead, our work introduces gradually increasing offsets, where the precise time of violation is unclear and thus, temporal shifts in the ERP curves occur. To our surprise, even the baseline conditions showed weak error negativity, which is reduced over time (see Figure 6.5) and therefore, may be caused by the visualization itself [306] (i.e., the simple floating hands combined with reduced depth perception in VR). Other measures than ERPs to detect a mismatch also exist and could be used for comparison with our results. For instance, a haptic delay has been shown to significantly increase beta and theta band activity [8]. Another interesting aspect is the indistinguishability between the steady and baseline conditions, when compared to the questionnaire responses. We assume that the questionnaire responses were given by participants averaging their experiences in the steady condition, leading to a score in-between those of the baseline and mixed conditions. In contrast, the EEG results perhaps give a more direct estimate of how the illusion was perceived, also showcasing the quick adaptation. Lastly, determining the time (RQ4.4) when a violation gets noticed is a crucial aspect for many techniques that aim to increase DTs for illusions in VR (e.g., by utilizing eye blinks and saccades [381]).

ERPs are challenging without a temporal anchor

Questionnaire data provide insights about users' perception

6.1.5.3 Beyond Noticeable Hand Redirection

Our long-term goal is to tailor illusions to an individual's perceptual boundary. In this work, we used offsets that are obvious to the user; however, more conservative studies found that the DTs are much smaller [384]. Therefore, we need to investigate offsets that are not only above but also around and below participants' thresholds. Future work needs to investigate whether unnoticeable offsets trigger similar effects, allowing us to differentiate them. In our vision, VR designers and systems could tell when they reach the perceptual boundary of an individual and stay just below this threshold. Therefore, we need to study more realistic settings and scenarios to investigate the approach's robustness. Finally, it would also be interesting to apply this method for different types of hand-based illusions such as our *Pseudo-Haptic Resistance* technique or *Redirected Touching* [193].

Basis for our next experiment

6.1.6 Conclusion & Contributions

In this section, we contributed to RQ4 by investigating physiological and interaction data to implicitly detect noticeable HR offsets between users' movements of their real and virtual hand (C4). We designed an experiment to (1) determine noticeable HR gains and (2) collect EEG, ECG, EDA, RSP and interaction data for no HR and noticeable HR. Given the novelty of our research, we conducted an exploratory data analysis to identify the most promising features.

*ERP peak negativity
appears to be a
promising candidate*

First, our experiment with 22 participants contributes to **C3** by extending the previously determined DT [144, 384] for gain-based HR. Our EEG results suggest that especially ERPs and their peak negativity are reliable in detecting strong visuo-proprioceptive offsets, making them a valid candidate for further explorations. This is in line with previous results on error detection in VR [118, 303] and research on semantic body violations [265]. Thus, they show great potential to address **RQ4**, and therefore, we further investigate them in the next Section 6.2.

*Redirected
movements have
distinct properties*

Furthermore, movement phases appeared to be directly affected by noticeable hand offsets. Already in Section 4.2, we found evidence that in the presence of visuo-haptic redirection, movement accuracy, and time change according to the magnitude of the redirection. These findings challenged existing movement prediction models that are crucial for planning and executing effective redirection [131]. This resulted in a new string of research, e.g., by Gonzalez and Follmer [134] who propose a new model for the trajectory prediction of the hand during redirected reaching in VR. Given these promising results, we use these measures in Section 6.2.

*ECG and RSP have
limited utility for
our research goal*

Next, we provide the first investigation of ECG and RSP in relation to the illusion of HR, and we recommend indices that may be studied in further experiments. While these physiological measures may work for separating conditions with vs. without noticeable HR, their slow-responding nature make them unsuitable for our research goal **RQ4**, because we require measures that immediately react according to user perception of the interaction. Nevertheless, there is still great potential to predict affective states [171] or potentially even monitor users' felt long-term presences inside an IVE.

*Quick adaption to
HR offsets*

Interestingly, we found that participants quickly adapted to larger offsets within 4–8 trials, and whether an offset remains unnoticed primarily depends on their most recent experience with this kind of interaction. This provides evidence for previous observations of Kohli [193], reporting a quick adaption to offsets. Moreover, our results provide valuable insights into the relationship between learned SCs, their temporal validity and robustness, and suggest that the current calibration of the user is another factor to consider during ongoing VR applications (**C3**). It appears that participants used their initial hand movements to establish a mental model of their environment, which may also explain why participants with previous VR experience had lower thresholds in Section 5.1.

Finally, we contribute a data set containing EEG, ECG, EDA, RSP, and interaction features that the community can be built upon (**C6**). There

is great value in such data sets considering the financial and time required to collect this type of data. In summary, we conducted a first investigation of implicit measures that allow us to detect if an individual experienced noticeable HR, laying the foundation for the next section that attempts to differentiate between HR offsets of different magnitudes corresponding to perceptual boundaries of individuals based on a single movement.

*Open source data set
including all
collected features*

6.2 TAILORING ILLUSIONS TO USERS' SENSITIVITY



Figure 6.14: We use the psychophysical method of constant stimuli to determine participants' perceptual boundaries for horizontal HR Below, At and Above their individual DT. Next, we collect movement, eye gaze and EEG data, compute features, and analyze them using frequentist and Bayesian statistics. Finally, we trained a multimodal classifier using Random Forest to predict if participants are exposed to HR of different magnitudes corresponding to their perceptual boundaries based on a single movement.

6.2.1 Introduction

In this section, we take the next step by exploring the potential of combining the most promising movement and EEG features from the previous Section 6.1, together with eye gaze data to distinguish whether an applied HR offset is Below, At, or Above a user's individual DT. In our vision, this method allows constant monitoring of a user's tolerance to the exposed hand-based illusion throughout an entire VR experience. This would allow the VR system to adjust the magnitude of employed HR dynamically, depending on context, interaction, and an individual's sensitivity to visuo-proprioceptive conflicts. Therefore, the goal is to understand if movement, eye gaze, and EEG data can substitute a DT experiment, eventually allowing continuous adaptation of HR offsets.

6.2.2 Experiment

We designed a 2-part experiment, investigating whether gradual horizontal HR around an individual's perceptual boundary can be detected using the three modalities: movement, eye gaze and EEG. In part 1, we used the psychophysical method of constant stimuli analogous to Steinicke et al. [322] and Zenner and Krüger [384] to model the discrimination performance for each participant for the specific type of HR (i.e., gradual horizontal offsets of the virtual hand to the right; see Figure 6.1: left). The results were used in part 2 of the experiment, which was tailored to each participant with HR offsets that corresponded to their perceptual boundaries. This way, we ensured that each participant was exposed to the same magnitude of perceived offsets. In part 1 and 2, we applied a 2IFC (see Section 2.5.2.2), where participants were instructed to perform two consecutive hand move-

*Our Vision for
hand-based illusions*

*Tailor experiment to
each participant's
sensitivity to offsets*

ments, hitting a virtual target with their index finger. During the first movement, no HR was applied, whereas in the second movement, we either applied HR of different magnitudes or no HR. Participants were asked to compare both movements and report if they felt a difference between them by responding to the $1AFC$ question: “Both movements felt the same” (see Figure 6.14: left). Participants could respond by using the ‘Yes’ or ‘No’ button on a presenter stick in their non-dominant hand [385]. Subsequently, the participants returned their hand to the initial location and continued with the task. The location of the virtual target (a red sphere) always appeared in the same location, 30 cm in front of them [97]. Although the task remains the same in part 1 and 2 of the experiment, the underlying methodologies and objectives differ substantially.

PART 1—DETERMINE PERCEPTUAL BOUNDARIES To tailor part 2 to each participant’s individual perceptual boundary, we modeled their discrimination performance in distinguishing movements with HR vs. no HR. To do so, we conducted a psychophysical threshold experiment, fitting the psychometric quick function [181] through our collected sample by optimizing the parameters α and β . We define these probabilities as perceptual boundaries, Below (25% probability), At (50% probability), and Above (75% probability), with At representing the CDT. This means that there is a 50% chance that a participant can detect the presence of HR, respectively, for 25% and 75%. 75% (Above) is often used as a less conservative threshold in the literature, whereas 25% (Below) was chosen to include a sample below the DT, investigating if participants respond to the offset even without consciously noticing it. To model the discrimination performance for each individual, a sufficient amount of data is needed. Based on the HR literature and our pilot tests, we arrived at the following configuration for the method of constant stimuli. We tested offsets ranging from 0 cm to 7 cm in increments of 1 cm, resulting in eight stimuli. Participants experienced each stimulus eight times (= 64 stimuli; in total 128 hand movements) to improve the robustness of the fitting and allow for consistency checks.

*Modelling
participants’
discrimination
performance for
horizontal HR*

PART 2—COLLECTING DATA AT PERCEPTUAL BOUNDARIES In the second part, participants performed a total of eight rounds of the discrimination task while only exposed to HR offsets Below, At, Above as well as no HR (Base). Each round consisted of 16 stimuli trials, $4 \times$ Below, At, Above, and Base, presented in a randomized order, resulting in 32 reaching movements per round. The $2IFC$ method allowed us to include a sufficient amount of ground truth reaching movements (no HR) on what participants experience to be ‘normal’ [118], which must be captured in VR [78]. In this part of the experi-




Figure 6.15: Effect of horizontal HR and our events used during aimed reaching movements in our experiment.

ment, the main focus was on collecting movement, eye gaze, and EEG data at participants' perceptual boundaries.

6.2.2.1 Participants

We recruited 18 right-handed participants (six females, twelve males), aged 18–31 (mean = 25.05; SD = 3.05) from the general public and the local university. We asked participants not to consume alcohol or caffeine twelve hours before the study. Participants had a range of different educational and professional backgrounds, including media informatics, computer science, education, pharmacy, cybersecurity, entrepreneurship, biomedical engineering, data science, and artificial intelligence. All participants reported normal or corrected-to-normal vision and did not report any known health issues which might impair their perception or proprioception. Nine participants had never used VR before, six had used it a few times (1–5 times a year), no one reported using it often (6–10 times a year), and three others used it regularly (more than 10 times a year). Participants not associated with our institution received €30 as remuneration for participating in the experiment. The study was approved by the University's Ethics Board.

Participants were asked to not consume alcohol or caffeine twelve hours before the study

6.2.2.2 Apparatus

We used a simple virtual environment consisting of a table, the experimental setup, and an instruction screen, which was implemented in Unity3D (v.2022.2.0). We included an androgynous representation of the virtual hand [300] to prevent unwanted effects such as a drift in DT [262]. The experimental logic was implemented using the Unity Experiment Framework (UXF v.2.4.3) [42], the Unity Staircase Procedure Toolkit [379] and our *VRQuestionnaireToolkit* [94]. Participants remained seated on a chair throughout the experiment with a table in front of them. They wore an HMD, an EEG headset and a Vive tracker attached to their dominant hand with a finger spline to fixate their

Extending previous testbed to ensure comparability between experiments

index finger. We used the HTC VIVE Pro Eye tracking system (SRanipal SDK) to capture eye movements. The experiment ran on an XMG PRO One offering an Intel® Core i7-10870H CPU, 32 GB RAM and an Nvidia® GeForce RTX 3070.

EEG SETUP EEG data were captured from 32 actively amplified electrodes using BrainAmp DC amplifiers from BrainProducts like in the previous experiment. Electrodes were placed according to the international 10–20 system, using the nasion/inion as reference points. To establish a connection between the electrodes and the scalp, conductive gel was applied and the impedance of all active electrodes was reduced to 5–10 kOhm before the experiment started [118]. The EEG data were sampled at a rate of 500 Hz.

Same EEG setup as in the previous experiment

6.2.2.3 Experimental Protocol

Participants arrived at the location and first received a general introduction to the study, i.e., we showed them the setup and explained the EEG headset to ensure that they were comfortable with it. Next, we gathered participants' consent and asked them to fill in a demographic questionnaire. We then started with the procedure of attaching the EEG electrodes to the heads of our participants. This procedure was carried out with two experimenters, one identified as male and one as female, to improve the comfort of our participants and to reduce the preparation time to about 40 min. Subsequently, participants were placed in the IVE and guided through an open-ended practice round, showing them the effect of horizontal HR. By doing so, we allowed them to familiarize themselves with the system and the task. Once they felt comfortable, we moved to the first part of the experiment, where we modeled their discrimination performance.

Introduction and preparation phase

Participants were told to sit comfortably and to move their hand to the target position at a consistent and comfortable speed. The system monitored that they stayed within a reasonable time range. Once their virtual index fingers reached the goal position, their finger needed to remain in that position for one second before the 1AFC question appeared. Participants were required to stay within a 5 mm radius for the dwell time indicator to remain active. Participants and experimenters were not allowed to talk to avoid interrupting the continuous docking task or introducing artifacts in the data. This part of the experiment took 40–45 min.

Collect data for determining individual sensitivity to HR

Next, participants took a longer break (about 15 min), while the experimenters configured part 2 of the experiment. In part 2, participants performed the same task tailored to their perceptual boundaries. After each of the eight rounds, participants took a break to reduce the effects of proprioceptive fatigue [280]. On average, the data

Exposing participants to offset corresponding to their perceptual boundaries

collection took 35–40 min, during which participants were not allowed to remove the VR headset to avoid moving the EEG electrodes. In total, the experiment was about 2.5 h, and we provided complimentary snacks and water.

6.2.2.4 Data Collection

We collected data from six sources: a pre-study questionnaire for demographic information; EEG, eye tracking and movement data; system logs (including trial times, object position and orientation, and velocity); and we collected participants' responses to the 1AFC question in VR [94]. To synchronize our data streams with VR interactions and the events, we again used the lab streaming layer (LSL)[200].

6.2.2.5 Events, Pre-Processing & Analysis

Our data from part 2 of the study were split into epochs corresponding to the conditions, the trials, and the events within them. We pre-processed, filtered, and analyzed the data using the methods described below. An overview of the events can be seen in Figure 6.15. The *trialStart* event is triggered after the participants successfully held the start position for 1 sec and moved 5 mm away from the start, and the *trialEnd* event as soon as they reached the target.

*Movement time,
peakvelocity and
transition points*

GAZE AND MOVEMENT DATA. We statistically analyzed our data after verifying the parametric test assumptions at $\alpha = .05$. We performed RM ANOVAs and applied Greenhouse–Geisser corrections when the assumption of sphericity was violated. In the presence of a main effect, we performed post hoc pairwise comparison t-tests adjusted using the Bonferroni-Holm method. In addition, we conducted a Bayesian analysis using JASP⁵ following Wagenmakers et al. [353]'s method. We exported our previously used measures *totalTime* and *peakVelocity* [289]. We extracted the transition points between ballistic and correction phases according to Liu et al. [221] in the time (*transitionPointTime*) and spatial domain (*transitionPointDistance*). We statistically analyze the two gaze features, *#handFixations* and *durationHandFixations* previously used by Lavoie et al. [203]. We define *#handFixations* as the number of gaze intersects where the virtual hand is fixated for at least 60 ms. *DurationHandFixations* is the total duration of hand fixations that are ≥ 60 ms.

*Eye gaze fixation
and duration with
the virtual hand*

EEG. Our analyses focused on the midline electrodes FCz, Cz, and Pz that have been successfully used to detect error and mismatches in VR [118, 119, 303], and our previous experiment in the section before.

⁵ JASP webpage: <https://jasp-stats.org/>. Last accessed: Nov 1, 2024

EVENT-RELATED POTENTIALS We followed Gehrke et al. [118, 119] approach to extract single-trial ERPs. After applying a band-pass filter from 0.1 to 15 Hz, ERPs were extracted around three event markers coupled to the hand movement: *peakAcceleration*, *peakVelocity*, and *transitionPointTime*. ERPs were baseline corrected by subtracting the average amplitude of the last 100 ms preceding the trial start. To ascertain the effects, the linear mixed-effects model 'eegfeature ~ condition * modality + 1|participantID' was fit at each time point. Effects were assessed using likelihood ratio tests for the main effects with Benjamini-Hochberg *p*-value correction for false discovery rate [27]. For post-hoc analyses, we specifically focus on the time window between 150–250 ms following salient moments of the movement phase with respect to HR. This time window is of key interest as a negative going deflection in the ERP here has frequently been linked to the detection of mismatches between what is predicted to be picked up by the sensory organs and what they actually sample [116]. In a confirmatory step, we defined the peak velocity of the reaching motion as a salient moment of the movement, as it marks the beginning of the correction phase towards reaching the target [289]. Specifically in HR, we believe this moment is a good approximation at which the HR offset may be consciously experienced [307].

Extracted ERPs at several temporal anchors

To address the specific aspects of unnoticeable hand-based illusions, we further propose time-frequency decomposed EEG data, which has yet to be explored as a novel direction in characterizing and understanding brain responses to HR. Spectral features may hold significant promise for a continuous metric describing participants' individual perceptual boundaries, as they do not require as precise a temporal anchor for meaningful feature extraction as do ERPs. One metric based on spectral features is the ratio of frontal theta power and parietal alpha power. This ratio has been shown to correlate with a subjective rating of workload or an increased cognitive load [72, 122]. We explore this novel metric as a correlate of cognitively processing HR since predicting the correlation between one's own movements and the HR gain likely increases spatial processing demands.

When does the disruption occur?

EVENT-RELATED SPECTRAL PERTURBATION (ERSP) First, we set out to confirm that our experimental task elicited robust spectral brain modulations. Hence, the evoked spectral response was compared to a baseline. To this end, grand-average event-related spectral perturbations (ERSP) were computed using the 'newtimef' function in EEGLAB (3 to 100 Hz in logarithmic scale, using a wavelet transformation with 3 cycles for the lowest frequency and a linear increase with frequency of 0.5 cycles). In order to account for different trial

ERSP as a novel approach to detecting HR of different magnitudes

segment duration and maintain the time-frequency resolution across participants, the spectrograms were linearly time-warped to the median times of the movement, i.e., the median time of movement onset, max. acceleration, max. velocity, and the onset of the correction phase. Then, a spatio-temporal cluster test (using MNE-python [138]) was conducted in comparison to power values in a -300 to -100 ms pre-trial baseline window. Lastly, we focus our analyses on one specific spectral feature: the ratio of theta band power at electrode FCz and alpha band power at Pz. This ratio has been shown to correlate with a subjective rating of workload or an increased cognitive load [72, 122]. Effects were assessed analogous to the ERP analyses.

6.2.3 Results

6.2.3.1 Part 1—Determine Perceptual Boundaries

We computed the thresholds at 25%, 50% and 75% detectability based on the fittings of the psychometric function. The results for 75% detectability are depicted in Figure 6.16, suggesting that the participant provided consistent responses. Plots for 25% and 50% can be found in the supplementary materials. However, for P09, the fitting did not reach convergence. As a result, we could not compute the DTs, which means that the participant could not continue part 2 of the experiment. This could have happened for various reasons. For example, the participant perhaps did not really understand the study task, or the HR offsets tested were too small. However, for the remaining 17 participants, we were able to compute DTs shown in Figure 6.17. The horizontal HR 50% DTs obtained are comparable to those of the existing literature [207, 384]. Furthermore, Figure 6.17 (left) supports Hartfill et al. [144]'s and our previous findings towards personalized DTs, because thresholds differ substantially across par-

*DTs are comparable
to literature*

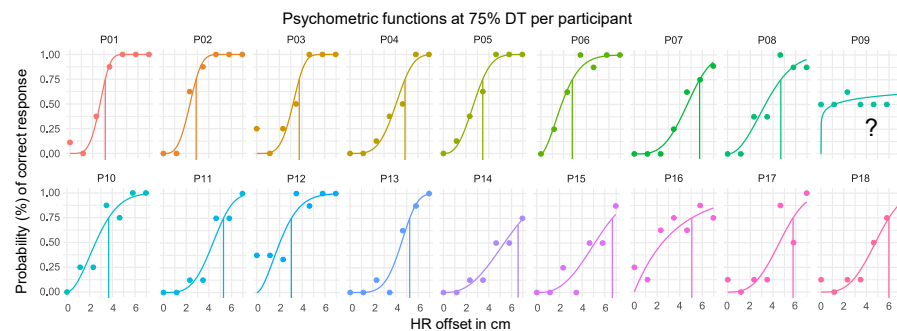


Figure 6.16: Plotted psychometric functions with 75% probability of a correct response (i.e., 75% DT) for each participant. The modeled discrimination performance shows an S-shaped curve typical for human perception. For P09 marked with a '?', we could not compute DTs, because all stimuli were perceived as equal according to the discrimination performance.




Figure 6.17: Obtained thresholds corresponding to individuals' perceptual boundaries. *Below* shows significantly lower thresholds than *At* and *Above*. *At* has significantly lower thresholds than *Above*.

ticipants, but are consistently high or low for each individual (see Section 4.2). This further demonstrates the need for novel approaches to tackle this problem, supporting our overarching research objective.

Threshold experiments can be subject to noise and are very sensitive to their configuration (#repetitions, #steps, etc.). For example, it could be that 25% and 50% result in DT clusters that overlap and are perceived as more or less the same. Therefore, verifying that our HR DTs are perceptually different is a prerequisite for part 2. We statistically analyzed the resulting thresholds and found a main effect ($F(1.036) = 105.7$, $p < .001$, $\eta_p^2 = .869$, $BF_{incl} > 1000$) of condition on the DTs. Post-hoc tests revealed that *Below* has significantly lower thresholds than *At* ($p < .001$, $d = -1.663$, $BF_{10} > 1000$) and *Above* ($p < .001$, $d = -3.524$, $BF_{10} > 1000$). Similarly, *At* showed lower thresholds than *Above* ($p < .001$, $d = -1.861$, $BF_{10} > 1000$) with strong positive correlations ($p < .001$, $BF_{10} > 280$, with $\rho > .8$) between *Below*, *At*, *Above* based on the DT. **As a result, we can confirm that our obtained HR DTs are perceptually different.**

Perceptual HR boundaries are significantly different

6.2.3.2 Part 2—Distinguish Perceptual Boundaries

To ensure that participants did not suffer from fatigue, we first visualized their discrimination performance for the eight rounds in Figure 6.18. The graph suggests that there is no notable shift in *Base*, *Below*, *At*, *Above* over the eight rounds. Bayesian analysis provided strong evidence for the absence of an effect between study round and the four conditions, *Base* ($F(7) = 25.584$, $p = .520$, $\eta_p^2 = .064$, $BF_{excl} = 10.6$), *Below* ($F(7) = 0.664$, $p = .702$, $\eta_p^2 = .049$, $BF_{excl} = 16.5$), *At* ($F(7) = 0.443$, $p = .872$, $\eta_p^2 = .033$, $BF_{excl} = 24.4$) and *Above* ($F(7) = 0.780$, $p = .605$, $\eta_p^2 = .057$, $BF_{excl} = 13.1$). **Thus, we conclude that individuals' thresholds remain consistent throughout part 2 of**

DTs remained robust throughout the experiment




Figure 6.18: Participants' discrimination performance at Base, Below, At and Above their individual DTs over the eight study rounds. The horizontal red line shows the expected probability of a correct response for Base = 0%, Below = 25%, At = 50% and Above = 75%. The boxplots display participants' responses throughout part 2 of the experiment. Visual inspection and Bayesian analysis suggest that there is no noticeable difference in discrimination performance, e.g., caused by fatigue.

time \ peakVelo	Base	Below	At	Above
Base	-	$p = 1.0$; $B_{10} = 0.403$	$p = 1.0$; $B_{10} = .279$	$p = 1.0$; $B_{10} = .305$
Below	$p = .028$; $B_{10} = 3.6$	-	$p = 1.0$; $B_{10} = .270$	$p = 1.0$; $B_{10} = .260$
At	$p = .001$; $B_{10} = 106.7$	$p = .208$; $B_{10} = .481$	-	$p = 1.0$; $B_{10} = .259$
Above	$p < .001$; $B_{10} > 1000$	$p < .001$; $B_{10} = 263.2$	$p < .001$; $B_{10} = 311.7$	-
transitT \ transitD	Base	Below	At	Above
Base	-	$p < .001$; $B_{10} = 183.2$	$p < .001$; $B_{10} > 1000$	$p < .001$; $B_{10} > 1000$
Below	$p = .437$; $B_{10} = .745$	-	$p = .005$; $B_{10} = 17.3$	$p < .001$; $B_{10} > 1000$
At	$p = .116$; $B_{10} = 1.4$	$p = .437$; $B_{10} = .433$	-	$p = .008$; $B_{10} = 6.2$
Above	$p < .001$; $B_{10} = 311.5$	$p < .001$; $B_{10} = 185.4$	$p < .001$; $B_{10} = 193.4$	-

Figure 6.19: P -values and Bayesian factors for the four movement features *totalTime*, *peakVelocity*, *transitionPointTime* and *transitionPointDistance*.

the experiment, allowing us to link our analysis back to the established perceptual boundaries.

MOVEMENT DATA. We extracted and analyzed the four features, *totalTime*, *peakVelocity*, *transitionPointTime* and *transitionPointDistance*. We found evidence for a main effect on *totalTime* ($F(3) = 25.584$, $p < .001$, $\eta_p^2 = .615$, $BF_{incl} > 1000$), *transitionPointTime* ($F(3) = 15.792$, $p < .001$, $\eta_p^2 = .497$, $BF_{incl} > 1000$) and *transitionPointDistance* ($F(3) = 40.493$, $p < .001$, $\eta_p^2 = .717$, $BF_{incl} > 1000$), but not for *peakVelocity* ($F(3) = 0.336$, $p = .724$, $\eta_p^2 = .021$, $BF_{incl} = 0.111$). Post-hoc tests showed significant differences between movements without HR and any other condition. Transition points from the ballistic to the correction phase appeared significantly earlier, and hand movements took significantly longer when horizontal HR was applied. The latter effect is in the opposite direction to what we already found for gain-based HR. This is an interesting finding that can be explained by the HR direction. Gain-based HR effectively leads to shorter movements because less physical distance is required to reach the virtual target, in contrast to horizontal HR, which increases physical movement distance (see Figure 6.15), and most likely also task difficulty. Figure 6.19 reports the test statistics, which are in line with previous findings that investigate the effects of redirected movements [134, 207]—but we ex-

Trajectory not only predicts noticeability, but also different magnitudes of HR




Figure 6.20: *#HandFixations* occurred significantly more frequently in the Base condition and Below the threshold than Above it. Besides looking more frequently to the hand, participants also spent more time looking at the hand in the Base condition than At and Above the threshold.

tend them to distinct perceptual boundaries around the noticeability threshold. **Thus, our findings support (H1), confirming the potential of movement data to differentiate between HR of magnitudes corresponding to individuals' sensitivity to visuo-proprioceptive offsets.**

GAZE DATA. Next, we analyze the two features, *#handFixations* and *durationHandFixations* depicted in Figure 6.20. We found evidence for a main effect for both *#handFixations* ($F(2.322) = 5.616$, $p = .005$, $\eta_p^2 = .260$, $BF_{incl} = 16.141$) and *durationHandFixations* ($F(1.9) = 6.510$, $p = .005$, $\eta_p^2 = .289$, $BF_{incl} = 35.118$), depending on the condition. Post-hoc pairwise comparisons for *#handFixations* showed significant differences between Base and At ($p = .002$, $d = .929$, $BF_{10} = 7.453$) as well as Below and Above ($p = .018$, $d = .742$, $BF_{10} = 3.401$). *DurationHandFixations* showed significant differences between Base and At ($p = .002$, $d = .916$, $BF_{10} = 8.306$) as well as Base and Above ($p = .002$, $d = .938$, $BF_{10} = 5.462$). Bayesian analysis provided evidence for the absence of an effect between Base and Below for *#handFixations* ($BF_{10} = 0.311$), and between At and Above for *durationHandFixation* ($BF_{10} = 0.252$). **Contrary to (H2), participants looked at their virtual hand more frequently and for longer during hand movements without HR than in any other condition. We believe that this could be the result of the nature of the task, which we further discuss in Section 6.2.5. Nevertheless, *#handFixations* and *durationHandFixations* clearly separate movements without HR from movements with HR at individuals' 75% DT.**

Participants look at their hand more frequently and for longer duration during normal hand movements

EEG DATA. First, we examine ERPs by plotting the mean ERP amplitudes for the electrodes FCz and Cz (available in supplementary materials of the corresponding paper) located above frontal cortical areas analog to [118, 119, 265] and our previous study. However, we did not observe the typical ERP amplitude following prediction

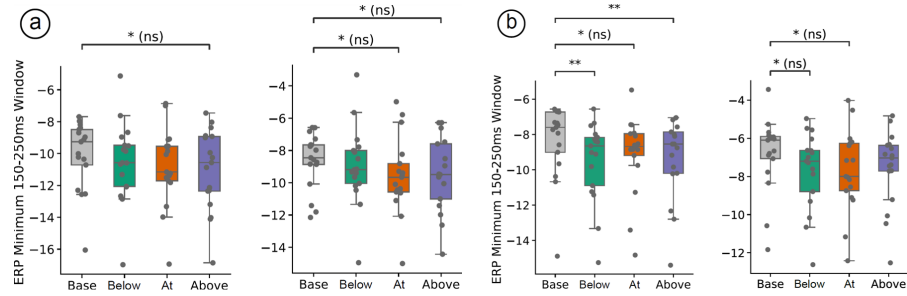


Figure 6.21: Amplitude minima in the 150–250 ms time window for each condition at electrode FCz (a) and Cz (b) following peak acceleration (left) and peak velocity (right) respectively.

violations in VR that often exhibit a negative component followed by a strong positive deflection. Given the absence of a distinct event, this is not surprising because semantic violations could, in fact, appear at different points during the interaction, given the nature of HR around the noticeability level. We examined the peak negativity in the 150–250 ms window following salient moments of the movement (see Figure 6.21). We found a main effect of HR for FCz at both *peakAcceleration* ($\chi^2 = 15.7$, $p = .001$) and *peakVelocity* ($\chi^2 = 16.3$, $p < .001$). Similar main effects were observed at Cz for *peakAcceleration* ($\chi^2 = 32.5$, $p < .001$) and *peakVelocity* ($\chi^2 = 23.1$, $p < .001$). Post-hoc tests revealed significant differences between Base and Below ($p = .005$) as well as Base and Above ($p = .001$) at electrode Cz. There were no other significant differences between conditions after p -adjustments. While some differentiation between the baseline and HR conditions was observed, suggesting a trend towards increasingly larger peak amplitude from modest negativity at Base, all the way to the strongest at Above, the pattern was inconclusive.

The spectral equivalent to ERP is ERSP, which provides another view on the effects of HR around the DT. For this type of analysis, we included the Pz electrode because it is located above the parietal lobe of the brain, responsible for movement guidance. Figure 6.23 shows the grand average ERSP cluster permutation test against the pre-stimulus baseline. The black contours outline the significant activity in the spectral power, differing from the baseline. At electrode FCz an initial burst in theta power occurred with the onset of movement and lasted until the peak velocity was reached. We observed a desynchronization in the beta range between 20–30 Hz lasting throughout the movement phase. Interestingly, a synchronization between 35–40 Hz first appeared at maximum acceleration, lasting until the end of the trial. At Pz, the dominant spectral feature was a desynchronization in the alpha band, appearing at maximum acceleration and lasting until the end of the trial. Peak strength of the

*Suffer from lack of
temporal anchor*

*Only weak evidence
for ERPs*

*We measured the
perceptual effects of
HR*

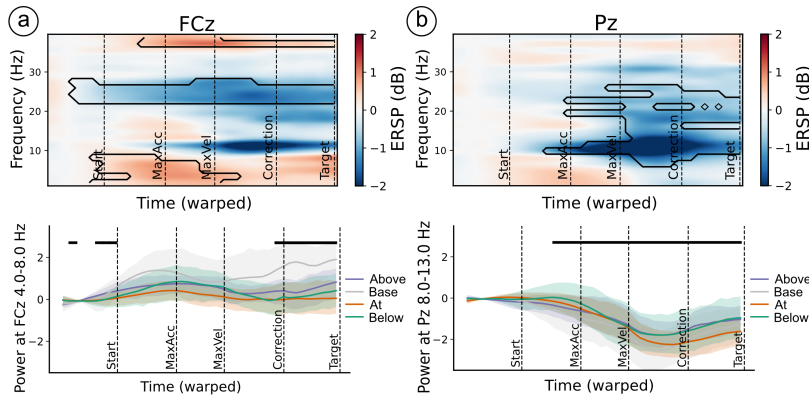


Figure 6.22: Top: Event-related spectral perturbations at electrodes FCz (a) and Pz (b). Changes in power from a -300 to -100 ms pre-stimulus baseline are marked by a black contour for significance. Bottom: Band power in theta (4–8 Hz) frequency range for electrode FCz (a) and alpha (8–13 Hz) for electrode Pz. Significant time points for the main effect *condition* are marked by black bars.




Figure 6.23: Ratio of power in theta (4–8 Hz) frequency range at electrode FCz divided by power in alpha (8–13 Hz) frequency range at electrode Pz. Significant time points for the main effect *condition* are marked by black bars.

desynchronization was between maximum velocity and the onset of the correction phase. Taken together, we consider these findings to be validations of the recorded data, clearly demonstrating task-related spectral dynamics. Post-hoc tests for theta and alpha bands showed significant differences between Base and all other conditions (see Figure 6.22).

Finally, we calculated the ratio of theta FCz and alpha Pz, which is commonly used to measure cognitive load [72, 122]. Similarly, we found a main effect and a post-hoc test showed a significant difference between Base and any other condition. **Nevertheless, we can only partially confirm (H3), because contrary to previous studies [118, 303], we did not observe the same distinct ERP signatures. However, our analysis of the peak error negativity showed promising results at the *peakAcceleration* event, especially at Cz. As a re-**

More severe HR can make the task for cognitively demanding

sult, *peakAcceleration* is an interesting marker for further exploration. Our ERSP analysis allowed us to distinguish between movements under the influence of HR from movements with no redirection applied. However, we did not find evidence that would suggest that, based on the presented features, we can easily differentiate between HR of different magnitudes.

6.2.3.3 Summary

Experiment elicit
robust perceptual
boundaries

Hand movement and
eye gaze data are
promising

In part 1 of the experiment, we established participants' perceptual boundaries Below, At, and Above their personal DTs using the method of constant stimuli. We verified that the thresholds obtained are perceptually different from each other and that the participants did not suffer from proprioceptive drift or fatigue in part 2 of the experiment. Our analysis showed that movement time and the transition points from ballistic to correction phase can be used to distinguish between all three perceptually different HR offsets (**H1**). Furthermore, participants looked significantly more often and longer at their virtual hand when no HR was applied than Above the DT (**H2**). Finally, ERP peak error negativity and the ERSP results showed great potential to detect the presence of HR, even at the unnoticeable Below level (**H3**).

6.2.4 Predicting Perceptual Boundaries

Predicting HR
magnitudes based on
a single hand
movement

To better understand the potential of our proposed method, we combine the three modalities: movement, gaze, and EEG by training a multimodal classifier. Here, our goal was to predict whether users were exposed to no HR vs. HR Below, At or Above their DTs based on a single trial. Unlike Si-Mohammed et al. [303], we did not perform a per-participant analysis but aggregated our collected samples into one data set, which we made publicly available to the community in the supplementary materials of the paper. This way, researchers can train their own models or formulate new research hypotheses. The data set contains 4352×77 data points.

6.2.4.1 Features

Following our statistical analysis, we used the features *totalTime*, *transitionPointTime*, *transitionPointDistance*, *#handFixations*, *durationHandFixations*, *Cz_amplitude_min*, *FCz_amplitude_min* and *FCz_theta/Pz_alpha_ratio*. Then, we normalized all features using z-scoring. Since *FCz_theta/Pz_alpha_ratio* is a continuous high-dimensional feature, we computed skewness, median, interquartile range (iqr), kurtosis, cumulated frequency (cumfreq_3), the 10th and the 90th quantile (quantile_10 and quantile_90).




Figure 6.24: Confusion matrix for 10-fold cross-validation of our multimodal classifier. The classifier can distinguish Base from any other condition. It appears that between movements under the influences of HR At and Below the DT are challenging to predict.

6.2.4.2 Training

We performed a 10-fold cross-validation by shuffling the data and splitting them in a stratified way, preserving the relative imbalance in the data set. Samples of each class in the train and the test set are removed until the data set is balanced. Then, we trained a classifier using Random Forest. We provide our model and code base in the supplementary materials of the paper.

6.2.4.3 Results

Our results show that without any optimization, we can achieve an overall accuracy of 40.682% and a mean F1 score of 39.359% at a theoretical probability of 25%, with a confusion matrix shown in Figure 6.24. We computed Combrisson and Jerbi [59]'s adjusted chance level of 36.184% at $p < .001$. This method takes the number of classes and samples into account, where if the accuracy is higher than the adjusted chance level, the result is statistically significant by a p -value. Since our classifier exceeds this probability, we can conclude that we can predict the correct class with an accuracy significantly higher than chance level. In particular, movements without HR can be correctly predicted with an accuracy of 63.2%. It appears that there is ambiguity between movements Below, At, and Above the DT. Here, the classifier performance seems rather weak under the influences of HR At (28.0%) and Below (29.2%) the DT, while Above can be predicted with an accuracy of 40.4%. The permutation feature importance depicted in Figure 6.25 suggests that no single feature dominated the prediction, but highlights that *totalTime*, *#handFixations*, *transitionPointDistance* and *durationHandFixations* are above average in terms of prediction power.

Without a sophisticated learning model and hyperparameter tuning

Can predict presence and absence of HR

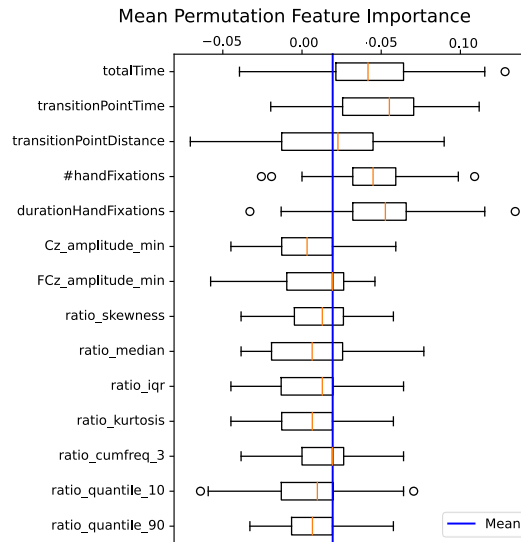


Figure 6.25: Shows mean permutation feature importance, suggesting that the features *totalTime*, *#handFixations*, *transitionPointDistance* and *durationHandFixations* have the most predictive power.

6.2.4.4 Summary

We trained a multimodal classifier using Random Forest and performed a 10-fold cross-validation, achieving an overall classification accuracy of about 40%. This is significantly higher than the adjusted chance level and demonstrates that movement, gaze, and EEG data collected Below, At or Above participants' individual DTs can be used to distinguish no HR from any other condition. However, for movements under the influence of HR Below, At and Above the DT, it is challenging to separate them from one another. All features contributed to the prediction of the perceptual limits, with *totalTime*, *transitionPointDistance*, *#handFixations* and *durationHandFixations* contributing most significantly. This marks a substantial step forward in tailoring illusions to individuals' perceptual boundaries.

*Method can predict
the correct class
significantly above
chance level*

6.2.5 Discussion

6.2.5.1 Predicting Perceptual Boundaries of HR

Considering all results, we provide evidence for being able to distinguish between hand movements with (At and Above) from without HR. Focusing on this binary classification problem would yield more impressive prediction accuracy, but it is insufficient to address our vision of predicting the perceptual boundaries of users. While we can differentiate At and Above from no HR, we acknowledge that our current model performs rather poorly between HR offsets of different magnitudes. Our results still support our research goal, because offsets Below can be safely used (practically

*Challenging to
differentiate between
HR of varying
magnitudes*

not crucial to detect), and appear to be indistinguishable from no HR, following our Bayesian features analysis. The effects start to occur when reaching At, but seem to be more prominent at Above.

We envision that users would start the VR experience with no offset and the system would increase it (Below) until it detects the unnoticeable limit (At, Above). Theoretically, it could dynamically adjust offsets within that range (Below) and detect potential shifts (At = no HR). However, based on our data, we cannot yet differentiate if HR offsets correspond to At or Above. Given the ambiguity in the Below condition, we would argue that achieving an overall accuracy of 40% for a 4-stage classifier on this type of data, without a sophisticated learning model and any hyperparameter tuning, is promising for a novel approach. We want to emphasize that the participants experienced horizontal HR offsets of different magnitudes in part 2 of the experiment and that our analysis was performed across all participants. Therefore, we can confidently say that our method is tailored to individuals' perception rather than fixed offset magnitudes. Most likely, better prediction accuracy can be achieved when training and evaluating on a per-participant level similar to Si-Mohammed et al. [303]; however, obtaining the necessary per-user data through a controlled psychophysical experiment in advance defeats the purpose of our method.

*Continuously monitor
and update HR
offsets*

*Training a classifier
for each individual*

It is important to note that the perceptual boundaries Below, At, or Above of the individual DTs should not be considered as distinct effects. For example, the average horizontal 25% threshold corresponds to an offset of 2.58 cm, while the 50% threshold was 3.56 cm at the target location of the hand, 30 cm in front of the participants. Thus, the type of effect always remains the same (i.e., virtual hand offset to the right), and only its magnitude changes. In light of this, achieving an overall accuracy of about 40% in single-trial classification, given the limited amount of data that can reasonably be collected in a psychophysical experiment, demonstrates the impressive potential of our approach. We expect that with more data, the robustness and prediction accuracy will improve further. Therefore, we recommend future work to build on our foundation and the data set provided, adding other types of redirection [381, 384] and interaction [17, 19, 133, 193], or without informing participants [26]. Hence, researchers can simply retrain our multimodal classifier using our resources, extending beyond our current setup to test its validity in diverse and more applied VR settings. To this end, there also seems to be a tendency towards movement and gaze features being the most effective. This opens up exciting opportunities because these are much easier to monitor than EEG data. However, it remains to be explored how these modalities perform in demanding VR scenarios.

*Results are
promising*

6.2.5.2 Generalizability & Limitations of the Method

*Validation in more
realistic and complex
IVES necessary*

*Considering
different types of
hand-based illusions*

As with any novel method, the main question remains whether our results are directly related to the effects of HR. To the best of our knowledge, we controlled for as many variables as possible to isolate potential effects by (1) calibrating individuals' DTs [144] and (2) verifying that the threshold did not drift over the course of the experiment [280]. In the next iteration of this research, we aim to investigate the robustness of the method in more realistic and complex IVEs. For example, to ensure that we operate at participants' perceptual boundaries, we used an established methodology [322, 384] and informed the participants about the presence of HR. As a result, some effects may be related to the procedure itself rather than the interaction under the influence of HR. For example, participants may have spent more time looking at their virtual hand when no HR was applied because they observed their hand closely to detect a potential offset, in contrast to the 75% Above threshold condition, where participants noticed the offset relatively early and therefore returned to their natural behavior, i.e., looking at the target [203]. Additionally, we used a specific type of HR (i.e., horizontal offsets to the right) at a fixed virtual distance and only looked at hand movements performed by participant's dominant hand. Thus, the generalizability of the method to bi-manual interactions [133], other HR algorithms [236, 261, 381, 385, 387], or greater visuo-haptic integration [80] remains to be explored.

*Awareness of HR
most likely affects
measures*

Furthermore, by informing the participants about the procedure and designing the task around HR detection, we used a very conservative approach. For our first exploration, this was needed to ensure comparability between participants, but it is far from any real VR experience. For example, Benda et al. [26] found that the detectability of HR differs greatly when participants are informed of its presence. The application of HR techniques without informing users is more realistic and practically relevant for VR design. As a result, much larger HR offsets can be used without disrupting the VR experience, which may help to improve the power of our method. Finally, understanding how other variables and VR interactions affect movement, gaze, and EEG features is crucial to assessing the potential of the method for constant monitoring in immersive VR experiences, going beyond substituting a threshold experiment. Ultimately, we rely on further research to validate our method.

6.2.5.3 Practicality and Utility of the Method

Our method relies on tracking hand movements, gaze, and participants' EEG. The first two measures can be monitored with most modern HMDs and do not require additional trackers, such as the one

we used in our experiment. The bottleneck of the method is the acquisition of EEG data because it is a time-consuming, tedious, and uncomfortable procedure for users. Calibrating a DT for one type of interaction takes about 10 min using a state-of-the-art DT experiment, which is equivalent to 25% of the time it took two experimenters to position gel-based EEG electrodes. However, this is only a hardware limitation because companies such as Galea⁶ offer HMDs with integrated physiological sensing capabilities for not just EEG, but also ECG and EDA. Ultimately, this would allow ubiquitous data collection inside IVEs, and in contrast to calibrating just a single DT for one type of interaction, it enables us to constantly monitor participants' perceptual sensitivity and adapt if necessary. In this way, the system could collect more data in varying environments, improving the robustness and overall accuracy. With this promising potential on the horizon, our investigation marks a significant step forward with serious implications for the broad spectrum of perceptual illusion techniques in VR, pushing toward immersive sensory experiences that feel indistinguishable from reality.

Physiological measures can be collected ubiquitously

6.2.6 Conclusion & Contributions

In this section, we mainly contributed to **RQ4** by investigating EEG, eye gaze, and interaction data to implicitly detect offsets corresponding to Below, AT and Above individuals' HR DTs based on a single hand movement (**C4**). This provides the foundation for our multimodal classifier, which ultimately aims to distinguish between HR of different magnitudes.

We conducted a 2-part experiment with 18 participants collecting EEG, eye gaze and, interaction data. First, we estimated participants' DTs for horizontal HR using the psychophysical method of constant stimuli. Our results contribute to **C3** by extending the previously determined DTs [144, 384] for this type of redirection technique.

Confirming previous DTs from literature

The results of part 2 suggest that, unlike Section 6.1, ERPs and their peak negativity show only anecdotal evidence to distinguish between different HR magnitudes. This may be a result of the absence of a distinct event, compared to [118, 265, 303] and our previous investigation during strong HR in Section 6.1. Our ERSP analysis allowed us to distinguish between movements under the influence of HR and movements with no redirection applied. However, we did not find evidence that would suggest that, based on the features presented, we can easily differentiate between HR of different magnitudes.

Movements under the influence of HR can be detected

6 Galea webpage: <https://galea.co/>. Last accessed: Nov 1, 2024

*Eye gaze and
movement features
hold great predictive
power*

In this experiment, we included eye gaze data because it is a fast-responding measure. The number of gazes intersecting the virtual hand and their duration was significantly lower during the presence of HR At and Above than without HR, while there was evidence for the absence of an effect between no HR and Below the DTs of individuals. Although we observed a numerical trend towards consistently lower number of intersects and their duration with increasing HR offsets (Figure 6.20), we did not find statistical evidence that would suggest that we can easily differentiate between HR At and Above the detectability level. In contrast, movement data, especially movement duration and transition points from ballistic to correction phase, can be used to distinguish between no HR and HR Below, At and Above the DT. This extends our findings on strong gain-based HR (Section 6.1) to a different type of redirection, demonstrating the features' predictive power, regardless of the applied technique.

*Multimodal
classifier to
dynamically tailor
HR offsets*

When combining the modalities through training a multimodal classifier using a basic Random forest without hyperparameter tuning, we achieved an overall prediction accuracy of about 40% for all four HR magnitudes. Given the ambiguity between Below and no HR, achieving an overall accuracy of 40% for a 4-stage classifier on this type of data, without a sophisticated learning model is promising for a novel approach. Considering all results, we provide evidence for the ability to distinguish between hand movements with HR (At and Above) from without HR. Although we can differentiate At and Above from no HR, we acknowledge that our current prediction model performs rather poorly between HR offsets of different magnitudes. We cannot yet differentiate if a HR offset corresponds to At or Above DT. To this end, we contribute a first step towards a method that dynamically tailors HR to individuals' perceptual boundaries, which would elevate the practical utility of the large landscape of illusions in VR—potentially even beyond hand-based techniques.

*Important first step,
but not done yet*

*Effort towards open
science*

Finally, we open source a data set containing EEG, eye gaze and hand movement features that the community can be built upon. In addition, we also provide the code base and learning parameter of our multimodal classifier (C6).

6.3 SUMMARY

In this chapter, our aim was to propose a novel method that can predict individuals' perceptual boundaries for a simple HR illusion using implicit user data (RQ4). To explore this, we conducted two experiments with a total of 40 participants that (1) investigate the potential of physiological and interaction data to differentiate

between noticeable gain-based and no HR, and (2) use the most promising measures to predict horizontal HR offsets that correspond to users' perceptual boundaries Below, AT and Above their DT. Our findings demonstrate the potential and limitations of the proposed method while providing a solid foundation for future work (C4).

Section 6.1 primarily investigates the effects of noticeable gain-based HR during continuous reaching movements on several physiological measures and movement data. In line with previous work on error prediction in VR [118, 303], we found ERPs and their peak negativity to correlate with the presence of strong HR. The transition points between ballistic and correction movement phases occurred significantly earlier under the influences of HR. Our exploratory data analysis on the ECG and RSP provided many insights, but ultimately, these types of data have limited utility for our research goal due to their slow-responding nature. Interestingly, we also found evidence for quick adaptation to strong positional hand offsets, only taking around 4–8 hand movements. Our results informed the design of the second experiment in Section 6.2 that attempts to bridge the gap between psychophysical threshold experiments and implicit measures.

*Bridging the gap
between threshold
experiments and
implicit measures*

Here, we determined users' individual sensitivity to gradual offsets between their virtual and real hand caused by horizontal HR. In contrast to our previous work, we not only determined a DT but rather modeled participants' discrimination performance across the entire spectrum of HR offsets by fitting a psychometric function. Hence, we select three offsets corresponding to Below (25%), AT (50%) and Above (75%) an individuals DT. In the second part of the experiment, we exposed participants only to offsets corresponding to their ability to detect them while collecting EEG, eye gaze, and movement data. Our ERP analysis could not confirm the predictive power of peak amplitudes for HR offsets around the detectability level. However, the ERSP analysis allowed us to distinguish between movements under the influence of HR from movements with no redirection applied. In contrast, movement features remain significantly different similar to our results in Section 6.1, even between HR offsets of different magnitudes. The added eye gaze features *#HandFixations* and *DurationHandFixation* showed a numerical trend between HR magnitudes and significant differences between no HR and Above.

*Physiological and
interaction data at
users' perceptual
boundaries*

In the next step, we trained a multimodal classifier, which achieved an overall prediction accuracy of about 40% for all four HR magnitudes. Our approach can effectively distinguish between hand movements with (At and Above) from without HR, but our current prediction model has limitations between HR offsets of different

*Predicting
perceptual
boundaries*

magnitudes. This work enhances the state of the art by being the first of its kind with the demonstrated potential of predicting HR offsets significantly above chance level without explicit user feedback.

As part of this chapter, we open-sourced two data sets, our code base for training the multimodal classifier, as well as raw data and analysis scripts to facilitate future work in the field, enabling the research community to more easily build upon our results (C6).

CONCLUSION

In this chapter, we summarize the most important parts of this dissertation in light of the major contributions stemming from our four research questions **RQ1–4**. We also outline the minor contributions that emerged from our investigations. Finally, we reflect on the overarching goal of this dissertation and provide recommendations for future research directions.

7.1 SUMMARY & DISSERTATION CONTRIBUTIONS

The goal of this dissertation is to improve haptic feedback for hand-based interactions with virtual objects through proxies in VR (**RQ1**), by applying undetectable hand-based illusion techniques (**RQ2**), understanding their application limits (**RQ3**) and tailoring them to individuals' perceptual boundaries (**RQ4**). Our motivation is grounded in the field of haptics in VR, which would ultimately allow Lisa to be fully immersed in a virtual world where she can practice with tools and experience a variety of scenarios first-hand in a safe and responsible way. To contribute to this vision, we conducted eleven lab studies with a total of 206 participants, built four proxy-based research prototypes and eight test beds for our experiments, and presented several practical use cases and applications that demonstrate the implication and relevance of our findings.

*Summary of
executed work*

From our analysis of the current landscape of haptic concepts and devices in our literature survey (Chapter 2), we learned that haptic devices often focus on a single modality of haptic feedback and lack direct embodiment to account for the wide variety of haptic sensations that occur during whole-hand exploration and manipulation of virtual objects. Thus, we first set out to investigate novel approaches for proxy design that address these shortcomings.

7.1.1 Major Contributions to RQ1

In Section 3.1.1, we explored reconfiguration as a method for proxy design by developing the *VoxelHap* toolkit. It gives full embodiment to virtual objects in the form of physical replicas, rendering tactile and kinesthetic haptic feedback for hand-based interactions (**C1**). *VoxelHap* utilizes elements from both passive and active haptics that can be used individually or in combination, making it a mixed haptic feedback device. Users can change the haptic properties of their proxies through manual offline reconfiguration by using a combination

*Manual proxy
reconfiguration to
change its haptic
properties*

*Design tool assists
users during proxy
construction process*

*VoxelHap provides
multimodal haptic
feedback*

of *Voxels* and *Plates*. This process is facilitated through our design tool based on the .hpdf. Our user studies showed that *VoxelHap* can provide a variety of tactile and kinesthetic haptic sensations while allowing direct and embodied whole-hand interactions (C1). With this, we go beyond most VR controllers, which are often limited to rendering one type of haptic feedback or lack full embodiment, which is essential to promote natural and intuitive interactions [91]. Thus, we stress the need for more holistic approaches to haptics. Compared to traditional VR controllers, *VoxelHap* proxies achieve higher shape and size similarity with virtual objects. In turn, this increased users' expectations, which may have caused them to become more sensitive to differences in the similarity between the virtual and proxy object. Our findings contribute to the theory of uncanny valley of haptics [29], suggesting that the more closely haptics in VR reassembles their real-world counterpart, the more sensitive users become to mismatches. In its current implementation, *VoxelHap* is limited to manual reconfiguration, which slows down its performance, preventing seamless interactions between multiple virtual objects.

*Compressed
3D-printed
metamaterials can
enhance tactile
perception*

In Section 3.2, we contribute a dynamic approach by using 3D-printed metamaterials with visual material overlays to simulate tactile experiences during one-finger exploration of a surface material (C1). We designed and fabricated metamaterial structures that are able to change their tactile surface properties, e.g., their hardness and roughness, upon lateral compression. Given the novelty of the approach for proxy-based haptic feedback, we contribute the design of five different metamaterial patterns. Our initial user study, which focused on haptic perception, identified two metamaterial designs that successfully convey different levels of roughness and (partially) hardness at varying levels of compression. We combined these metamaterial patterns with visual material overlays to investigate their potential as passive haptic proxies for simulating visuo-haptic material experiences in VR. We demonstrated that our metamaterial designs create different tactile states through compression that can communicate different tactile sensations. Our approach significantly enhances the state of the art, as it allows dynamic shape-change of the metamaterial upon compression. As a result, it can provide a wide range of tactile sensations without requiring additional fabrication or manual reconfiguration.

Research Question 1

How can we design proxies that can change their perceived kinesthetic and tactile properties?

Theoretical Contributions

- Concept of *VoxelHap*, a construction toolkit allowing users to build proxies for VR that provide tactile and kinesthetic haptic feedback for hand interactions.
- Investigation on how tactile and kinesthetic feedback through proxies affect users' perception of the VR interaction.
- Concept of 3D-printed metamaterial compression in combination with visual texture overlays to enhance tactile perception in VR.
- Investigation of users' perception of our metamaterials during visuo-haptic interaction.

Design Contributions

- Design of *VoxelHap* concept consisting of *Voxels* and *Plates*, the mechanical mechanisms for Base-, Rotation- and VibrationVoxel as well as Connection-, Texture-, Weight- and ShapePlates. Designs for the assembly workflow that uses *.hpdf* for proxy construction, interactions with *VoxelHap* proxies and tactile and kinesthetic renderings.
- Discussion of lessons learned for VR proxy design.
- Design of five metamaterial patterns with emerging tactile features upon compression that we open-sourced in a repository.

Technical Contributions

- Implementation and fabrication of *VoxelHap* prototype, including hardware consisting of low-cost microelectronics and actuators as well as low and high-level software. A visualization tool to assist proxy assembly, a test bed for two user evaluations, and a demonstration containing six interactive VR experiences. We open-source our materials, i.e., code base, 3D models, and schematics in a repository.
- Fabrication of five metamaterial patterns using conventional FDM 3D printing, a metamaterial actuation device, and a test bed for a user evaluation.

Together, our contributions pave the way for multipurpose proxies as metamaterials may be used to augment existing controllers and devices—even in combination with *VoxelHap*. Finally, we open-sourced all our materials to encourage future work to build upon our results (C6).

7.1.2 Major Contributions to RQ2

Although our proposed approaches advance the field of proxy-based haptic feedback for hand-based interactions, we noticed hardware limitations that cannot yet be easily overcome. To address this issue, we investigated software-based techniques where designers can trigger haptic feedback based on visual manipulations. In Chapter 4

Novel hand-based
illusion techniques

and Section 5.1, we established three novel undetectable illusion techniques: *Pseudo-Haptic Resistance*, *Visuo-Haptic Translation* and *Visuo-Haptic Rotation* through a series of rigorous psychophysical lab experiments. All techniques exploit the visual-dominance phenomenon by gradually offsetting visual from proprioceptive information to create the effects. Our experiments ensured that: (1) we can effectively trigger perceptual illusions and (2) users do not notice the manipulations of their interactions within the IVE. As a result, they help to overcome the physical limitations of proxy design without disrupting the VR experience (C2).

Illusions enhance
tactile and
kinesthetic haptic
feedback

In Section 4.1, we presented *Pseudo-Haptic Resistance*, which allows the system to change the perceived resistance of a knob upon rotational interactions, which can be achieved by visually slowing down users' real-world rotations of the knob. As a result, users need to cover more distance to reach the target location, which translates into the perception of more resistance. We also provide a model that describes the relationship between rotational C/D gain and perceived physical resistance. Thus, VR designers can more effectively simulate haptics for a variety of rotational interactions. Next, Section 4.2 introduced *Visuo-Haptic Translation*, enabling a single physical proxy slider to act as a stand-in for multiple virtual sliders of varying lengths. This can be achieved by visually scaling up or down users' real-world translations of the slider. Consequently, the haptic limits of virtual sliders are adequately represented, tricking Lisa's perception into believing that the proxy and virtual slider match. Here, we do not create a '*pseudo-force*' as per definition of Lécuyer [227], but instead present a way to alter the perceived properties of proxy to make it more versatile. *Visuo-Haptic Rotation* established in Section 5.1 demonstrates this principle for a different kind of interaction, rotation with a tool inside virtual containers of varying sizes embodied through a single container proxy.

As part of our investigation, we developed two fully functional hardware prototypes, a knob and slider (C1), that allow monitoring of the interaction through their integrated sensing capabilities, e.g., they can internally store the physical slider states for various virtual objects. Both the slider and knob can reset themselves to the state of the virtual object upon hand interaction. To illustrate the working principle of these novel techniques, we designed and implemented a total of five interactive VR experiences that demonstrate the possibilities of the presented techniques (C5).

Research Question 2

How can we design unnoticeable hand-based illusions that expand the scope of proxy-based interactions?

Theoretical Contributions

- Proposal of a novel concept of *Pseudo-Haptic Resistance*, allowing designers to create the perception of physical resistance based on C/D gain-based HR during rotational interactions.
- Investigation on the rotational C/D gain necessary to change the perceived physical resistance and to what extent C/D gains remain unnoticeable to users, including a first model on how rotational C/D gains translate to the perceived physical resistance.
- Proposal of a novel concept of *Visuo-Haptic Translation*, allowing designers to remap a single proxy slider to multiple virtual sliders of varying lengths by applying bidirectional linear C/D gains during interactions with the proxy slider.
- Extension of the *Visuo-Haptic Translation* concept to *Visuo-Haptic Rotation*. The technique allows designers to remap a single proxy container to multiple virtual containers of varying diameters by applying bidirectional rotational C/D gains.
- Quantification of the conservative lower bound for unnoticeable bidirectional, linear and rotational C/D gains.

Design Contributions

- Design of several IVEs, virtual objects, and interactions that deliver immersive VR experiences through our proposed illusion techniques in combination with proxies.
- Investigation of the side effects of applying VR illusion techniques that exploit visuo-proprioceptive mismatches, helping designers to make informed decisions.

Technical Contributions

- Development of a prototype knob and slider proxy with sensing and actuation capabilities using off-the-shelf low-cost hardware.
- Implementation of three psychophysical VR experiments to obtain DTs for rotational interactions, the minimum rotational C/D change required to create the sensation of physical resistance using the adaptive staircase methods, as well as the relationship between rotational C/D gain and perceived resistance using the method of adjustments.
- Development of three interactive VR experiences displaying our developed *Pseudo-Haptic Resistance* technique.
- Implementation of two psychophysical VR experiments to obtain estimates for the DTs of *Visuo-Haptic Translation* and *Visuo-Haptic Rotation* using the adaptive staircase method.
- Development of an interactive VR-DJ experience that combines *Pseudo-Haptic Resistance* and *Visuo-Haptic Translation* with on-the-fly HR [52] to resemble multiple virtual sliders and dials on a DJ-desk.

*Many factors affect
the unnoticeable
offset*

While we were studying the techniques, we found that the amount of visual-proprioceptive offset and, thus, the detectability of our hand-based illusions greatly differ between users. Moreover, factors such as the traveled distance had measurable effects on the unnoticeable offset. Thus, we used **RQ3** to question the robustness and generalizability of our, but also any other hand-based illusion technique presented in the literature. Already in Section 4.2, we found that in the presence of visuo-haptic redirection, movement, accuracy, and time change according to the magnitude of the redirection. These findings challenged existing movement prediction models that are crucial for planning and executing effective redirection [131]. This resulted in a new string of research, e.g., by Gonzalez and Follmer [134], who proposed a new model for the trajectory prediction of the hand during redirected reaching in VR. Given these promising results, we use these measures in Section 6.2.

7.1.3 Major Contributions to RQ3

*Haptic dominance
occurs when visual
stimulus becomes
implausible*

Building on our observations from the previous research questions, we contribute a thorough investigation of potential variables that influence the detectability and robustness of hand-based illusions (**C3**). In Section 3.2, we start by decoupling haptic (tactile) from visual information, which resulted in mixed material perception, allowing us to shift the perception of the hardness and roughness of a material. However, the visuo-haptic combinations were often implausible due to great mismatches (e.g., visual fabric and flat uncompressed metamaterial). As a result, we found that participants relied more on their haptic sense when facing large mismatches [157]. Interestingly, participants were also very sensitive to shape and size differences when interacting with *VoxelHap* proxies (Section 3.1.1). Although we did not quantify the effect with our experiments, we provided considerations for proxy design and highlighted the limits of visual-dominance during visual-haptic mismatches.

*Our results
generalize to other
VR experiences due
to their conservative
nature*

By introducing our novel hand-based illusion techniques for VR in Chapter 4, we make an attempt to quantify the effect of visuo-proprioceptive offsets, specifically gradual positional offsets. We applied psychophysical methods to determine conservative DT estimates for our techniques to give designers a lower bound that they can safely apply in VR interaction design without risking detection. These DTs were determined in very conservative settings, where participants were made aware of the technique and their task was to detect the offset as early as possible. Thus, any practical VR experience can rely on our reported DTs. However, we surprisingly found that DTs are significantly different for slider manipulations of different lengths. As a result, we wanted to understand what limits

or even extends the amount of unnoticeable offset to enhance the practical utility of illusion techniques. Therefore, we used Section 5.1 to study the potential effects of aspects fundamental to any interaction. For instance, we looked at whether how users hold objects may influence their proprioceptive accuracy. Our experiments were designed to identify magnitudes of practically relevant effects, i.e., the offset was large enough to create a tangible difference, as minor differences would be merely of theoretical nature. We quantify the unnoticeable offset for four different grasping types, varying object masses and three movement trajectories, which is of great value for VR designers who want to include unnoticeable hand-based illusion techniques. Our frequentist and Bayesian analysis of the DTs provide evidence for the absence of a tangible effect on the detectability of HR, contributing to the understanding of multisensory integration. However, we found that restricted linear movement (through a slider) allows for significantly greater offsets than unrestricted linear movement. This poses the question of whether additional proprioceptive cues caused by, e.g., haptic limits of sliders lead to greater visuo-haptic integration [262]. As a result, users are less likely to detect offsets because they rely more on their visual perception. However, the opposite effect was found when adding additional relative force-feedback through linear stretching in Chapter 4. Here, DTs significantly decreased during linear stretching compared to linear translation. Nevertheless, it must be noted that stretching provides a very strong distance cue and since there was no adequate visual rendering, participants may have relied more on the proprioception during sensory integration due to the implausible nature of the visualization inside the IVE.

We investigated a large range of factors relevant to any interaction with proxies

The brain integrates the sensory inputs, and vision tends to dominate

Given these findings about how proprioception affects users' sensitivity to offsets, in Section 5.2, we asked the question if only rendering a mid-air floating disconnected virtual hand influences the perceptual thresholds of hand-based illusion techniques. Up to this point, hand-based illusion techniques, including ours, that provide perceptual limits for the unnoticeable offset were established using this type of visualization. However, more prominent and complete avatars provide more visual cue anchoring and may, therefore, have a significant impact on the perception of virtual/real hand offsets. To investigate whether the results from previous studies hold up with the increasing trend towards more embodiment avatars, we designed four avatar completeness levels to study the effects of embodiment on the detectability of horizontal HR. Again, we quantify the amount of offset for this type of illusion, contributing towards a holistic picture of the extent to which sensory mismatches remain undetected.

Embodiment does not seem to practically affect the detectability of HR

Research Question 3

Which factors related to the user, interaction, and proxy extend or limit the unnoticeable offset during hand-based illusions?

Theoretical Contributions

- Quantification of the conservative HR DTs for various grasping and proxy types and weights, interactions and HR techniques.
- Investigation of the dominating factors during visuo-proprioceptive conflicts related to users' background, interaction and proxy.
- Conceptualisation of four avatar completeness levels: Hand, Hand+Arm, FullBody and FullBodyTexture.
- Insights into multisensory integration during visuo-haptic mismatches.
- Insights on the robustness of our VR illusions during visuo-proprioceptive conflicts in the presence of proprioceptive cues.
- Investigation of users' adaptation process to visuo-proprioceptive offsets.

Design Contributions

- Insights on the factors that limit or extend human sensitivity to visuo-proprioceptive mismatches, which helps design proxy objects, IVEs, and interactions, including hand-based illusions, with the aim to remain unnoticeable.
- Designs of IVEs, virtual objects, and interactions that illustrate the practical implications of our findings for VR design.

Technical Contributions

- Implementation of two psychophysical (+ experiments from novel techniques proposed in **RQ2** and **RQ4**) VR experiments to obtain estimates for the DTs of a wide variety of interactions using the adaptive staircase method.
- Development of an interactive VR cooking experience that demonstrates a design process with visuo-haptic techniques.

Our results show a higher sense of embodiment felt towards more complete avatars and provide strong evidence that the existing hand-based illusion techniques and their perceptual thresholds can be applied regardless of the avatar representation.

Next, in Section 6.1, our analysis of EEG, EDA, and movement data suggested that the experience with the interaction itself affects the detectability of HR. In our study, users quickly adapted to (larger) offsets within 4–8 interactions. This knowledge is of great value to any VR designers because it provides insights into how much time is required to develop full body ownership and suggests that the current calibration of the user is another factor to consider during ongoing VR applications (**C3**). Notably, it took participants 40–50 sec to complete eight reaching movements, which is comparable to

*Updating
sensorimotor
contingencies*

the previously reported onset time of the body ownership illusion [172]. It appears that users' initial movements serve the purpose of establishing a mental model of their environment, which could also explain why participants with previous VR experience had lower thresholds in Section 5.1, because they can rely on their mental model of the IVE.

Another aspect we explored through our studies is that while DTs seem to differ drastically, they do so in a systematic way, i.e., thresholds of individuals appear to be consistently high or low. Through our questionnaires, we looked at whether this relates to factors known to affect proprioceptive accuracy, such as participation in sports or users' age. However, our sample size and diversity of participants were severely limited. As a result, despite prior VR experience, our analysis did not reveal a relationship between any of the demographic information and participants' ability to detect visual-proprioceptive offsets. Our findings, alongside the continuous stream of research in the field, suggest a complex relationship between environment, task, interaction, and individuals' sensitivity. Therefore, we formulated **RQ4** with the ambiguous goal of establishing a novel method that allows continuous monitoring of users' current sensitivity to visual-proprioceptive offsets.

*Sensitivity to offsets
could be linked to
users' backgrounds*

7.1.4 Major Contributions to RQ4

Our contribution to this research question is two-fold. In Section 6.1 we first looked at the potential of physiological and interaction data to distinguish VR experiences with noticeable gain-based HR from experiences with no HR applied. Given the individual difference in the ability to detect visuo-proprioceptive hand offsets, we tailored this experiment to each individual user by determining participants' DT for gain-based HR using a psychophysical interleaved staircase procedure. In the main part of the experiment, participants performed a continuous reaching task with and without personalized HR offsets. We collected physiological (EEG, ECG, EDA, RSP) and interaction data that, given the novelty of the approach, were analyzed using a mix of statistical and exploratory research methods.

*Novel approach to
tailor experiments to
individuals'
sensitivity to offsets*

We contribute an investigation of ERPs and their peak negativity, showing that they are reliable in detecting strong visuo-proprioceptive offsets, making them a valid candidate for further explorations in Section 6.2. Similarly, movement phases were directly affected by noticeable hand offsets. Next, we contribute a first investigation of ECG and RSP in relation to an HR illusion, and we recommend indices that may be studied in further experiments. While these physiological measures may work for separating condi-

*Distinguish
experiences with and
without noticeable
HR*

tions with vs. without noticeable HR, their slow-responding nature makes them unsuitable for our research goal **RQ4** because we rely on measures that immediately respond according to user perception of the interaction in the IVE.

*A novel method to
hand-based illusions
of different
magnitudes*

*Gaze and movement
features worked best*

In the second experiment Section 6.2, we model participants' perceptual ability to discriminate movements with and without different magnitudes of horizontal HR using the psychophysical method of constant stimuli. Through this approach, we determine the perceptual boundaries Below, At, and Above their DT. Based on our previous experiment in Section 6.1, we only included EEG, interaction data and added eye gaze due to its quick responding nature. Participants then repeatedly performed reaching movements while being exposed to no HR or HR offsets corresponding to their perceptual boundaries. Our analysis suggests that, unlike Section 6.1, ERPs and their peak negativity showed only a weak trend between different HR magnitudes, which may be a result of the absence of a distinct event compared to the literature [118, 265, 303]. Our ERSP analysis and the number of gazes intersecting the virtual hand allowed us to distinguish between movements under the influence of HR from movements with no HR applied. Moreover, movement duration and transition points from ballistic to correction phase can be used to distinguish between HR at the perceptual boundaries. This extends our findings on strong gain-based HR (Section 6.1) to a different type of redirection, validating our previous contributions and confirming the features' predictive power, regardless of the applied HR technique (**C4**).

*Detecting perceptual
boundaries by
combining
modalities*

On our mission to answer **RQ4**, we then combined the modalities through training a multimodal classifier using a random forest ML model without any hyperparameter tuning. We contribute the first classifier in the literature that attempts to distinguish a hand-based illusion technique at varying magnitudes based on a single exposure. We achieved an overall prediction accuracy of about 40% for all four HR magnitudes, which is promising for a 4-stage classifier on this type of data without a sophisticated learning model. Although we can differentiate HR offsets from no HR, we acknowledge that our current prediction model needs further refinement for HR offsets of different magnitudes. Our results still support (**RQ4**) because noticeable HR offsets can be detected significantly above chance level (**C4**).

Research Question 4

How can we tailor unnoticeable hand-based illusions to individual users' perceptual boundaries?

Theoretical Contributions

- Investigation of physiological (EEG, ECG, EDA, RSP) and interaction data to distinguish VR experiences that apply strong gain-based HR vs. no HR.
- Proposal of a novel method, combining EEG, eye gaze, and interaction data to distinguish between horizontal HR Below, At, and Above individuals' DTs based on a single interaction.

Design Contributions

- Insights on the potential of EEG, ECG, EDA, RSP, and interaction data to determine noticeable HR.
- Insights on the potential of leveraging EEG, eye tracking, and interaction data to detect visuo-proprioceptive offsets, which can ultimately help adapt HR offsets to individual sensitivity in an *on-the-fly* fashion.

Technical Contributions

- Implementation of a 2-part psychophysical VR experiment to (1) obtain estimates for individuals' gain-based HR DTs using the adaptive staircase method and (2) collect physiological and interaction data under noticeable HR.
- Open-source a data set containing EEG, ECG, EDA, RSP, and interaction features during the presence and absence of noticeable HR.
- Implementation of a psychophysical VR experiment to model participants' perceptual sensitivity to horizontal HR offset using the method of constant stimuli.
- Open-source a data set with EEG, eye gaze, and interaction data collected at individuals' perceptual boundaries for horizontal HR on a single movement basis.
- Train and open-source a multimodal classifier to separate hand movements at different HR magnitudes corresponding to individuals' perceptual boundaries.

Finally, conducting this type of research requires significant resources (expensive equipment and time), which may prevent research labs from contributing to this promising domain. To facilitate future explorations, e.g., by applying more sophisticated learning methods or running different types of analysis, we contribute two data sets containing all collected features from both experiments in Section 6.1 and Section 6.2 that the research community can be built upon. In addition, we also provide the code base and learning parameter of our multimodal classifier to increase transparency (C6).

Contributing to open science movement

7.1.5 Further System Contributions

Throughout our investigations, we developed several prototypes that helped achieve our research goals but did not directly contribute to the main research questions. We would like to dedicate this section to two special contributions that received great attention in the HCI/VR community and the public. Both contributions stem from our general motivation to facilitate XR research and make illusion techniques in VR accessible to a wider audience.

VRQuestionnaire-Toolkit and MoVRI



VRQuestionnaireToolkit

Videos about the work and prototypes presented in this chapter are available online and can be accessed through the QR codes. Images and parts of the text in this chapter, as well as the presented figures, tables, ideas, concepts, implementations, applications and uses cases, studies and experiments, results, discussions, and conclusions, have been published previously in:

[94] **Martin Feick**, Niko Kleer, Anthony Tang, and Antonio Krüger. *In Proceedings of ACM UIST 2020 Adjunct*. The Virtual Reality Questionnaire Toolkit.



MoVRI

[102] **Martin Feick**, André Zenner, Simon Seibert, Oscar Javier Ariza Nunez, David Wagmann, Juliana Helena Keller, Anton Wittig and Antonio Krüger. *Demo at Saarland Informatics Campus 2023*. MoVRI: The Museum of Virtual Reality Illusions.

7.1.5.1 VRQuestionnaireToolkit

Enabling the community to collect measures in virtual environments

First, we would like to highlight our contributions with the *VRQuestionnaireToolkit*, which enables the research community to easily collect subjective measures within IVEs *in situ*. Even though researchers execute experiments in IVEs, they mostly rely on paper-based or screen-based questionnaires. Thus, participants need to exit the virtual world to self-report their experience. Putze et al. [276] and Schwind et al. [299] found that removing subjects from the IVE before filling out questionnaires introduces a systematic bias in

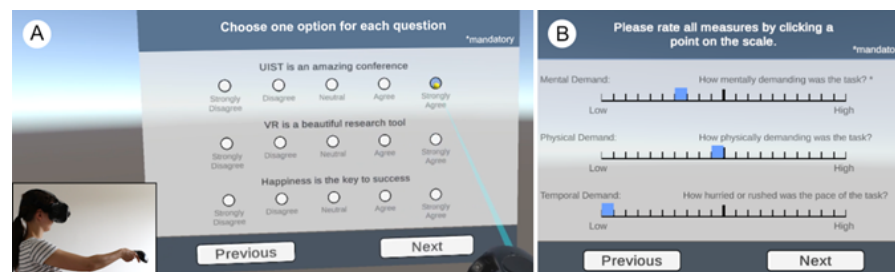


Figure 7.1: VRQuestionnaireToolkit's UI: User selects radio button (A) using pointer techniques. (B) NASA TLX using LinearGrid.

their subjective responses due to them experiencing a break in presence. Therefore, we contribute a highly customizable and reusable open-source toolkit in the form of a Unity3D asset that can be easily integrated into existing Unity3D projects. Our questionnaire design (see Figure 7.1) incorporates many findings from 3D interaction design and is, therefore, the first open-source toolkit whose design is based on solid academic ground. For instance, we chose pointing as a selection technique and world-anchoring rather than using a trackpad and body-anchoring [7]. The toolkit comes with a pre-installed set of standard questionnaires such as NASA TLX, SSQ, and SUS Presence questionnaire, but can be easily customized and extended. Our system aims to lower the entry barrier to use questionnaires in VR and to significantly reduce the development time and cost needed to run pre-, in-between- and post-study questionnaires. To this day, we still support and improve the toolkit, which proved to be an extremely valuable technical contribution to the research community.

Comes with a set of standardized questionnaires

We originally developed the *VRQuestionnaireToolkit* to conduct our own research in this dissertation, but it quickly outgrew this purpose and became a widely used tool in the community. Since it is a Unity3D asset, it can be used with a variety of VR platforms, such as HTC Vive or Meta Quest. Additionally, other researchers reported using it in mobile applications to conduct remote studies during the COVID-19 pandemic or even integrated it in AR to work with Microsoft HoloLens [170]. At this point, renaming the toolkit to *XRQuestionnaireToolkit* would better describe its broad application space.

The VRQuestionnaireToolkit is used beyond VR

7.1.5.2 MoVRI—The Museum of VR Illusions

Second, we would like to present *MoVRI—The Museum of VR Illusions*. Many VR illusion techniques, including the ones presented in this dissertation, have been documented and published in research papers. But it is often difficult and cumbersome for people, in particular those new to VR, to understand and imagine how the techniques work and how the effects feel. An easy way to try out




Figure 7.2: *MoVRI* showcases four illusion techniques, *Impossible Spaces* [330], *Redirected Touching* [193], *Haptic Retargeting* [17, 52], and *Pseudo-Haptic Weight* [288] that can be experienced in different exhibition rooms. The system can be calibrated to the physical space available, with larger spaces allowing for more realistic museum tours.

*Allowing people to
experience illusions
firsthand*

famous VR illusion techniques does not exist today and existing demos are usually specialized on only a single technique. It is this gap that *MoVRI* fits in by allowing users interested in VR illusion techniques to explore different virtual “exhibition rooms”, each demonstrating and explaining a different VR illusion technique. As a result, *MoVRI* empowers novices and VR experts to experience VR illusions firsthand that would otherwise remain inaccessible.

*MoVRI is an
education tool that
not only shows but
also explains illusion*

To achieve this, each illusion technique is implemented in two ways, *experience* and *explanation* mode, because often, the purpose of VR illusions is to remain unnoticeable for users. In experience mode, a user performs interactions by playing a mini-game with, for instance, an unnoticeable hand redirection illusion, while in explanation mode, the real and the warped virtual hand position alongside a textual (and audio) explanation is displayed in the VR environment (see Figure 7.3). How this happens depends on the experienced illusion. For example, since many of our exhibited techniques utilize real/virtual hand offsets, we implemented a live-streaming of the HMD’s front-facing camera into the virtual environment. This allows visitors to compare their real hand position with the warped virtual hand position (see Figure 7.3). In addition, the illusion techniques are then explained by a museum guide through text and audio.

*MoVRI uses a plot
twist; the museum
itself uses an illusion*

The clue about *MoVRI* is that exploring the museum works with an illusion, suggesting a much larger virtual space than physically available—which will be revealed to the visitor at the very end of the museum tour. The museum consists of exhibition rooms and an elevator to travel to different floors (or levels). An aisle around the room can be used to access up to three exhibition rooms on one floor. The *MoVRI* uses a modular approach, i.e., we offer a template room that can be used when including a new illusion technique. *MoVRI* can be tailored to the physical space available by running through a calibration routine; however, the setup works best if more physical space is available—smaller setups prevent compelling illusions and may trigger claustrophobia. To this end, we provide three different types of rooms: (1) the foyer, where each visitor starts; (2) placeholder rooms that can be filled with any content; and (3) exhibition rooms that showcase VR illusions. To demonstrate many of the well-established hand-based illusion techniques, proxies are needed [256]; these are often placed on furniture. To prevent visitors from bumping into such requisites, which are only relevant for specific illusions, *MoVRI* uses placeholders such as virtual fountains, plants, or couches [304]. During calibration, placeholders are automatically scaled to the required size. *MoVRI* is customizable with respect to physical constraints and room layout.

*MoVRI offers great
scalability*




Figure 7.3: *Haptic Retargeting* explanation mode shows real (top) and warped virtual (bottom) hand by using the HMD's front camera. Hand movements with (A) one-to-one mapping and (B) under the influences of horizontal HR.

We displayed *MoVRI* on several occasions and received overwhelmingly positive feedback, especially from users who have not had the chance to experience illusions in VR before. The museum was also used to educate university and high school students on the topic and hopefully sparked their interest and inspired them to pursue their studies or even careers in this research field. Our long-term vision for *MoVRI* is that due to its modular approach, it can act as a platform to share novel VR illusion techniques within the community or during the review process to promote transparency and, thus, contribute to better research practices.

MoVRI received considerable attention during public demonstrations

7.2 FUTURE WORK & CONCLUSION

We already proposed and discussed ideas for future work in the individual project sections of this dissertation. Here, we provide additional high-level recommendations and research directions for future work before concluding this dissertation.

7.2.1 Future Directions for Proxy Design

With *VoxelHap*, we focused on addressing the challenge of *Similarity* [256] with the idea of having a one-to-one mapping, i.e., a replica, between the properties of the proxy and the virtual object, for the most relevant virtual objects inside an IVE. However, this cannot be scaled to all virtual objects inside an IVE—which would result in users having to physically reconstruct the IVE, defeating the purpose of VR [158]. With our approach to 3D-printed metamaterial actuation, we contribute to multipurpose proxies, specifically focusing on touch interactions. Through this, we promoted more interdisciplinary research with domains such as soft robotics, fabrication, material sciences, and shape-changing devices. Our work on compressing 3D-printed metamaterials to enhance tactile perception demonstrates that these research fields offer great potential for the development of novel haptic devices. We also recommend further research looking into how we can leverage everyday objects that are already present in

Use advancements from interdisciplinary fields

*Exploit ubiquitous
devices that we
already use*

a user's environment as proxies [63, 150]. With the constant advancements in computer vision technology, we also believe that analyzing and/or predicting the physical properties of users' surroundings can be used to dynamically create and update mappings for proxy-based haptics and various kinds of interactions. This is especially relevant since our environments have become more ubiquitous, and smart devices that already have sensing and actuation capabilities provide an inexpensive way for interaction design with proxies [229].

*Combine software-
and hardware-based
approaches*

To address the problem of *Collocation* [256], we used existing redirection techniques in combination with our stationary slider and knob proxies in Chapter 4. Matthews et al. [237] built on our results by presenting the first fully-integrated proxy interface with multiple sliders, buttons and knobs, seamlessly redirecting users to the appropriate proxy given the virtual target control. Here, *Haptic Retargeting* allowed these proxies to act as stand-ins for multiple virtual sliders and knobs with varying properties. However, the technique also underlies the same perceptual limits and thus, it cannot be scaled up indefinitely, limiting the interaction space. Here, concepts such as robotic graphics [240] or encounter-type haptics [132] can help to overcome the limitations of purely visual techniques and stationary proxies. Proxy interfaces could quickly reposition themselves based on the predicted target in 3D space [375]. Yet, the virtual target prediction is another research area that needs further improvement to enhance proxies' application space. To this day, this is mostly done by applying heuristics based on eye gaze, distance, and trajectory from the hand to the virtual object [52, 237]. Research in this direction should be extended to prediction models based on semantic information about the environment and motion models [57, 375] to effectively combine hand-based illusion with proxies.

*Understanding
users' intentions is
key to illusion*

*Approaches to
haptics in VR can
complement each
other*

Our approach to address the physical limitations and expand the scope of proxies-based interactions relied on visual manipulations of users' interactions inside the IVE. However, there is a large body of work that could provide great alternatives or additions to the use of hand-based illusions with proxies. For example, devices worn on the fingertip that can render high-resolution tactile feedback [301] but lack kinesthetic haptic rendering capability would greatly complement one another. Similarly, body-worn devices such as EMS [225] or TES [15] could help to address the limited force feedback capabilities of smaller actuators that can be found in most proxy objects. Finally, a somewhat controversial direction is to draw inspiration from the field of neurohaptics. Here, researchers apply motor cortex stimulation to trigger haptic sensations through chip implants [268]. While these invasive approaches are still in their early days and are primarily applied to treat medical conditions such as

post-stroke pain or phantom pain, they may become relevant for VR research in the future. Their noninvasive technological counterparts, transcranial direct current stimulation and TMS have already been used to stimulate local brain areas. For example, Tanaka et al. [336] presented a proof-of-concept prototype using TMS to create various haptic sensations, such as tapping and pressing onto users' hands and feet. Yet this approach lacks resolution and conveys limited haptic sensations, but it highlights an interesting direction for future work to enhance haptic interaction with and without proxies.

7.2.2 Future Directions for Hand-based Illusions

Together with ongoing work in this field, we contributed three novel hand-based illusion techniques for VR. However, as a growing research field, there is still room for innovative techniques that can simulate haptic sensation as discussed in Chapter 4 and Chapter 5. Here, the combination with advancements in the field of proxies appears to be a promising direction [132, 372, 386]. Yet, many techniques focus on manipulating users' visual perception but ignore other senses, such as auditory. For example, we believe that sound renderings that correspond to the interaction and IVE are worth studying [40]. Furthermore, there is great potential in investigating and especially combining multisensory manipulations to achieve novel perceptual effects. Through this, researchers could also increase the perceived realism, which may help to hide the presence of illusion techniques as suggested by Ogawa et al. [262].

*Beyond
manipulating the
visual sensory
modality*

In this dissertation, we provide very conservative lower bounds for the range of unnoticeable offsets; however, there exists a trade-off between detectability and effectiveness. For example, larger offsets are desirable because they result in stronger perceptual effects. Thus, during our studies, researchers already explored ways to increase DT by using tendon stimulation [261], exploiting change-blindness effects through blinks [385] and saccades [381] or introducing distracting factors [83, 384]. There is great potential in exploring similar techniques with other technologies such as EMS [225], MMS [337] or even by including other sensory modalities. Also, we would argue that there is a need for a conceptualization of the space through, for instance, a validated taxonomy that investigates sensory modalities and ways to stimulate them according to the interaction, as well as their potential to hide offsets. As a result, VR designers would have access to an effective toolbox that, depending on the environment, interaction, and the user, seamlessly applies undetectable hand-based illusions in VR experiences. Another interesting direction could be looking at how hand-based illusion techniques can become part of

*How can we
effectively hide
offsets?*

the interaction design itself. For example, to mechanically reconfigure a proxy without users even noticing.

*Address limitations
of current HR
algorithms*

This shows the need for new and innovative redirection algorithms that meet these demands. For example, similar to the existing shape-aware HR [236], physics-aware [338] HR could be used to create realistic behavior during hand interactions with virtual objects and proxies inside the IVE while preserving high levels of presence. There is also still a need to improve the horizontal HR algorithm, as the yaw rotation of the virtual hand is distorted. This is a result of users aiming for a virtual target with their hand rotated towards the target. Instead of updating the yaw rotation of the hand while reaching toward the target under the influence of horizontal HR, it feels natural to maintain the initially established hand rotation. This becomes problematic when grasping or interacting with virtual objects embodied by a physical proxy as users' virtual hand approaches the proxy from a different angle than they anticipate [17]. To this day, it still remains an unaddressed challenge and as a result, we mostly opted for gain-based HR in our studies. To address this problem, we recommend future work to collect real-world grasping data and apply target prediction models to compensate for the discrepancy to facilitate grasping and interacting with virtual objects embodied by a proxy. Finally, an interesting direction is to extend the idea of *Retargeted Self-Haptics* [87], redirecting hands to the user's own body to create haptic feedback. This could be scaled up to multiple users, similar to what has been done with *TurkDeck* [53] and proxy-based haptic interaction. Ultimately, systems could leverage the diversity of users' clothes, jewelry, and the devices they carry to simulate touch sensations; however, this also raises ethical and privacy concerns that must be considered.

*Expand to other
body parts*

Another interesting direction to explore is how redirection techniques can be used for other body parts, such as the user's head, trunk, or feet. Especially the latter could be used for foot-based interfaces such as *Kickables* [291], but also for traditional interactions such as the operation of pedals or during sports activities such as climbing [195] or playing soccer. Alongside extending the efforts towards bi-manual hand illusions [133], researchers could also look into illusions with multiple hands [71] or bodies [244], which open up new exciting interaction paradigms. However, one limitation might be that while users' hands are often in sight, other body parts usually remain in the peripheral view, which limits the application space of vision-based illusion techniques. This is also a general limitation of the techniques presented in this dissertation because they only work if users pay attention to their hands during interactions. That being said, we expect that once body ownership is established and

users adapt to offsets in SCs, they can effectively use their virtual embodiment even without continuous visual feedback.

In our studies in Chapter 4 and Chapter 5, we used static hand models, augmented users' hands with tracking marker and capacitive sensing to ensure high agency and visual-haptic integration [343]. In recent years, more sophisticated hand-tracking technology based on computer vision algorithms has emerged. While early approaches for hand and finger tracking relied on heavy marker augmentation, HMDs such as the Meta Quest 2 only require built-in cameras to achieve robust and accurate hand tracking. This will facilitate the use of hand-based illusion techniques, increase accessibility, and thus pave the way to become more practical in VR interaction design.

*Advancements in
hand tracking
technology*

We presented several use cases and applications for our developed proxies and hand-based illusion techniques in Chapter 3, Chapter 4 and Chapter 5. Here, we primarily focused on the domains of training, entertainment, and simulation; however, there exist several other application domains that can greatly benefit from hand-based illusion techniques. For example, in our ongoing collaboration in the rehabilitation area, visuo-proprioceptive offsets in VR can be used to correct movement trajectories or improve patients' range of motion and balance according to the prescribed training protocol by the physiotherapists. Another application domain could be simulating environments where movements '*feel*' differently, e.g., underwater or in space. We believe there exist many unexplored application domains that may have a high impact on society, and we recommend that future work carefully consider scenarios outside of the presented scope.

*Identify new
application domains*

7.2.3 Individualizing Hand-based Illusions

Over the course of this dissertation, we observed that users' expectations and sensitivity to mismatches differ, and despite providing useful estimates for HR DT, there does not appear to be a one-fits-all solution. On our mission to personalize VR experiences, we focused on hand-based illusions to gather insights into how users experience visuo-proprioceptive offsets in VR. However, the results presented in Chapter 6 can only be seen as a first step because the classifier's performance needs to be improved, and validation of the approach beyond a controlled lab environment is missing. The lack of data sets, in particular, made it challenging to validate our method. Thus, we strongly advocate extending the efforts toward open science, especially considering the resources needed to collect such data. Only when joining scientific efforts can we push the frontiers of hand-based illusions for VR.

*On the quest to
continuous
hand-based illusions
tailored to
individuals*

*Reaching
undetectability is
often not necessary*

*Can Weber fractions
explain sensitivity to
jumps in offset?*

We would also like to highlight that undetectability is often not necessary. While it may be crucial for VR training and simulation applications to ensure that the acquired motor skills transfer into the real-world [164, 257], VR has the superpower to allow users to go beyond what is real [3]. Here, the most important aspect is to not disrupt the immersive experience. For this, the context or plot [165] in which body illusions occur plays a key role. For example, imagine Lisa is playing a VR game embodied by an avatar that has extendable arms. Thus, in this situation, there is nothing surprising about extended reach, even though it is far from *realistic*. In fact, because she adapted to it, it would be disruptive to the experience if suddenly, she could not interact with distant objects anymore. This is also why many interaction techniques, such as the *Go-Go* technique [272], allows users to maintain high body ownership without impacting presence in the IVE. Consequently, the question becomes how much we can diverge from the user's current model of the world without causing disruption. Our investigations considered the real world as the reference, but we believe that it is crucial to extend our results to varying contexts in future work. For example, we could imagine that visuo-proprioceptive hand offsets could behave similarly to the theory of Weber fractions [259], where the maximum unnoticeable change in offset from one movement to the next movement is proportional to the current offset. This could be validated through a psychophysical experiment similar to the ones outlined in this dissertation, providing many insights into the adaptation of hand offsets during continuous redirection. To this end, we only focused on noticeability because it is well-defined in the psychophysical literature with the scientific tools available to measure it. Furthermore, while unnoticeability ensures undisruptiveness, we acknowledge that it severely limits the kinds of interactions and *awesome* things users can do and experience in VR. As a result, our investigations are through a conservative lens, and we strongly recommend future work to go beyond this.

7.2.4 Long-Term Exposure to Hand-based Illusions

*Deteriorated
proprioception may
harm users*

One of our major concerns while developing our hand-based illusion techniques is their potential impact on human proprioceptive abilities. Considering these effects is of great importance since poor proprioception can be harmful. For example, dysfunction of proprioception results in impaired balance and is one of the leading causes of many incidents, such as a higher possibility of falls [108]. When exposing participants to visual hand offsets, they quickly adapted to it (Section 6.1) and as a result, their real-world movements did not correspond to what they were visually experiencing. While this might be unproblematic for a short time period, for example, in the context of our studies, we were wondering if there exist any long-term

effects on users' proprioceptive accuracy. This is of great importance as, in the future, the ratio between time spent in the real vs. virtual world may shift, e.g., because of new professions that may emerge from the continuous movement towards the Metaverse¹. Suddenly, people will be spending their entire working day in IVEs. Therefore, understanding the potential long-term effects of perceptual illusions, even beyond hand-based techniques, is crucial, as this would greatly impact the safety of users. We recommend future work look into this topic, also for other types of body illusions, such as redirected walking, as manipulating the vestibular system could have serious consequences [348]. Nevertheless, our proposed method (C4) could also be used bidirectionally, i.e., users could be brought back to their real-world SCs, and once we detect that they recalibrated to this state, they can leave the virtual environment. This way, the method could be considered as a form of reality mediator. We demonstrated the potential of hand-based illusions, but we only want to promote their use if designers, users, and systems can apply them in an ethically correct, i.e., a safe and responsible way. This is relevant not only to the field of illusions but also to the whole domain of haptic feedback in VR. As we advance the field and it becomes more realistic, future research should consider the potential physical and mental harm it may cause and look at ways how this can be prevented [58].

It is important to consider the consequences of long-term exposure

Travelling between realities

7.2.5 Conclusion

In this dissertation, we set out to *explore the boundaries of hand-based illusions to enhance haptics in VR* with the help of four core research questions. We contributed four proxy-based approaches that rely on manual reconfiguration and dynamic actuation to change proxies' haptic properties. We demonstrated that in combination with visual overlays, our proxies can effectively produce tactile and kinesthetic haptic effects during hand interactions with virtual objects. Our user studies showed that this positively affected user experience inside IVEs. Given the physical limitations of hardware devices, we investigated software-based approaches, creating haptic effects only by applying visuo-proprioceptive offsets during hand interactions with proxies. Here, we contributed three novel hand-based illusions that can simulate resistance during rotations of a virtual knob, change the perceived length of virtual sliders and the perceived diameter of virtual containers. We studied how much offset between the visual and proprioceptive sensory modality can go unnoticeable while still enabling the system to change users' perceived haptics. The extent to which these hand-based illusions remained unnoticeable depends on several (unknown) factors. Thus, we conducted a systematic analysis of factors related to user representation, interactions, and properties of the

¹ Meta webpage: <https://tinyurl.com/37jz3ctc>. Last accessed: Nov 1, 2024

proxy, studying their role in visuo-proprioceptive conflicts through a series of psychophysical experiments. With this, we contributed many insights but, most importantly, demonstrated that DTs have limited practical utility. They vary significantly, even throughout a VR experience, which urges the need for an alternative approach to adapt hand-based illusions to the environment, context, and users' sensitivity to those offsets. Hence, we proposed a novel method that utilizes implicit physiological and interaction data with the aim of predicting if users were exposed to HR of different magnitudes, corresponding to their individual sensitivity. Our results demonstrate the potential of the approach, an important step into an unexplored territory. Finally, this dissertation also contributes towards transparency, accessibility, and reproducibility in science because we open-source a large part of the developed hardware, software, and two data sets.

LIST OF FIGURES

Figure 1.1	Illusion Techniques for Virtual Reality	4
Figure 1.2	Overview of Dissertation	16
Figure 2.1	VR History	18
Figure 2.2	VR Headsets	19
Figure 2.3	CAVE and HMD for VR	20
Figure 2.4	Reality-Virtuality Continuum	21
Figure 2.5	Olfactory and Gustatory in VR	27
Figure 2.6	Sensorimotor Contingencies Model	31
Figure 2.7	Examples of Optical Illusions	33
Figure 2.8	Grasp Type Taxonomy	35
Figure 2.9	Bayesian Inference during Visuo-Proprioceptive Conflicts	37
Figure 2.10	Visual-Haptic Reality-Virtuality Continuum . .	39
Figure 2.11	Examples of World-Grounded Active Haptics	41
Figure 2.12	Examples of Body-Grounded Active Haptics .	42
Figure 2.13	Examples of Ungrounded Active Kinesthetic Haptics	43
Figure 2.14	Examples of Ungrounded Active Tactile Haptics	44
Figure 2.15	Invisible Haptics	44
Figure 2.16	Examples of Emerging Haptic Concepts	45
Figure 2.17	Examples of Passive Haptic Feedback	46
Figure 2.18	Challenges of Colocation and Similarity	47
Figure 2.19	Examples of Dynamic Passive Haptic Feedback	49
Figure 2.20	Examples of Encounter-type Haptics	51
Figure 2.21	Encounter-Type Haptics at Room-Scale	52
Figure 2.22	Reconfiguration of Proxies	55
Figure 2.23	Fabrication of Proxies	57
Figure 2.24	Illustration of Hand Redirection	61
Figure 2.25	Examples of Pseudo-Haptic Effects	63
Figure 2.26	Examples of Redirected Touching Techniques .	65
Figure 2.27	Examples of Haptic Retargeting	66
Figure 2.28	Examples of Redirection Techniques	67
Figure 2.29	Detection Thresholds in Different Settings . . .	70
Figure 2.30	Psychophysical Methods	71
Figure 2.31	Brain Areas and EEG Eletrode Setup	74
Figure 2.32	Event-Related Potentials in VR	75
Figure 2.33	Hand Movement Velocity in VR	76
Figure 3.1	VoxelHap Toolkit	79
Figure 3.2	Overview of a BaseVoxel	80
Figure 3.3	Schematics of a BaseVoxel	81
Figure 3.4	3D-printed Shell with PCB	82

Figure 3.5	VoxelHap ShapePlates	83
Figure 3.6	Texture-, Weight- and ConnectionPlate	84
Figure 3.7	Assembly Process of VoxelHap Proxies	85
Figure 3.8	VoxelHap Rattle Proxy	86
Figure 3.9	VoxelHap Dial Lock	86
Figure 3.10	VoxelHap Salt Shaker	87
Figure 3.11	Illusions Enhance VoxelHap	87
Figure 3.12	VoxelHap Mug	88
Figure 3.13	Record Vibrations from Object Interactions	89
Figure 3.14	VoxelHap Experiment 1 Results	94
Figure 3.15	VoxelHap Experiment 2 Results	96
Figure 3.16	VoxelHap Experiment 1 and 2 Comparison	96
Figure 3.17	User Preference and Vibrotactile Feedback	97
Figure 3.18	Metamaterial with Visual Textures Overlays	101
Figure 3.19	Metamaterial Deformation Simulation	103
Figure 3.20	Metamaterial Designs 1 to 5	104
Figure 3.21	Metamaterial Actuation Setup	106
Figure 3.22	Digital Rendering of the Experimental Setup	110
Figure 3.23	Results Haptic Roughness and Hardness	111
Figure 3.24	Results Haptic Baseline	116
Figure 3.25	Results Visual Baseline	117
Figure 3.26	Results Haptic Baseline vs. Visuo-Haptics	118
Figure 3.27	Visual Baseline vs. Visuo-Haptic Comparison	119
Figure 3.28	Average Similarity for each Metamaterial	120
Figure 3.29	Results Similarity Ratings	120
Figure 3.30	Results Metamaterials + Visual Overlays	121
Figure 4.1	Pseudo-Haptic Resistance	129
Figure 4.2	Displays Proxy Knob	131
Figure 4.3	Experimental Threshold Task	134
Figure 4.4	Visualization of Method of Adjustments	135
Figure 4.5	Results from Threshold Experiments	138
Figure 4.6	Configured Physical Resistance for C/D Gains	139
Figure 4.7	Bayes Factors and P-Values for C/D Gains	140
Figure 4.8	Relationship Torque and C/D gain	142
Figure 4.9	Interactive VR-DJ Application	142
Figure 4.10	Steering Wheel and Door Handle Application	143
Figure 4.11	Visuo-Haptic Translation	146
Figure 4.12	C/D Effects on Translation and Stretching	147
Figure 4.13	Example of Interleaved-Staircase	149
Figure 4.14	Visuo-Haptic Translation Study Setup	150
Figure 4.15	Visuo-Haptic Translation Study Task	151
Figure 4.16	Results CDT for Translation and Stretching	154
Figure 4.17	Individual Differences	155
Figure 4.18	Effect of C/D Completion Times	156
Figure 4.19	Relationship Virtual Distance and Time	157
Figure 4.20	Extension of VR-DJ Application	158

Figure 5.1	Factors Affecting Hand-Based Illusions	165
Figure 5.2	Display Four Grasping Types	167
Figure 5.3	Visuo-Haptic Translation and Rotation	168
Figure 5.4	Experiment: Grasp, Motion and Proxy	169
Figure 5.5	Grasp and Motion: Experiment 1 Apparatus	171
Figure 5.6	Effect of C/D Ratio Manipulations	172
Figure 5.7	Grasp and Motion: Experiment 1 Results	174
Figure 5.8	Grasp and Proxy: Experiment 2 Task	175
Figure 5.9	Grasp and Proxy: Experiment 2 Results	177
Figure 5.10	Restricted vs. Unrestricted Linear Movement	179
Figure 5.11	Interactive Cooking Application	181
Figure 5.12	Avatar Completeness and Detectability of HR	185
Figure 5.13	Avatars Body Poses Participants Adopted	187
Figure 5.14	Example of a Threshold Procedure	188
Figure 5.15	Perceived Embodiment for Avatar Levels	191
Figure 5.16	CDT of Avatar Completeness Levels	192
Figure 5.17	CDT and Perceived Embodiment	192
Figure 6.1	Physiological and Interaction Data during HR	199
Figure 6.2	Study Conditions and Counterbalancing	202
Figure 6.3	Distribution of DTs for Gain-based HR	208
Figure 6.4	Avatar Embodiment Questionnaire	208
Figure 6.5	Mean ERP Amplitudes for Conditions	209
Figure 6.6	Negativity Peak Amplitudes	209
Figure 6.7	Time and Distance to Reach Target	212
Figure 6.8	Normalized Movement Phases	213
Figure 6.9	EDA Analysis	214
Figure 6.10	Varying RSP Responses from Participants	215
Figure 6.11	Dimensional Reduction RSP	215
Figure 6.12	PCA for Respiration	216
Figure 6.13	Dimensional Reduction ECG	217
Figure 6.14	Overview of Investigated Features	222
Figure 6.15	Effect of Horizontal HR and Events	224
Figure 6.16	Discrimination Performance of Participants	228
Figure 6.17	Threshold Distribution for Horizontal HR	229
Figure 6.18	Consistency in Discrimination Performance	230
Figure 6.19	Analysis of Movement Data	230
Figure 6.20	Analysis of Eye Gaze Data	231
Figure 6.21	ERP Amplitudes and Minima	232
Figure 6.22	Analysis of Spectral EEG Features	233
Figure 6.23	Analysis of Spectral EEG Ratio	233
Figure 6.24	Confusion Matrix for Multimodal Classifier	235
Figure 6.25	Feature Importance of Multimodal Classifier	236
Figure 7.1	The VRQuestionnaireToolkit	254
Figure 7.2	The Museum of VR Illusions	255
Figure 7.3	Haptic Retargeting Inside MoVRI	257

LIST OF TABLES

Table 1	Summary of VoxelHap's Technical Evaluation	91
Table 2	CDT for Translation and Stretching	153
Table 3	Grasp Type Classification	167

ABBREVIATIONS

.hpdf	Haptic Proxy Description Format
1AFC	1-Alternative-Forced Choice
2AFC	2-Alternative-Forced Choice
2IFC	2-Interval-Forced Choice
AHF	Active Haptic Feedback
AR	Augmented Reality
C/D	Control/Display
CAVE	Cave Automatic Virtual Environment
CDT	Conservative Detection Threshold
DoF	Degrees of Freedom
DPHF	Dynamic Passive Haptic Feedback
DT	Detection Threshold
EMS	Electrical Muscle Stimulation
ECG	Electrocardiogram
EDA	Electrodermal Activity
EEG	Electroencephalogram
ERP	Event-Related Potential
ETHF	Encounter-Type Haptic Feedback
FoV	Field of View
HCI	Human-Computer Interaction
HMD	Head-Mounted Display
HR	Hand Redirection
HRV	Heart Rate Variability
IVE	Immersive Virtual Environment
JND	Just-Noticeable-Difference
ML	Machine Learning
MMS	Magnetic Muscle Stimulation
PCA	Principal Component Analysis
PHF	Passive Haptic Feedback
RSP	Respiratory Rate
SCL	Skin Conductance Level
SCR	Skin Conductance Response

SCs	Sensorimotor Contingencies
SSQ	Simulator Sickness Questionnaire
TES	Tendon Electrical Stimulation
TMS	Transcranial Magnetic Stimulation
VR	Virtual Reality

BIBLIOGRAPHY

- [1] Muhammad Abdullah, Minji Kim, Waseem Hassan, Yoshihiro Kuroda, and Seokhee Jeon. "HapticDrone: An Encountered-Type Kinesthetic Haptic Interface with Controllable Force Feedback: Initial Example for 1D Haptic Feedback." In: *Adjunct Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST'17 Adjunct. ACM, 2017. DOI: [10.1145/3131785.3131821](https://doi.org/10.1145/3131785.3131821).
- [2] Parastoo Abtahi and Sean Follmer. "Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173724](https://doi.org/10.1145/3173574.3173724).
- [3] Parastoo Abtahi, Sidney Q. Hough, James A. Landay, and Sean Follmer. "Beyond Being Real: A Sensorimotor Control Perspective on Interactions in Virtual Reality." In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI'22. ACM, 2022. DOI: [10.1145/3491102.3517706](https://doi.org/10.1145/3491102.3517706).
- [4] Parastoo Abtahi, Benoit Landry, Jackie Yang, Marco Pavone, Sean Follmer, and James A. Landay. "Beyond The Force: Using Quadcopters to Appropriately Objects and the Environment for Haptics in Virtual Reality." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300589](https://doi.org/10.1145/3290605.3300589).
- [5] Merwan Achibet, Benoît Le Gouis, Maud Marchal, Pierre-Alexandre Léziart, Ferran Argelaguet, Adrien Girard, Anatole Lécuyer, and Hiroyuki Kajimoto. "FlexiFingers: Multi-Finger Interaction in VR Combining Passive Haptics and Pseudo-Haptics." In: *Proceedings of the 2017 IEEE Symposium on 3D User Interfaces*. 3DUI'17. IEEE, 2017. DOI: [10.1109/3DUI.2017.7893325](https://doi.org/10.1109/3DUI.2017.7893325).
- [6] Adilzhan Adilkhanov, Amir Yelenov, Ramakanth Singal Reddy, Alexander Terekhov, and Zhanat Kappasov. "VibeRo: Vibrotactile Stiffness Perception Interface for Virtual Reality." In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 2785–2792. DOI: [10.1109/LRA.2020.2972793](https://doi.org/10.1109/LRA.2020.2972793).
- [7] Dmitry Alexandrovsky, Susanne Putze, Michael Bonfert, Sebastian Höffner, Pitt Michelmann, Dirk Wenig, Rainer Malaka, and Jan David Smeddinck. "Examining Design Choices of Questionnaires in VR User Studies." In: *Proceedings of the 2020*

- CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376260](https://doi.org/10.1145/3313831.3376260).
- [8] Haneen Alsuradi, Wanjoo Park, and Mohamad Eid. "Mid-frontal Theta Power Encodes the Value of Haptic Delay." In: *Scientific Reports* 12.1 (2022). 8869. DOI: [10.1038/s41598-022-12911-0](https://doi.org/10.1038/s41598-022-12911-0).
 - [9] Judith Amores, Anna Fuste, and Robert Richer. "Deep Reality: Towards Increasing Relaxation in VR by Subtly Changing Light, Sound and Movement Based on HR, EDA, and EEG." In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI EA'19. ACM, 2019. DOI: [10.1145/3290607.3311770](https://doi.org/10.1145/3290607.3311770).
 - [10] Karl Andersson. *Manipulating Control-Display Ratios in Room-Scale Virtual Reality (Master's Thesis)*. KTH Royal Institute of Technology, 2017.
 - [11] Gabrielle E. Ansems, Trevor J. Allen, and Uwe Proske. "Position Sense at the Human Forearm in the Horizontal Plane during Loading and Vibration of Elbow Muscles." In: *The Journal of Physiology* 5763.2 (2006), pp. 445–455. DOI: [10.1113/jphysiol.2006.115097](https://doi.org/10.1113/jphysiol.2006.115097).
 - [12] Bruno Araujo, Ricardo Jota, Varun Perumal, Jia Xian Yao, Karan Singh, and Daniel Wigdor. "Snake Charmer: Physically Enabling Virtual Objects." In: *Proceedings of the 10th ACM Conference on Tangible, Embedded and Embodied Interaction*. TEI'16. ACM, 2016. DOI: [10.1145/2839462.2839484](https://doi.org/10.1145/2839462.2839484).
 - [13] Ferran Argelaguet, Ludovic Hoyet, Michaël Trico, and Anatole Lécuyer. "The Role of Interaction in Virtual Embodiment: Effects of the Virtual Hand Representation." In: *Proceedings of the 2016 IEEE Virtual Reality*. VR'16. IEEE, 2016. DOI: [10.1109/VR.2016.7504682](https://doi.org/10.1109/VR.2016.7504682).
 - [14] Ferran Argelaguet, David Antonio Gómez Jáuregui, Maud Marchal, and Anatole Lécuyer. "Elastic Images: Perceiving Local Elasticity of Images through a Novel Pseudo-Haptic Deformation Effect." In: *ACM Transactions on Applied Perception* 10.3 (2013). 17. DOI: [10.1145/2501599](https://doi.org/10.1145/2501599).
 - [15] Oscar Javier Ariza Nunez, André Zenner, Frank Steinicke, Florian Daiber, and Antonio Krüger. "Holitouch: Conveying Holistic Touch Illusions by Combining Pseudo-Haptics With Tactile and Proprioceptive Feedback During Virtual Interaction With 3DUIs." In: *Frontiers in Virtual Reality* 3 (2022). 879845. DOI: [10.3389/frvir.2022.879845](https://doi.org/10.3389/frvir.2022.879845).

- [16] Jatin Arora, Aryan Saini, Nirmita Mehra, Varnit Jain, Shwetank Shrey, and Aman Parnami. "VirtualBricks: Exploring a Scalable, Modular Toolkit for Enabling Physical Manipulation in VR." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300286](https://doi.org/10.1145/3290605.3300286).
- [17] Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. "Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI'16. ACM, 2016. DOI: [10.1145/2858036.2858226](https://doi.org/10.1145/2858036.2858226).
- [18] Yuki Ban, Takashi Kajinami, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. "Modifying an Identified Curved Surface Shape using Pseudo-Haptic Effect." In: *Proceedings of the 2012 IEEE Haptics Symposium*. HAPTICS'12. IEEE, 2012. DOI: [10.1109/HAPTIC.2012.6183793](https://doi.org/10.1109/HAPTIC.2012.6183793).
- [19] Yuki Ban, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. "Modifying an Identified Position of Edged Shapes using Pseudo-Haptic Effects." In: *Proceedings of the 18th ACM Symposium on Virtual Reality Software and Technology*. VRST'12. ACM, 2012. DOI: [10.1145/2407336.2407353](https://doi.org/10.1145/2407336.2407353).
- [20] Yuki Ban, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. "Displaying Shapes with Various Types of Surfaces Using Visuo-Haptic Interaction." In: *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology*. VRST'14. ACM, 2014. DOI: [10.1145/2671015.2671028](https://doi.org/10.1145/2671015.2671028).
- [21] Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. "TeslaTouch: Electro-vibration for Touch Surfaces." In: *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*. UIST'10. ACM, 2010. DOI: [10.1145/1866029.1866074](https://doi.org/10.1145/1866029.1866074).
- [22] Alejandro Beacco, Ramon Oliva, Carlos Cabreira, Jaime Gallego, and Mel Slater. "Disturbance and Plausibility in a Virtual Rock Concert: A Pilot Study." In: *Proceedings of the 2021 IEEE Virtual Reality and 3D User Interfaces*. VR'21. IEEE, 2021. DOI: [10.1109/VR50410.2021.00078](https://doi.org/10.1109/VR50410.2021.00078).
- [23] Robert J. van Beers, Daniel M. Wolpert, and Patrick Haggard. "When Feeling is More Important than Seeing in Sensorimotor Adaptation." In: *Current Biology* 12.10 (2002), pp. 834–837. DOI: [10.1016/s0960-9822\(02\)00836-9](https://doi.org/10.1016/s0960-9822(02)00836-9).
- [24] Iseult A. M. Beets, Marc Macé, Raf L. J. Meesen, Koen Cuypers, Oron Levin, and Stephan P. Swinnen. "Active versus Passive Training of a Complex Bimanual Task: Is Prescriptive Proprio-

- ceptive Information Sufficient for Inducing Motor Learning?" In: *PLoS One* 7.5 (2012). e37687. DOI: [10.1371/journal.pone.0037687](https://doi.org/10.1371/journal.pone.0037687).
- [25] Brett Benda, Shaghayegh Esmaili, and Eric D. Ragan. "Determining Detection Thresholds for Fixed Positional Offsets for Virtual Hand Remapping in Virtual Reality." In: *Proceedings of the 2020 IEEE International Symposium on Mixed and Augmented Reality*. ISMAR'20. IEEE, 2020. DOI: [10.1109/ISMAR50242.2020.00050](https://doi.org/10.1109/ISMAR50242.2020.00050).
- [26] Brett Benda, Benjamin Rheault, Yanna Lin, and Eric D. Ragan. "Examining Effects of Technique Awareness on the Detection of Remapped Hands in Virtual Reality." In: *IEEE Transactions on Visualization and Computer Graphics* 30.5 (2024), pp. 2651–2661. DOI: [10.1109/TVCG.2024.3372054](https://doi.org/10.1109/TVCG.2024.3372054).
- [27] Yoav Benjamini and Yosef Hochberg. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." In: *Journal of the Royal Statistical Society: Series B (Methodological)* 57.1 (1995), pp. 289–300. DOI: [10.1111/j.2517-6161.1995.tb02031.x](https://doi.org/10.1111/j.2517-6161.1995.tb02031.x).
- [28] Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. "NormalTouch and TextureTouch: High-Fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers." In: *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology*. UIST'16. ACM, 2016. DOI: [10.1145/2984511.2984526](https://doi.org/10.1145/2984511.2984526).
- [29] Christopher C. Berger, Mar Gonzalez-Franco, Eyal Ofek, and Ken Hinckley. "The Uncanny Valley of Haptics." In: *Science Robotics* 3.17 (2018). eaar7010. DOI: [10.1126/scirobotics.aar7010](https://doi.org/10.1126/scirobotics.aar7010).
- [30] Joanna Bergström, Aske Mottelson, and Jarrod Knibbe. "Re-sized Grasping in VR: Estimating Thresholds for Object Discrimination." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST'19. ACM, 2019. DOI: [10.1145/3332165.3347939](https://doi.org/10.1145/3332165.3347939).
- [31] Ilias Bergström, Sérgio Azevedo, Panos Papiotis, Nuno Saldanha, and Mel Slater. "The Plausibility of a String Quartet Performance in Virtual Reality." In: *IEEE Transactions on Visualization and Computer Graphics* 23.4 (2017), pp. 1352–1359. DOI: [10.1109/TVCG.2017.2657138](https://doi.org/10.1109/TVCG.2017.2657138).
- [32] Carlos Bermejo, Lik Hang Lee, Paul Chojeki, David Przewozny, and Pan Hui. "Exploring Button Designs for Mid-Air Interaction in Virtual Reality: A Hexa-Metric Evaluation of Key Representations and Multi-Modal Cues." In: *Proceedings*

- of the ACM on Human-Computer Interaction 5.EICS (2021). 194. DOI: [10.1145/3457141](https://doi.org/10.1145/3457141).
- [33] Dustin Beyer, Serafima Gurevich, Stefanie Mueller, Hsiang-Ting Chen, and Patrick Baudisch. "Platener: Low-Fidelity Fabrication of 3D Objects by Substituting 3D Print with Laser-Cut Plates." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI'15. ACM, 2015. DOI: [10.1145/2702123.2702225](https://doi.org/10.1145/2702123.2702225).
 - [34] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. "Design and Fabrication of Materials with Desired Deformation Behavior." In: *ACM Transactions on Graphics* 29.4 (2010). 63. DOI: [10.1145/1778765.1778800](https://doi.org/10.1145/1778765.1778800).
 - [35] Matthew Botvinick and Jonathan Cohen. "Rubber Hands 'Feel' Touch that Eyes See." In: *Nature* 391.6669 (1998). 756. DOI: <https://doi.org/10.1038/35784>.
 - [36] Elodie Bouzbib, Gilles Bailly, Sinan Haliyo, and Pascal Frey. "Can I Touch This?: Survey of Virtual Reality Interactions via Haptic Solutions: Revue de Littérature des Interactions en Réalité Virtuelle par le biais de Solutions Haptiques." In: *Proceedings of the 32nd Conference on l'Interaction Homme-Machine*. IHM'21. ACM, 2022. DOI: [10.1145/3450522.3451323](https://doi.org/10.1145/3450522.3451323).
 - [37] Elodie Bouzbib, Claudio Pacchierotti, and Anatole Lécuyer. "When Tangibles Become Deformable: Studying Pseudo-Stiffness Perceptual Thresholds in a VR Grasping Task." In: *IEEE Transactions on Visualization and Computer Graphics* 29.5 (2023), pp. 2743–2752. DOI: [10.1109/TVCG.2023.3247083](https://doi.org/10.1109/TVCG.2023.3247083).
 - [38] Virginia Braun and Victoria Clarke. "Thematic Analysis." In: *APA Handbook of Research Methods in Psychology, Vol 2: Research Designs: Quantitative, Qualitative, Neuropsychological, and Biological*. APA Handbooks in Psychology. American Psychological Association, 2012. Chap. 4, pp. 57–71. DOI: [10.1037/13620-004](https://doi.org/10.1037/13620-004).
 - [39] Andrew J. Bremner and Charles Spence. "The Development of Tactile Perception." In: *Advances in Child Development and Behavior* 52 (2017), pp. 227–268. DOI: [10.1016/bs.acdb.2016.12.002](https://doi.org/10.1016/bs.acdb.2016.12.002).
 - [40] Jean-Pierre Bresciani, Marc O. Ernst, Knut Drewing, Guillaume Bouyer, Vincent Maury, and Abderrahmane Kheddar. "Feeling What You Hear: Auditory Signals can Modulate Tactile Tap Perception." In: *Experimental Brain Research* 162.2 (2005), pp. 172–180. DOI: [10.1007/s00221-004-2128-2](https://doi.org/10.1007/s00221-004-2128-2).

- [41] Willem-Paul Brinkman, Allart RD Hoekstra, and René van Egmond. "The Effect of 3D audio and other Audio Techniques on Virtual Reality Experience." In: *Studies in Health Technology and Informatics* 219 (2015), pp. 44–48. DOI: [10.3233/978-1-61499-595-1-44](https://doi.org/10.3233/978-1-61499-595-1-44).
- [42] Jack Brookes, Matthew Warburton, Mshari Alghadier, Mark Mon-Williams, and Faisal Mushtaq. "Studying Human Behavior with Virtual Reality: The Unity Experiment Framework." In: *Behavior Research Methods* 52.2 (2020), pp. 455–463. DOI: [10.3758/s13428-019-01242-0](https://doi.org/10.3758/s13428-019-01242-0).
- [43] Jas Brooks, Noor Amin, and Pedro Lopes. "Taste Retargeting via Chemical Taste Modulators." In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. UIST'23. ACM, 2023. DOI: [10.1145/3586183.3606818](https://doi.org/10.1145/3586183.3606818).
- [44] Jas Brooks, Shan-Yuan Teng, Jingxuan Wen, Romain Nith, Jun Nishida, and Pedro Lopes. "Stereo-Smell via Electrical Trigeminal Stimulation." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI'21. ACM, 2021. DOI: [10.1145/3411764.3445300](https://doi.org/10.1145/3411764.3445300).
- [45] Ian M. Bullock, Joshua Z. Zheng, Sara De La Rosa, Charlotte Guertler, and Aaron M. Dollar. "Grasp Frequency and Usage in Daily Household and Machine Shop Tasks." In: *IEEE Transactions on Haptics* 6.3 (2013), pp. 296–308. DOI: [10.1109/TOH.2013.6](https://doi.org/10.1109/TOH.2013.6).
- [46] Saniye Tugba Bulu. "Place Presence, Social Presence, Co-Presence, and Satisfaction in Virtual Worlds." In: *Computers & Education* 58.1 (2012), pp. 154–161. DOI: [10.1016/j.compedu.2011.08.024](https://doi.org/10.1016/j.compedu.2011.08.024).
- [47] Eric Burns, Sharif Razzaque, Abigail T. Panter, Mary C. Whitton, Matthew R. McCallus, and Frederick P. Brooks. "The Hand is Slower than the Eye: A Quantitative Exploration of Visual Dominance over Proprioception." In: *IEEE Proceedings Virtual Reality 2005*. IEEE, 2005. DOI: [10.1109/VR.2005.1492747](https://doi.org/10.1109/VR.2005.1492747).
- [48] Eric Burns, Sharif Razzaque, Abigail T. Panter, Mary C. Whitton, Matthew R. McCallus, and Frederick P. Brooks. "The Hand Is More Easily Fooled than the Eye: Users Are More Sensitive to Visual Interpenetration than to Visual-Proprioceptive Discrepancy." In: *Presence: Teleoperators and Virtual Environments* 15.1 (2006), pp. 1–15. DOI: [10.1162/pres.2006.15.1.1](https://doi.org/10.1162/pres.2006.15.1.1).
- [49] Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. "UltraHaptics: Multi-point Mid-air Haptic Feedback for Touch Surfaces." In: *Proceedings of the 26th*

- Annual ACM Symposium on User Interface Software and Technology*. UIST'13. ACM, 2013. DOI: [10.1145/2501988.2502018](https://doi.org/10.1145/2501988.2502018).
- [50] Géry Casiez, Daniel Vogel, Ravin Balakrishnan, and Andy Cockburn. "The Impact of Control-Display Gain on User Performance in Pointing Tasks." In: *Human-Computer Interaction* 23.3 (2008), pp. 215–250. DOI: [10.1080/07370020802278163](https://doi.org/10.1080/07370020802278163).
 - [51] Lung-Pan Cheng, Sebastian Marwecki, and Patrick Baudisch. "Mutual Human Actuation." In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST'17. ACM, 2017. DOI: [10.1145/3126594.3126667](https://doi.org/10.1145/3126594.3126667).
 - [52] Lung-Pan Cheng, Eyal Ofek, Christian Holz, Hrvoje Benko, and Andrew D. Wilson. "Sparse Haptic Proxy: Touch Feedback in Virtual Environments Using a General Passive Prop." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.3025753](https://doi.org/10.1145/3025453.3025753).
 - [53] Lung-Pan Cheng, Thijs Roumen, Hannes Rantzsch, Sven Köhler, Patrick Schmidt, Robert Kovacs, Johannes Jasper, Jonas Kemper, and Patrick Baudisch. "TurkDeck: Physical Virtual Reality Based on People." In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*. UIST'15. ACM, 2015. DOI: [10.1145/2807442.2807463](https://doi.org/10.1145/2807442.2807463).
 - [54] Inrak Choi, Heather Culbertson, Mark R. Miller, Alex Olwal, and Sean Follmer. "Grabity: A Wearable Haptic Interface for Simulating Weight and Grasping in Virtual Reality." In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST'17. ACM, 2017. DOI: [10.1145/3126594.3126599](https://doi.org/10.1145/3126594.3126599).
 - [55] Inrak Choi and Sean Follmer. "Wolverine: A Wearable Haptic Interface for Grasping in VR." In: *Adjunct Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology*. UIST'16 Adjunct. ACM, 2016. DOI: [10.1145/2984751.2985725](https://doi.org/10.1145/2984751.2985725).
 - [56] Inrak Choi, Yiwei Zhao, Eric J. Gonzalez, and Sean Follmer. "Augmenting Perceived Softness of Haptic Proxy Objects through Transient Vibration and Visuo-Haptic Illusion in Virtual Reality." In: *IEEE Transactions on Visualization and Computer Graphics* 27.12 (2021), pp. 4387–4400. DOI: [10.1109/TVCG.2020.3002245](https://doi.org/10.1109/TVCG.2020.3002245).
 - [57] Aldrich Clarence, Jarrod Knibbe, Maxime Cordeil, and Michael Wybrow. "Unscripted Retargeting: Reach Prediction for Haptic Retargeting in Virtual Reality." In: *Proceedings of the 2021 IEEE Virtual Reality and 3D User Interfaces*. VR'21. IEEE, 2021. DOI: [10.1109/VR50410.2021.00036](https://doi.org/10.1109/VR50410.2021.00036).

- [58] Gaëlle Clavelin, Jan Gugenheimer, and Mickael Bouhier. "Potential Risks of Ultra Realistic Haptic Devices in XR." In: *Proceedings of the Workshop on Novel Challenges of Safety, Security and Privacy in Extended Reality at the ACM Conference on Human Factors in Computing Systems*. SSPXR'22. ACM, 2022.
- [59] Etienne Combrisson and Karim Jerbi. "Exceeding Chance Level by Chance: The Caveat of Theoretical Chance Levels in Brain Signal Classification and Statistical Assessment of Decoding Accuracy." In: *Journal of Neuroscience Methods* 250 (2015), pp. 126–136. DOI: [10.1016/j.jneumeth.2015.01.010](https://doi.org/10.1016/j.jneumeth.2015.01.010).
- [60] Tom Cornsweet. *Visual Perception*. Academic press, 2012. DOI: [10.1016/B978-0-12-189750-5.X5001-5](https://doi.org/10.1016/B978-0-12-189750-5.X5001-5).
- [61] Carolina Cruz-Neira, Daniel J Sandin, Thomas A DeFanti, Robert V Kenyon, and John C Hart. "The CAVE: Audio Visual Experience Automatic Virtual Environment." In: *Communications of the ACM* 35.6 (1992), pp. 64–73. DOI: [10.1145/129888.129892](https://doi.org/10.1145/129888.129892).
- [62] Mark R. Cutkosky. "On Grasp Choice, Grasp Models, and the Design of Hands for Manufacturing Tasks." In: *IEEE Transactions on Robotics and Automation* 5.3 (1989), pp. 269–279. DOI: [10.1109/70.34763](https://doi.org/10.1109/70.34763).
- [63] Florian Daiber, Donald Degraen, André Zenner, Tanja Döring, Frank Steinicke, Oscar Ariza, and Adalberto L. Simeone. "Everyday Proxy Objects for Virtual Reality." In: *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI EA'21. ACM, 2021. DOI: [10.1145/3411763.3441343](https://doi.org/10.1145/3411763.3441343).
- [64] Henrique Galvan Debarba, Ronan Boulic, Roy Salomon, Olaf Blanke, and Bruno Herbelin. "Self-Attribution of Distorted Reaching Movements in Immersive Virtual Reality." In: *Computers & Graphics* 76 (2018), pp. 142–152. DOI: <https://doi.org/10.1016/j.cag.2018.09.001>.
- [65] Donald Degraen, Martin Feick, Serdar Durdyev, and Antonio Krüger. "Prototyping Surface Slipperiness using Sole-Attached Textures during Haptic Walking in Virtual Reality." In: *Proceedings of the 23rd International Conference on Mobile and Ubiquitous Multimedia*. MUM'24. ACM, 2024. DOI: [10.1145/3701571.3701589](https://doi.org/10.1145/3701571.3701589).
- [66] Donald Degraen, Bruno Fruchard, Frederik Smolders, Emmanouil Potetsianakis, Seref Güngör, Antonio Krüger, and Jürgen Steimle. "Weirding Haptics: In-Situ Prototyping of Vibrotactile Feedback in Virtual Reality through Vocalization." In: *Proceedings of the 34th Annual ACM Symposium on User*

- Interface Software and Technology*. UIST'21. ACM, 2021. DOI: [10.1145/3472749.3474797](https://doi.org/10.1145/3472749.3474797).
- [67] Donald Degraen, Michal Piovarči, Bernd Bickel, and Antonio Krüger. "Capturing Tactile Properties of Real Surfaces for Haptic Reproduction." In: *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology*. UIST'21. ACM, 2021. DOI: [10.1145/3472749.3474798](https://doi.org/10.1145/3472749.3474798).
- [68] Donald Degraen, Anna Reindl, Akhmajon Makhsadov, André Zenner, and Antonio Krüger. "Envisioning Haptic Design for Immersive Virtual Environments." In: *Companion Publication of the 2020 ACM Designing Interactive Systems Conference*. DIS'20 Companion. ACM, 2020. DOI: [10.1145/3393914.3395870](https://doi.org/10.1145/3393914.3395870).
- [69] Donald Degraen, André Zenner, and Antonio Krüger. "Enhancing Texture Perception in Virtual Reality Using 3D-Printed Hair Structures." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300479](https://doi.org/10.1145/3290605.3300479).
- [70] A. Lee Dellon. "The Moving Two-Point Discrimination Test: Clinical Evaluation of the Quickly Adapting Fiber/Receptor System." In: *The Journal of Hand Surgery* 3.5 (1978), pp. 474–481. DOI: [https://doi.org/10.1016/S0363-5023\(78\)80143-9](https://doi.org/10.1016/S0363-5023(78)80143-9).
- [71] Diane Dewez, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet. "Do You Need Another Hand? Investigating Dual Body Representations During Anisomorphic 3D Manipulation." In: *IEEE Transactions on Visualization and Computer Graphics* 28.5 (2022), pp. 2047–2057. DOI: [10.1109/TVCG.2022.3150501](https://doi.org/10.1109/TVCG.2022.3150501).
- [72] Gianluca Di Flumeri et al. "EEG-Based Mental Workload Neuro-metric to Evaluate the Impact of Different Traffic and Road Conditions in Real Driving Settings." In: *Frontiers in Human Neuroscience* 12 (2018), p. 509. DOI: [10.3389/fnhum.2018.00509](https://doi.org/10.3389/fnhum.2018.00509).
- [73] Zoltan Dienes. "Using Bayes to Get the Most out of Non-significant Results." In: *Frontiers in Psychology* 5 (2014). 781. DOI: [10.3389/fpsyg.2014.00781](https://doi.org/10.3389/fpsyg.2014.00781).
- [74] Alan Dix, Janet E. Finlay, Gregory D. Abowd, and Russell Beale. *Human-Computer Interaction (3rd Edition)*. Prentice-Hall, Inc., 2003.
- [75] Lionel Dominjon, Anatole Lécuyer, J-M Burkhardt, Paul Richard, and Simon Richir. "Influence of Control/Display Ratio on the Perception of Mass of Manipulated Objects in Virtual Environments." In: *IEEE Proceedings Virtual Reality 2005*. IEEE, 2005. DOI: [10.1109/VR.2005.1492749](https://doi.org/10.1109/VR.2005.1492749).

- [76] Tim Düwel, Martin Feick, and Antonio Krüger. "Considering Interaction Types for Geometric Primitive Matching." In: *Everyday Proxy Objects for Virtual Reality Workshop of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI'21 Workshops. ACM, 2021.
- [77] Elham Ebrahimi, Andrew Robb, Leah S. Hartman, Christopher C. Pagano, and Sabarish V. Babu. "Effects of Anthropomorphic Fidelity of Self-avatars on Reach Boundary Estimation in Immersive Virtual Environments." In: *Proceedings of the 15th ACM Symposium on Applied Perception*. SAP'18. ACM, 2018. DOI: [10.1145/3225153.3225170](https://doi.org/10.1145/3225153.3225170).
- [78] Darragh Egan, Sean Brennan, John Barrett, Yuansong Qiao, Christian Timmerer, and Niall Murray. "An Evaluation of Heart Rate and ElectroDermal Activity as an Objective QoE Evaluation Method for Immersive Virtual Reality Environments." In: *2016 Eighth International Conference on Quality of Multimedia Experience*. QoMEX'16. IEEE, 2016. DOI: [10.1109/QoMEX.2016.7498964](https://doi.org/10.1109/QoMEX.2016.7498964).
- [79] Mohamed El Beheiry, Sébastien Doutreligne, Clément Caporal, Cécilia Ostertag, Maxime Dahan, and Jean-Baptiste Masson. "Virtual Reality: Beyond Visualization." In: *Journal of Molecular Biology* 431.7 (2019), pp. 1315–1321. DOI: [10.1016/j.jmb.2019.01.033](https://doi.org/10.1016/j.jmb.2019.01.033).
- [80] Marc O. Ernst and Martin S. Banks. "Humans Integrate Visual and Haptic Information in a Statistically Optimal Fashion." In: *Nature* 415.6870 (2002), pp. 429–433. DOI: [10.1038/415429a](https://doi.org/10.1038/415429a).
- [81] Marc O. Ernst and Heinrich H. Bühlhoff. "Merging the Senses into a Robust Percept." In: *Trends in Cognitive Sciences* 8.4 (2004), pp. 162–169. DOI: [10.1016/j.tics.2004.02.002](https://doi.org/10.1016/j.tics.2004.02.002).
- [82] David Escobar-Castillejos, Julieta Noguez, Luis Neri, Alejandra Magana, and Bedrich Benes. "A Review of Simulators with Haptic Devices for Medical Training." In: *Journal of Medical Systems* 40.4 (2016). 104. DOI: [10.1007/s10916-016-0459-8](https://doi.org/10.1007/s10916-016-0459-8).
- [83] Shaghayegh Esmaeili, Brett Benda, and Eric D. Ragan. "Detection of Scaled Hand Interactions in Virtual Reality: The Effects of Motion Direction and Task Complexity." In: *Proceedings of the 2020 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'20. IEEE, 2020. DOI: [10.1109/VR46266.2020.00066](https://doi.org/10.1109/VR46266.2020.00066).
- [84] James Coleman Eubanks, Alec G. Moore, Paul A. Fishwick, and Ryan P. McMahan. "The Effects of Body Tracking Fidelity on Embodiment of an Inverse-Kinematic Avatar for Male Participants." In: *Proceedings of the 2020 IEEE International Symposium on Mixed and Augmented Reality*. ISMAR'20. IEEE, 2020. DOI: [10.1109/ISMAR50242.2020.00025](https://doi.org/10.1109/ISMAR50242.2020.00025).

- [85] Hyun Taek Kim Eunhee Chang and Byounghyun Yoo. "Virtual Reality Sickness: A Review of Causes and Measurements." In: *International Journal of Human-Computer Interaction* 36.17 (2020), pp. 1658–1682. DOI: [10.1080/10447318.2020.1778351](https://doi.org/10.1080/10447318.2020.1778351).
- [86] Kirill A. Fadeev, Alexey S. Smirnov, Olga P. Zhigalova, Polina S. Bazhina, Alexey V. Tumialis, and Kirill S. Golokhvast. "Too Real to Be Virtual: Autonomic and EEG Responses to Extreme Stress Scenarios in Virtual Reality." In: *Behavioural Neurology* 2020 (2020). 5758038. DOI: [10.1155/2020/5758038](https://doi.org/10.1155/2020/5758038).
- [87] Cathy Mengying Fang and Chris Harrison. "Retargeted Self-Haptics for Increased Immersion in VR without Instrumentation." In: *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology*. UIST'21. ACM, 2021. DOI: [10.1145/3472749.3474810](https://doi.org/10.1145/3472749.3474810).
- [88] Cathy Fang, Yang Zhang, Matthew Dworman, and Chris Harrison. "Wireality: Enabling Complex Tangible Geometries in Virtual Reality with Worn Multi-String Haptics." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376470](https://doi.org/10.1145/3313831.3376470).
- [89] Alessandro Febretti, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, JD Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, Dana Plepys, et al. "CAVE2: A Hybrid Reality Environment for Immersive Simulation and Information Analysis." In: *Proceedings Volume 8649, The Engineering Reality of Virtual Reality*. SPIE, 2013. DOI: [10.1117/12.2005484](https://doi.org/10.1117/12.2005484).
- [90] Gustav Theodor Fechner. *Elemente der Psychophysik*. Vol. 2. Breitkopf and Härtel, 1860.
- [91] Martin Feick, Scott Bateman, Anthony Tang, André Miede, and Nicolai Marquardt. "TanGi: Tangible Proxies for Embodied Object Exploration and Manipulation in Virtual Reality." In: *Proceedings of the 2020 IEEE International Symposium on Mixed and Augmented Reality*. ISMAR'20. IEEE, 2020. DOI: [10.1109/ISMAR50242.2020.00042](https://doi.org/10.1109/ISMAR50242.2020.00042).
- [92] Martin Feick, Cihan Biyikli, Kiran Gani, Anton Wittig, Anthony Tang, and Antonio Krüger. "VoxelHap: A Toolkit for Constructing Proxies Providing Tactile and Kinesthetic Haptic Feedback in Virtual Reality." In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. UIST'23. ACM, 2023. DOI: [10.1145/3586183.3606722](https://doi.org/10.1145/3586183.3606722).
- [93] Martin Feick, Donald Degraen, Fabian Hupperich, and Antonio Krüger. "MetaReality: Enhancing Tactile Experiences using Actuated 3D-printed Metamaterials in Virtual Reality." In: *Frontiers in Virtual Reality* 4 (2023). 1172381. DOI: [10.3389/frvir.2023.1172381](https://doi.org/10.3389/frvir.2023.1172381).

- [94] Martin Feick, Niko Kleer, Anthony Tang, and Antonio Krüger. "The Virtual Reality Questionnaire Toolkit." In: *Adjunct Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST'20 Adjunct. ACM, 2020. DOI: [10.1145/3379350.3416188](https://doi.org/10.1145/3379350.3416188).
- [95] Martin Feick, Niko Kleer, André Zenner, Anthony Tang, and Antonio Krüger. "Visuo-Haptic Illusions for Linear Translation and Stretching Using Physical Proxies in Virtual Reality." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI'21. ACM, 2021. DOI: [10.1145/3411764.3445456](https://doi.org/10.1145/3411764.3445456).
- [96] Martin Feick, Terrance Mok, Anthony Tang, Lora Oehlberg, and Ehud Sharlin. "Perspective on and Re-Orientation of Physical Proxies in Object-Focused Remote Collaboration." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173855](https://doi.org/10.1145/3173574.3173855).
- [97] Martin Feick, Kora P. Regitz, Anthony Tang, Tobias Jungbluth, Maurice Rekrut, and Antonio Krüger. "Investigating Noticeable Hand Redirection in Virtual Reality using Physiological and Interaction Data." In: *Proceedings of the 2023 IEEE Conference Virtual Reality and 3D User Interfaces*. VR'23. IEEE, 2023. DOI: [10.1109/VR55154.2023.00035](https://doi.org/10.1109/VR55154.2023.00035).
- [98] Martin Feick, Kora Persephone Regitz, Lukas Gehrke, André Zenner, Anthony Tang, Tobias Patrick Jungbluth, Maurice Rekrut, and Antonio Krüger. "Predicting the Limits: Tailoring Unnoticeable Hand Redirection Offsets in Virtual Reality to Individuals' Perceptual Boundaries." In: *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. UIST'24. ACM, 2024. DOI: [10.1145/3654777.3676425](https://doi.org/10.1145/3654777.3676425).
- [99] Martin Feick, Kora Persephone Regitz, Anthony Tang, and Antonio Krüger. "Designing Visuo-Haptic Illusions with Proxies in Virtual Reality: Exploration of Grasp, Movement Trajectory and Object Mass." In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI'22. ACM, 2022. DOI: [10.1145/3491102.3517671](https://doi.org/10.1145/3491102.3517671).
- [100] Martin Feick, Anthony Tang, and Antonio Krüger. "HapticPuppet: A Kinesthetic Mid-air Multidirectional Force-Feedback Drone-based Interface." In: *Adjunct Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. UIST'22 Adjunct. ACM, 2022. DOI: [10.1145/3526114.3558694](https://doi.org/10.1145/3526114.3558694).

- [101] Martin Feick, Xuxin Tang, Raul Garcia-Marti, Alexandru Luchianov, Roderick Huang, Chang Xiao, Alexa Siu, and Mustafa Doga Dogan. "Imprinto: Enhancing Infrared Inkjet Watermarking for Human and Machine Perception." In: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. CHI'25. ACM, 2025. DOI: [10.1145/3706598.3713286](https://doi.org/10.1145/3706598.3713286).
- [102] Martin Feick, André Zenner, Simon Seibert, Oscar Javier Ariza Nuñez, David Wagmann, Juliana Helena Keller, Anton Wittig, and Antonio Krüger. "MoVRI: The Museum of Virtual Reality Illusions." In: *Demo at Saarland Informatic Campus* (2023).
- [103] Martin Feick, André Zenner, Simon Seibert, Anthony Tang, and Antonio Krüger. "The Impact of Avatar Completeness on Embodiment and the Detectability of Hand Redirection in Virtual Reality." In: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. CHI'24. ACM, 2024. DOI: [10.1145/3613904.3641933](https://doi.org/10.1145/3613904.3641933).
- [104] Martin Feick, André Zenner, Oscar Javier Ariza Nunez, Anthony Tang, Biyikli Cihan, and Antonio Krüger. "Turn-It-Up: Rendering Resistance for Knobs in Virtual Reality through Undetectable Pseudo-Haptics." In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. UIST'23. ACM, 2023. DOI: [10.1145/3586183.3606722](https://doi.org/10.1145/3586183.3606722).
- [105] Thomas Feix, Ian M. Bullock, and Aaron M. Dollar. "Analysis of Human Grasping Behavior: Correlating Tasks, Objects and Grasps." In: *IEEE Transactions on Haptics* 7.4 (2014), pp. 430–441. DOI: [10.1109/TOH.2014.2326867](https://doi.org/10.1109/TOH.2014.2326867).
- [106] Thomas Feix, Ian M. Bullock, and Aaron M. Dollar. "Analysis of Human Grasping Behavior: Object Characteristics and Grasp Type." In: *IEEE Transactions on Haptics* 7.3 (2014), pp. 311–323. DOI: [10.1109/TOH.2014.2326871](https://doi.org/10.1109/TOH.2014.2326871).
- [107] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M. Dollar, and Danica Kragic. "The GRASP Taxonomy of Human Grasp Types." In: *IEEE Transactions on Human-Machine Systems* 46.1 (2016), pp. 66–77. DOI: [10.1109/THMS.2015.2470657](https://doi.org/10.1109/THMS.2015.2470657).
- [108] Ana Ferlinc, Ester Fabiani, Tomaz Velnar, and Lidija Gradisnik. "The Importance and Role of Proprioception in the Elderly: A Short Review." In: *Materia Socio-Medica* 31.3 (2019), pp. 219–221. DOI: [10.5455/msm.2019.31.219-221](https://doi.org/10.5455/msm.2019.31.219-221).
- [109] Tiare Feuchtner and Jörg Müller. "Extending the Body for Interaction with Reality." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.3025689](https://doi.org/10.1145/3025453.3025689).

- [110] Andy Field. *Discovering Statistics using IBM SPSS Statistics (4th Edition)*. SAGE Publications, 2013.
- [111] Adrien Fonnet and Yannick Prié. “Survey of Immersive Analytics.” In: *IEEE Transactions on Visualization and Computer Graphics* 27.3 (2021), pp. 2101–2122. DOI: [10.1109/TVCG.2019.2929033](https://doi.org/10.1109/TVCG.2019.2929033).
- [112] Daniel Freeman, Polly Haselton, Jason Freeman, Bernhard Spanlang, Sameer Kishore, Emily Albery, Megan Denne, Poppy Brown, Mel Slater, and Alecia Nickless. “Automated Psychological Therapy using Immersive Virtual Reality for Treatment of Fear of Heights: A Single-blind, Parallel-group, Randomised Controlled Trial.” In: *The Lancet Psychiatry* 5.8 (2018), pp. 625–632. DOI: [10.1016/S2215-0366\(18\)30226-8](https://doi.org/10.1016/S2215-0366(18)30226-8).
- [113] Scott Frees, G. Drew Kessler, and Edwin Kay. “PRISM Interaction for Enhancing Control in Immersive Virtual Environments.” In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 14.1 (2007), pp. 2–es. DOI: [10.1145/1229855.1229857](https://doi.org/10.1145/1229855.1229857).
- [114] Kaushal Gangwar, RPS Gangwar, et al. “Metamaterials: Characteristics, Process and Applications.” In: *Advance in Electronic and Electric Engineering* 4.1 (2014), pp. 97–106.
- [115] BoYu Gao, Joonwoo Lee, Huawei Tu, Wonjun Seong, and HyungSeok Kim. “The Effects of Avatar Visibility on Behavioral Response with or without Mirror-Visual Feedback in Virtual Environments.” In: *Proceedings of the 2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops. VRW’20*. IEEE, 2020. DOI: [10.1109/VRW50115.2020.00241](https://doi.org/10.1109/VRW50115.2020.00241).
- [116] Marta I. Garrido, James M. Kilner, Klaas E. Stephan, and Karl J. Friston. “The Mismatch Negativity: A Review of Underlying Mechanisms.” In: *Clinical Neurophysiology* 120.3 (2009), pp. 453–463. DOI: [10.1016/j.clinph.2008.11.029](https://doi.org/10.1016/j.clinph.2008.11.029).
- [117] Angelika Gedsun, Riad Sahli, Xing Meng, René Hensel, and Roland Bennewitz. “Bending as Key Mechanism in the Tactile Perception of Fibrillar Surfaces.” In: *Advanced Materials Interfaces* 9.4 (2022), p. 2101380. DOI: <https://doi.org/10.1002/admi.202101380>.
- [118] Lukas Gehrke, Sezen Akman, Pedro Lopes, Albert Chen, Avinash Kumar Singh, Hsiang-Ting Chen, Chin-Teng Lin, and Klaus Gramann. “Detecting Visuo-Haptic Mismatches in Virtual Reality using the Prediction Error Negativity of Event-Related Brain Potentials.” In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. CHI’19*. ACM, 2019. DOI: [10.1145/3290605.3300657](https://doi.org/10.1145/3290605.3300657).

- [119] Lukas Gehrke, Pedro Lopes, Marius Klug, Sezen Akman, and Klaus Gramann. "Neural Sources of Prediction Errors Detect Unrealistic VR Interactions." In: *Journal of Neural Engineering* 19.3 (2022), p. 036002. DOI: [10.1088/1741-2552/ac69bc](https://doi.org/10.1088/1741-2552/ac69bc).
- [120] Tobias Geib and Martin Feick. "Everyday Objects for Volumetric 3D Sketching in Virtual Reality." In: *Everyday Proxy Objects for Virtual Reality Workshop of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI'21 Workshops. ACM, 2021.
- [121] Boynton Geoff. *Psychology 333 – Sensation and Perception*. 2008. URL: <https://courses.washington.edu/psy333/>.
- [122] Alan Gevins and Michael E. Smith. "Neurophysiological Measures of Cognitive Workload during Human-Computer Interaction." In: *Theoretical Issues in Ergonomics Science* 4.1-2 (2003), pp. 113–131. DOI: [10.1080/14639220210159717](https://doi.org/10.1080/14639220210159717).
- [123] James J. Gibson. "Adaptation, After-effect and Contrast in the Perception of Curved Lines." In: *Journal of Experimental Psychology* 16.1 (1933), pp. 1–31. DOI: [10.1037/h0074626](https://doi.org/10.1037/h0074626).
- [124] Guilherme Gonçalves, Miguel Melo, Luís Barbosa, José Vasconcelos-Raposo, and Maximino Bessa. "Evaluation of the Impact of Different Levels of Self-representation and Body Tracking on the Sense of Presence and Embodiment in Immersive VR." In: *Virtual Reality* 26.1 (2022), pp. 1–14. DOI: [10.1007/s10055-021-00530-5](https://doi.org/10.1007/s10055-021-00530-5).
- [125] Jun Gong, Olivia Seow, Cedric Honnet, Jack Forman, and Stefanie Mueller. "MetaSense: Integrating Sensing Capabilities into Mechanical Metamaterial." In: *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology*. UIST'21. ACM, 2021. DOI: [10.1145/3472749.3474806](https://doi.org/10.1145/3472749.3474806).
- [126] Mar Gonzalez-Franco, Parastoo Abtahi, and Anthony Steed. "Individual Differences in Embodied Distance Estimation in Virtual Reality." In: *Proceedings of the 2019 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'19. IEEE, 2019. DOI: [10.1109/VR.2019.8798348](https://doi.org/10.1109/VR.2019.8798348).
- [127] Mar Gonzalez-Franco, Brian Cohn, Eyal Ofek, Dalila Burin, and Antonella Maselli. "The Self-Avatar Follower Effect in Virtual Reality." In: *Proceedings of the 2020 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'20. IEEE, 2020. DOI: [doi: 10.1109/VR46266.2020.00019](https://doi.org/10.1109/VR46266.2020.00019).
- [128] Mar Gonzalez-Franco and Jaron Lanier. "Model of Illusions and Virtual Reality." In: *Frontiers in Psychology* 8 (2017). 1125. DOI: [10.3389/fpsyg.2017.01125](https://doi.org/10.3389/fpsyg.2017.01125).

- [129] Mar Gonzalez-Franco, Eyal Ofek, Ye Pan, Angus Antley, Anthony Steed, Bernhard Spanlang, Antonella Maselli, Domna Banakou, Nuria Pelechano, Sergio Orts-Escolano, et al. "The Rocketbox Library and the Utility of Freely Available Rigged Avatars." In: *Frontiers in Virtual Reality* 1 (2020). 561558. DOI: [10.3389/frvir.2020.561558](https://doi.org/10.3389/frvir.2020.561558).
- [130] Mar Gonzalez-Franco and Tabitha C. Peck. "Avatar Embodiment. Towards a Standardized Questionnaire." In: *Frontiers in Robotics and AI* 5 (2018). 74. DOI: [10.3389/frobt.2018.00074](https://doi.org/10.3389/frobt.2018.00074).
- [131] Eric J. Gonzalez, Parastoo Abtahi, and Sean Follmer. "Evaluating the Minimum Jerk Motion Model for Redirected Reach in Virtual Reality." In: *Adjunct Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST'19 Adjunct. ACM, 2019. DOI: [10.1145/3332167.3357096](https://doi.org/10.1145/3332167.3357096).
- [132] Eric J. Gonzalez, Parastoo Abtahi, and Sean Follmer. "REACH+: Extending the Reachability of Encountered-Type Haptics Devices through Dynamic Redirection in VR." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST'20. ACM, 2020. DOI: [10.1145/3379337.3415870](https://doi.org/10.1145/3379337.3415870).
- [133] Eric J. Gonzalez and Sean Follmer. "Investigating the Detection of Bimanual Haptic Retargeting in Virtual Reality." In: *Proceedings of the 25th ACM Symposium on Virtual Reality Software and Technology*. VRST'19. ACM, 2019. DOI: [10.1145/3359996.3364248](https://doi.org/10.1145/3359996.3364248).
- [134] Eric J. Gonzalez and Sean Follmer. "Sensorimotor Simulation of Redirected Reaching using Stochastic Optimal Feedback Control." In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI'23. ACM, 2023. DOI: [10.1145/3544548.3580767](https://doi.org/10.1145/3544548.3580767).
- [135] Eric J. Gonzalez, Eyal Ofek, Mar Gonzalez-Franco, and Mike Sinclair. "X-Rings: A Hand-Mounted 360° Shape Display for Grasping in Virtual Reality." In: *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology*. UIST'21. ACM, 2021. DOI: [10.1145/3472749.3474782](https://doi.org/10.1145/3472749.3474782).
- [136] Geoffrey Gorisse, Olivier Christmann, Samory Houzangbe, and Simon Richir. "From Robot to Virtual Doppelganger: Impact of Visual Fidelity of Avatars Controlled in Third-person Perspective on Embodiment and Behavior in Immersive Virtual Environments." In: *Frontiers in Robotics and AI* 6 (2019). 8. DOI: [10.3389/frobt.2019.00008](https://doi.org/10.3389/frobt.2019.00008).

- [137] Klaus Gramann, Friederike U. Hohlefeld, Lukas Gehrke, and Marius Klug. "Human Cortical Dynamics during Full-body Heading Changes." In: *Scientific Reports* 11.1 (2021). 18186. DOI: [10.1038/s41598-021-97749-8](https://doi.org/10.1038/s41598-021-97749-8).
- [138] Alexandre Gramfort et al. "MEG and EEG Data Analysis with MNE-Python." In: *Frontiers in Neuroscience* 7 (2013). 267. DOI: [10.3389/fnins.2013.00267](https://doi.org/10.3389/fnins.2013.00267).
- [139] Daniel Groeger, Martin Feick, Anusha Withana, and Jürgen Steimle. "Tactlets: Adding Tactile Feedback to 3D Objects Using Custom Printed Controls." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST'19. ACM, 2019. DOI: [10.1145/3332165.3347937](https://doi.org/10.1145/3332165.3347937).
- [140] Jože Guna, Gregor Geršak, Iztok Humar, Maja Krebl, Marko Orel, Huimin Lu, and Matevž Pogačnik. "Virtual Reality Sickness and Challenges Behind Different Technology and Content Settings." In: *Mobile Networks and Applications* 25.4 (2020), pp. 1436–1445. DOI: [10.1007/s11036-019-01373-w](https://doi.org/10.1007/s11036-019-01373-w).
- [141] Lesley A. Hall and DI McCloskey. "Detections of Movements Imposed on Finger, Elbow and Shoulder Joints." In: *The Journal of Physiology* 335.1 (1983), pp. 519–533. DOI: [10.1113/jphysiol.1983.sp014548](https://doi.org/10.1113/jphysiol.1983.sp014548).
- [142] Francis Hamit. *Virtual Reality and the Exploration of Cyberspace*. Sams, 1994.
- [143] Surina Hariri, Nur Ain Mustafa, Kasun Karunanayaka, and Adrian David Cheok. "Electrical Stimulation of Olfactory Receptors for Digitizing Smell." In: *Proceedings of the 2016 Workshop on Multimodal Virtual and Augmented Reality*. MVAR'16. ACM, 2016. DOI: [10.1145/3001959.3001964](https://doi.org/10.1145/3001959.3001964).
- [144] Judith Hartfill, Jenny Gabel, Lucie Kruse, Susanne Schmidt, Kevin Riebandt, Simone Kühn, and Frank Steinicke. "Analysis of Detection Thresholds for Hand Redirection during Mid-Air Interactions in Virtual Reality." In: *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology*. VRST'21. ACM, 2021. DOI: [10.1145/3489849.3489866](https://doi.org/10.1145/3489849.3489866).
- [145] Takeru Hashimoto, Shigeo Yoshida, and Takuji Narumi. "MetamorphX: An Ungrounded 3-DoF Moment Display that Changes its Physical Properties through Rotational Impedance Control." In: *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. UIST'22. ACM, 2022. DOI: [10.1145/3526113.3545650](https://doi.org/10.1145/3526113.3545650).
- [146] Morton L. Heilig. *Sensorama Simulator*. U.S. Patent 3,050,870, 1962.

- [147] Seongkook Heo, Christina Chung, Geehyuk Lee, and Daniel Wigdor. "Thor's Hammer: An Ungrounded Force Feedback Device Utilizing Propeller-Induced Propulsive Force." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3174099](https://doi.org/10.1145/3173574.3174099).
- [148] Seongkook Heo, Jaeyeon Lee, and Daniel Wigdor. "PseudoBend: Producing Haptic Illusions of Stretching, Bending, and Twisting Using Grain Vibrations." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST'19. ACM, 2019. DOI: [10.1145/3332165.3347941](https://doi.org/10.1145/3332165.3347941).
- [149] Stephan Hertweck, Desiée Weber, Hisham Alwanni, Fabian Unruh, Martin Fischbach, Marc Erich Latoschik, and Tonio Ball. "Brain Activity in Virtual Reality: Assessing Signal Quality of High-Resolution EEG While Using Head-Mounted Displays." In: *Proceedings of the 2019 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'19. IEEE, 2019. DOI: [10.1109/VR.2019.8798369](https://doi.org/10.1109/VR.2019.8798369).
- [150] Anuruddha Hettiarachchi and Daniel Wigdor. "Annexing Reality: Enabling Opportunistic Use of Everyday Objects as Tangible Proxies in Augmented Reality." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI'16. ACM, 2016. DOI: [10.1145/2858036.2858134](https://doi.org/10.1145/2858036.2858134).
- [151] Ken Hinckley, Randy Pausch, John C. Goble, and Neal F. Kassell. "Passive Real-World Interface Props for Neurosurgical Visualization." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI'94. ACM, 1994. DOI: [10.1145/191666.191821](https://doi.org/10.1145/191666.191821).
- [152] Sara Ilstedt Hjelm. "Research+ Design: The Making of Brainball." In: *Interactions* 10.1 (2003), pp. 26–34. DOI: [10.1145/604575.604576](https://doi.org/10.1145/604575.604576).
- [153] Hunter G. Hoffman. "Physically Touching Virtual Objects using Tactile Augmentation Enhances the Realism of Virtual Environments." In: *Proceedings of the IEEE 1998 Virtual Reality Annual International Symposium*. IEEE, 1998. DOI: [10.1109/VRAIS.1998.658423](https://doi.org/10.1109/VRAIS.1998.658423).
- [154] Simon M. Hofmann, Felix Klotzsche, Alberto Mariola, Vadim Nikulin, Arno Villringer, and Michael Gaebler. "Decoding Subjective Emotional Arousal from EEG during an Immersive Virtual Reality Experience." In: *eLife* 10 (2021). e64812. DOI: [10.7554/eLife.64812](https://doi.org/10.7554/eLife.64812).

- [155] Hsin-Yu Huang, Chih-Wei Ning, Po-Yao Wang, Jen-Hao Cheng, and Lung-Pan Cheng. "Haptic-Go-Round: A Surrounding Platform for Encounter-Type Haptics in Virtual Reality Experiences." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376476](https://doi.org/10.1145/3313831.3376476).
- [156] Carmen Hull and Wesley Willett. "Building with Data: Architectural Models as Inspiration for Data Physicalization." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.3025850](https://doi.org/10.1145/3025453.3025850).
- [157] Akiko Iesaki, Akihiro Somada, Asako Kimura, Fumihisa Shibata, and Hideyuki Tamura. "Psychophysical Influence on Tactile Impression by Mixed-Reality Visual Stimulation." In: *Proceedings of the 2008 IEEE Virtual Reality Conference*. IEEE, 2008. DOI: [10.1109/VR.2008.4480793](https://doi.org/10.1109/VR.2008.4480793).
- [158] Brent Edward Insko. *Passive Haptics Significantly Enhances Virtual Environments (PhD Thesis)*. The University of North Carolina at Chapel Hill, 2001.
- [159] Alexandra Ion, Robert Kovacs, Oliver S. Schneider, Pedro Lopes, and Patrick Baudisch. "Metamaterial Textures." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173910](https://doi.org/10.1145/3173574.3173910).
- [160] Alexandra Ion, Ludwig Wall, Robert Kovacs, and Patrick Baudisch. "Digital Mechanical Metamaterials." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.3025624](https://doi.org/10.1145/3025453.3025624).
- [161] Kozo Ishizaki, Sridhar Komarneni, and Makoto Nanko. *Porous Materials: Process Technology and Applications*. Vol. 4. Springer Science & Business Media, 2013.
- [162] Harold Jeffreys. *Theory of Probability*. OUP Oxford, 1998.
- [163] Seokhee Jeon and Seungmoon Choi. "Haptic Augmented Reality: Taxonomy and an Example of Stiffness Modulation." In: *Presence: Teleoperators and Virtual Environments* 18.5 (2009), pp. 387–408. DOI: [10.1162/pres.18.5.387](https://doi.org/10.1162/pres.18.5.387).
- [164] Seunggon Jeon, Seungwon Paik, Ungyeon Yang, Patrick C. Shih, and Kyungsik Han. "The More, the Better? Improving VR Firefighting Training System with Realistic Firefighter Tools as Controllers." In: *Sensors* 21.21 (2021), p. 7193. DOI: [10.3390/s21217193](https://doi.org/10.3390/s21217193).
- [165] Jason Jerald. *The VR Book: Human-Centered Design for Virtual Reality*. ACM and Morgan & Claypool, 2015.

- [166] Li Jiang, Rohit Girotra, Mark R. Cutkosky, and Chris Ullrich. "Reducing Error Rates with Low-Cost Haptic Feedback in Virtual Reality-Based Training Applications." In: *Proceedings of the IEEE World Haptics Conference*. WHC'05. IEEE, 2005. DOI: [10.1109/WHC.2005.111](https://doi.org/10.1109/WHC.2005.111).
- [167] Corinne Jola, Angharad Davis, and Patrick Haggard. "Proprioceptive Integration and Body Representation: Insights into Dancers' Expertise." In: *Experimental Brain Research* 213 (2011), pp. 257–265. DOI: [10.1007/s00221-011-2743-7](https://doi.org/10.1007/s00221-011-2743-7).
- [168] Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. "The ReacTable: Exploring the Synergy between Live Music Performance and Tabletop Tangible Interfaces." In: *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*. TEI'07. ACM, 2007. DOI: [10.1145/1226969.1226998](https://doi.org/10.1145/1226969.1226998).
- [169] Joel Jordan and Mel Slater. "An Analysis of Eye Scanpath Entropy in a Progressively Forming Virtual Environment." In: *Presence: Teleoperators and Virtual Environments* 18.3 (2009), pp. 185–199. DOI: [10.1162/pres.18.3.185](https://doi.org/10.1162/pres.18.3.185).
- [170] Denise Kahl, Marc Ruble, and Antonio Krüger. "The Influence of Environmental Lighting on Size Variations in Optical See-through Tangible Augmented Reality." In: *Proceedings of the 2022 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'22. IEEE, 2022. DOI: [10.1109/VR51125.2022.00030](https://doi.org/10.1109/VR51125.2022.00030).
- [171] Apostolos Kalatzis, Laura Stanley, and Vishnunarayan Girishan Prabhu. "Affective State Classification in Virtual Reality Environments Using Electrocardiogram and Respiration Signals." In: *Proceedings of the 2021 IEEE International Conference on Artificial Intelligence and Virtual Reality*. AIVR'21. IEEE, 2021. DOI: [10.1109/AIVR52153.2021.00037](https://doi.org/10.1109/AIVR52153.2021.00037).
- [172] Andreas Kalckert and H.H. Ehrsson. "The Onset Time of the Ownership Sensation in the Moving Rubber Hand Illusion." In: *Frontiers in Psychology* 8 (2017), p. 344. DOI: [10.3389/fpsyg.2017.00344](https://doi.org/10.3389/fpsyg.2017.00344).
- [173] Alexander Kalus, Johannes Klein, Tien-Julian Ho, Lee-Ann Seegets, and Niels Henze. "MobileGravity: Mobile Simulation of a High Range of Weight in Virtual Reality." In: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. CHI'24. ACM, 2024. DOI: [10.1145/3613904.3642658](https://doi.org/10.1145/3613904.3642658).
- [174] Noriko Kamakura, Michiko Matsuo, Harumi Ishii, Fumiko Mitsuboshi, and Yoriko Miura. "Patterns of Static Prehension in Normal Hands." In: *American Journal of Occupational Therapy* 34.7 (1980), pp. 437–445. DOI: [10.5014/ajot.34.7.437](https://doi.org/10.5014/ajot.34.7.437).

- [175] Gaetano Kanizsa. "Margini Quasi-percettivi in Campi con Stimolazione Omogenea." In: *Rivista di Psicologia* 49.1 (1955), pp. 7–30.
- [176] Kasun Karunanayaka, Nurafiqah Johari, Surina Hariri, Hanis Camelia, Kevin Stanley Bielawski, and Adrian David Cheok. "New Thermal Taste Actuation Technology for Future Multi-sensory Virtual Reality and Internet." In: *IEEE Transactions on Visualization and Computer Graphics* 24.4 (2018), pp. 1496–1505. DOI: [10.1109/TVCG.2018.2794073](https://doi.org/10.1109/TVCG.2018.2794073).
- [177] Robert S. Kennedy, Norman E. Lane, Kevin S. Berbaum, and Michael G. Lilienthal. "Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness." In: *The International Journal of Aviation Psychology* 3.3 (1993), pp. 203–220. DOI: [10.1207/s15327108ijap0303_3](https://doi.org/10.1207/s15327108ijap0303_3).
- [178] Christian Keysers, Valeria Gazzola, and Eric-Jan Wagenmakers. "Using Bayes Factor Hypothesis Testing in Neuroscience to Establish Evidence of Absence." In: *Nature Neuroscience* 23.7 (2020), pp. 788–799. DOI: [10.1038/s41593-020-0660-4](https://doi.org/10.1038/s41593-020-0660-4).
- [179] Johan Kildal. "Kooboh: Variable Tangible Properties in a Handheld Haptic-Illusion Box." In: *Proceedings of Haptics: Perception, Devices, Mobility, and Communication*. EuroHaptics'12. Springer, 2012. DOI: [10.1007/978-3-642-31404-9_33](https://doi.org/10.1007/978-3-642-31404-9_33).
- [180] Konstantina Kilteni, Raphaella Groten, and Mel Slater. "The Sense of Embodiment in Virtual Reality." In: *Presence: Teleoperators and Virtual Environments* 21.4 (2012), pp. 373–387. DOI: [10.1162/PRES_a_00124](https://doi.org/10.1162/PRES_a_00124).
- [181] Frederick A.A. Kingdom and Nicolaas Prins. "Chapter 5 - Adaptive Methods." In: *Psychophysics: A Practical Introduction (2nd Edition)*. Academic Press, 2016, pp. 119–148. DOI: [10.1016/B978-0-12-407156-8.00005-0](https://doi.org/10.1016/B978-0-12-407156-8.00005-0).
- [182] Itaru Kitahara, Morio Nakahara, and Yuichi Ohta. "Sensory Properties in Fusion of Visual/Haptic Stimuli using Mixed Reality." In: *Advances in Haptics* (2010). DOI: [10.5772/8712](https://doi.org/10.5772/8712).
- [183] Yoshifumi Kitamura, Yuichi Itoh, and Fumio Kishino. "Real-Time 3D Interaction with ActiveCube." In: *Extended Abstracts of the 2001 CHI Conference on Human Factors in Computing Systems*. CHI EA'01. ACM, 2001. DOI: [10.1145/634067.634277](https://doi.org/10.1145/634067.634277).
- [184] Roberta L. Klatzky, Susan J. Lederman, and Catherine Reed. "There's More to Touch than Meets the Eye: The Salience of Object Attributes for Haptics with and Without Vision." In: *Journal of Experimental Psychology: General* 116.4 (1987), 356–369. DOI: [10.1037/0096-3445.116.4.356](https://doi.org/10.1037/0096-3445.116.4.356).

- [185] Roberta L. Klatzky, Dianne Pawluk, and Angelika Peer. "Haptic Perception of Material Properties and Implications for Applications." In: *Proceedings of the IEEE* 101.9 (2013), pp. 2081–2092. DOI: [10.1109/JPROC.2013.2248691](https://doi.org/10.1109/JPROC.2013.2248691).
- [186] Niko Kleer and Martin Feick. "A Study on the Influence of Task Dependent Anthropomorphic Grasping Poses for Everyday Objects." In: *Proceedings of the 2022 IEEE-RAS 21st International Conference on Humanoid Robots*. Humanoids'22. IEEE, 2022. DOI: [10.1109/Humanoids53995.2022.10000198](https://doi.org/10.1109/Humanoids53995.2022.10000198).
- [187] Niko Kleer, Martin Feick, Florian Daiber, and Michael Feld. "Towards Remote Expert Supported Autonomous Assistant Robots in Shopping Environments." In: *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. HRI '24. ACM, 2024. DOI: [10.1145/3610978.3640735](https://doi.org/10.1145/3610978.3640735).
- [188] Niko Kleer, Martin Feick, and Michael Feld. "Leveraging Publicly Available Textual Object Descriptions for Anthropomorphic Robotic Grasp Predictions." In: *Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IROS'22. IEEE, 2022. DOI: [10.1109/IR0S47612.2022.9981541](https://doi.org/10.1109/IR0S47612.2022.9981541).
- [189] Niko Kleer, Martin Feick, Amr Gomaa, Michael Feld, and Antonio Krüger. "Bridging the Gap to Natural Language-based Grasp Predictions through Semantic Information Extraction." In: *Proceedings of the 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IROS'24. IEEE, 2024.
- [190] Niko Kleer, Ole Keil, Martin Feick, Amr Gomaa, Tim Schwartz, and Michael Feld. "Incorporation of the Intended Task into a Vision-based Grasp Type Predictor for Multi-fingered Robotic Grasping." In: *Proceedings of the 2024 33rd IEEE International Conference on Robot and Human Interactive Communication*. ROMAN'24. IEEE, 2024. DOI: [10.1109/RO-MAN60168.2024.10731261](https://doi.org/10.1109/RO-MAN60168.2024.10731261).
- [191] Pascal Knierim, Thomas Kosch, Valentin Schwind, Markus Funk, Francisco Kiss, Stefan Schneegass, and Niels Henze. "Tactile Drones - Providing Immersive Tactile Feedback in Virtual Reality through Quadcopters." In: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. CHI EA '17. ACM, 2017. DOI: [10.1145/3027063.3050426](https://doi.org/10.1145/3027063.3050426).
- [192] Luv Kohli. "Exploiting Perceptual Illusions to Enhance Passive Haptics." In: *Proceedings of the Workshop on Perceptual Illusions in Virtual Environments at the IEEE Conference on Virtual Reality*. PIVE'09. IEEE, 2009.

- [193] Luv Kohli. "Redirected Touching: Warping Space to Remap Passive Haptics." In: *Proceedings of the 2010 IEEE Symposium on 3D User Interface*. 3DUI'10. 2010. DOI: [10.1109/3DUI.2010.5444703](https://doi.org/10.1109/3DUI.2010.5444703).
- [194] Luv Kohli, Mary C. Whitton, and Frederick P. Brooks. "Redirected Touching: The Effect of Warping Space on Task Performance." In: *Proceedings of the 2012 IEEE Symposium on 3D User Interfaces*. 3DUI'12. 2012. DOI: [10.1109/3DUI.2012.6184193](https://doi.org/10.1109/3DUI.2012.6184193).
- [195] Felix Kosmalla, André Zenner, Corinna Tasch, Florian Daiber, and Antonio Krüger. "The Importance of Virtual Hands and Feet for Virtual Reality Climbing." In: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI EA'20. ACM, 2020. DOI: [10.1145/3334480.3383067](https://doi.org/10.1145/3334480.3383067).
- [196] Robert Kovacs, Eyal Ofek, Mar Gonzalez Franco, Alexa Fay Siu, Sebastian Marwecki, Christian Holz, and Mike Sinclair. "Haptic PIVOT: On-Demand Handhelds in VR." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST'20. ACM, 2020. DOI: [10.1145/3379337.3415854](https://doi.org/10.1145/3379337.3415854).
- [197] Eric Krokos and Amitabh Varshney. "Quantifying VR Cyber-sickness using EEG." In: *Virtual Reality* 26.1 (2022), pp. 77–89. DOI: [10.1007/s10055-021-00517-2](https://doi.org/10.1007/s10055-021-00517-2).
- [198] Katherine J. Kuchenbecker, Joseph Romano, and William McMahan. "Haptography: Capturing and Recreating the Rich Feel of Real Surfaces." In: *Proceedings of Robotics Research: The 14th International Symposium ISRR*. Springer. 2011. DOI: [10.1007/978-3-642-19457-3_15](https://doi.org/10.1007/978-3-642-19457-3_15).
- [199] Eun Kwon, Gerard J. Kim, and Sangyoon Lee. "Effects of Sizes and Shapes of Props in Tangible Augmented Reality." In: *Proceedings of the 2009 8th IEEE International Symposium on Mixed and Augmented Reality*. ISMAR'09. IEEE, 2009. DOI: [10.1109/ISMAR.2009.5336463](https://doi.org/10.1109/ISMAR.2009.5336463).
- [200] LSL. *The Lab Streaming Layer (LSL)*. <https://github.com/sccn/labstreaminglayer>. Accessed: 2024-05-05. 2024.
- [201] Joseph J. LaViola Jr, Ernst Kruijff, Ryan P. McMahan, Doug Bowman, and Ivan P. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, 2017.
- [202] Bruno Laeng, Shoaib Nabil, and Akiyoshi Kitaoka. "The Eye Pupil Adjusts to Illusorily Expanding Holes." In: *Frontiers in Human Neuroscience* 16 (2022). 877249. DOI: [10.3389/fnhum.2022.877249](https://doi.org/10.3389/fnhum.2022.877249).

- [203] Ewen B. Lavoie, Aida M. Valevicius, Quinn A. Boser, Ognjen Kovic, Albert H. Vette, Patrick M. Pilarski, Jacqueline S. Hebert, and Craig S Chapman. "Using Synchronized Eye and Motion Tracking to Determine High-Precision Eye-Movement Patterns during Object-Interaction Tasks." In: *Journal of Vision* 18.6 (2018), p. 18. DOI: [10.1167/18.6.18](https://doi.org/10.1167/18.6.18).
- [204] Ewen Lavoie and Craig S. Chapman. "What's Limbs Got to Do With It? Real-world Movement Correlates with Feelings of Ownership over Virtual Arms during Object Interactions in Virtual Reality." In: *Neuroscience of Consciousness* 2020.1 (2021). niaa027. DOI: [10.1093/nc/niaa027](https://doi.org/10.1093/nc/niaa027).
- [205] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. *Research Methods in Human-Computer Interaction (2nd Edition)*. Morgan Kaufmann, 2017.
- [206] Flavien Lebrun, Sinan D Haliyo, Malika Auvray, David Gueorguiev, and Gilles Bailly. "Study of the Impact of Visual and Haptic Sensory Sensitivities on the Detection of Visuo-Haptic Illusions in Virtual Reality." In: *Proceedings of the 35th Conference on l'Interaction Humain-Machine. IHM '24*. ACM, 2024. DOI: [10.1145/3649792.3649803](https://doi.org/10.1145/3649792.3649803).
- [207] Flavien Lebrun, Sinan Haliyo, and Gilles Bailly. "A Trajectory Model for Desktop-Scale Hand Redirection in Virtual Reality." In: *Proceedings of the IFIP Conference on Human-Computer Interaction. INTERACT'21*. Springer, 2021. DOI: [10.1007/978-3-030-85607-6_8](https://doi.org/10.1007/978-3-030-85607-6_8).
- [208] Anatole Lécuyer, J-M Burkhardt, Sabine Coquillart, and Philippe Coiffet. "'Boundary of Illusion': An Experiment of Sensory Integration with a Pseudo-Haptic System." In: *Proceedings IEEE Virtual Reality 2001*. IEEE, 2001. DOI: [10.1109/VR.2001.913777](https://doi.org/10.1109/VR.2001.913777).
- [209] Anatole Lécuyer, Jean-Marie Burkhardt, and Laurent Etienne. "Feeling Bumps and Holes without a Haptic Interface: The Perception of Pseudo-Haptic Textures." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI'04*. ACM, 2004. DOI: [10.1145/985692.985723](https://doi.org/10.1145/985692.985723).
- [210] Anatole Lécuyer, Sabine Coquillart, Abderrahmane Kheddar, Paul Richard, and Philippe Coiffet. "Pseudo-Haptic Feedback: Can Isometric Input Devices Simulate Force Feedback?" In: *Proceedings IEEE Virtual Reality 2000 (Cat. No.00CB37048)*. IEEE, 2000. DOI: [10.1109/VR.2000.840369](https://doi.org/10.1109/VR.2000.840369).
- [211] Susan J. Lederman and Roberta L. Klatzky. "Haptic Perception: A Tutorial." In: *Attention, Perception, & Psychophysics* 71.7 (2009), pp. 1439–1459. DOI: [10.3758/APP.71.7.1439](https://doi.org/10.3758/APP.71.7.1439).

- [212] David Ledo, Steven Houben, Jo Vermeulen, Nicolai Marquardt, Lora Oehlberg, and Saul Greenberg. "Evaluation Strategies for HCI Toolkit Research." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173610](https://doi.org/10.1145/3173574.3173610).
- [213] Chi-Jung Lee, Hsin-Ruey Tsai, and Bing-Yu Chen. "Hair-Touch: Providing Stiffness, Roughness and Surface Height Differences Using Reconfigurable Brush Hairs on a VR Controller." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI'21. ACM, 2021. DOI: [10.1145/3411764.3445285](https://doi.org/10.1145/3411764.3445285).
- [214] Jaeyeon Lee, Mike Sinclair, Mar Gonzalez-Franco, Eyal Ofek, and Christian Holz. "TORC: A Virtual Reality Controller for In-Hand High-Dexterity Finger Interaction." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300301](https://doi.org/10.1145/3290605.3300301).
- [215] Marjorie R. Leek. "Adaptive Procedures in Psychophysical Research." In: *Perception & Psychophysics* 63.8 (2001), pp. 1279–1292. DOI: [10.3758/BF03194543](https://doi.org/10.3758/BF03194543).
- [216] Danny Leen, Raf Ramakers, and Kris Luyten. "StrutModeling: A Low-Fidelity Construction Kit to Iteratively Model, Test, and Adapt 3D Objects." In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST'17. ACM, 2017. DOI: [10.1145/3126594.3126643](https://doi.org/10.1145/3126594.3126643).
- [217] Markus Leyrer, Sally A. Linkenauger, Heinrich H. Bühlhoff, Uwe Kloos, and Betty Mohler. "The Influence of Eye Height and Avatars on Egocentric Distance Estimates in Immersive Virtual Environments." In: *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization*. APGV'11. ACM, 2011. DOI: [10.1145/2077451.2077464](https://doi.org/10.1145/2077451.2077464).
- [218] Nianlong Li, Han-Jong Kim, LuYao Shen, Feng Tian, Teng Han, Xing-Dong Yang, and Tek-Jin Nam. "HapLinkage: Prototyping Haptic Proxies for Virtual Hand Tools Using Linkage Mechanism." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST'20. ACM, 2020. DOI: [10.1145/3379337.3415812](https://doi.org/10.1145/3379337.3415812).
- [219] Jakub Limanowski. "Precision Control for a Flexible Body Representation." In: *Neuroscience & Biobehavioral Reviews* 134 (2022), p. 104401. DOI: [10.1016/j.neubiorev.2021.10.023](https://doi.org/10.1016/j.neubiorev.2021.10.023).
- [220] Robert W. Lindeman, John L. Sibert, and James K. Hahn. "Hand-Held Windows: Towards Effective 2D Interaction in Immersive Virtual Environments." In: *Proceedings IEEE Virtual Reality (Cat. No. 99CB36316)*. IEEE, 1999. DOI: [10.1109/VR.1999.756952](https://doi.org/10.1109/VR.1999.756952).

- [221] Lei Liu, Robert van Liere, Catharina Nieuwenhuizen, and Jean-Bernard Martens. "Comparing Aimed Movements in the Real World and in Virtual Reality." In: *Proceedings of the 2009 IEEE Virtual Reality Conference*. IEEE, 2009. DOI: [10.1109/VR.2009.4811026](https://doi.org/10.1109/VR.2009.4811026).
- [222] Benjamin Lok, Samir Naik, Mary Whitton, and Frederick P. Brooks. "Effects of Handling Real Objects and Self-Avatar Fidelity on Cognitive Task Performance and Sense of Presence in Virtual Environments." In: *Presence: Teleoperators and Virtual Environments* 12.6 (2003), pp. 615–628. DOI: [10.1162/105474603322955914](https://doi.org/10.1162/105474603322955914).
- [223] Benjamin Long, Sue Ann Seah, Tom Carter, and Sriram Subramanian. "Rendering Volumetric Haptic Shapes in Mid-air using Ultrasound." In: *ACM Transactions on Graphics* 33.6 (2014). 181. DOI: [10.1145/2661229.2661257](https://doi.org/10.1145/2661229.2661257).
- [224] Pedro Lopes, Alexandra Ion, and Patrick Baudisch. "Impacto: Simulating Physical Impact by Combining Tactile Stimulation with Electrical Muscle Stimulation." In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*. UIST'15. ACM, 2015. DOI: [10.1145/2807442.2807443](https://doi.org/10.1145/2807442.2807443).
- [225] Pedro Lopes, Sijing You, Lung-Pan Cheng, Sebastian Marwecki, and Patrick Baudisch. "Providing Haptics to Walls & Heavy Objects in Virtual Reality by Means of Electrical Muscle Stimulation." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.3025600](https://doi.org/10.1145/3025453.3025600).
- [226] G Lu, GQ Lu, and ZM Xiao. "Mechanical Properties of Porous Materials." In: *Journal of Porous Materials* 6 (1999), pp. 359–368. DOI: [10.1023/A:1009669730778](https://doi.org/10.1023/A:1009669730778).
- [227] Anatole Lécuyer. "Simulating Haptic Feedback Using Vision: A Survey of Research and Applications of Pseudo-Haptic Feedback." In: *Presence: Teleoperators and Virtual Environments* 18.1 (2009), pp. 39–53. DOI: [10.1162/pres.18.1.39](https://doi.org/10.1162/pres.18.1.39).
- [228] I. Scott MacKenzie. "Fitts' Law as a Research and Design Tool in Human-Computer Interaction." In: *Human-Computer Interaction* 7.1 (1992), pp. 91–139. DOI: [10.1207/s15327051hci0701_3](https://doi.org/10.1207/s15327051hci0701_3).
- [229] Akhmajon Makhsadov, Donald Degraen, André Zenner, Felix Kosmalla, Kamila Mushkina, and Antonio Krüger. "VRySmart: A Framework for Embedding Smart Devices in Virtual Reality." In: *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI EA'22. ACM, 2022. DOI: [10.1145/3491101.3519717](https://doi.org/10.1145/3491101.3519717).

- [230] Dominique Makowski, Tam Pham, Zen J. Lau, Jan C. Brammer, François Lespinasse, Hung Pham, Christopher Schölzel, and S. H. Annabel Chen. "NeuroKit2: A Python Toolbox for Neurophysiological Signal Processing." In: *Behavior Research Methods* 53.4 (2021), pp. 1689–1696. DOI: [10.3758/s13428-020-01516-y](https://doi.org/10.3758/s13428-020-01516-y).
- [231] Marzena Malińska, Krystyna Zużewicz, Joanna Bugajska, and Andrzej Grabowski. "Heart Rate Variability (HRV) during Virtual Reality Immersion." In: *International Journal of Occupational Safety and Ergonomics* 21.1 (2015), pp. 47–54. DOI: [10.1080/10803548.2015.1017964](https://doi.org/10.1080/10803548.2015.1017964).
- [232] Elena Marchiori, Evangelos Niforatos, and Luca Preto. "Analysis of Users' Heart Rate Data and Self-reported Perceptions to Understand Effective Virtual Reality Characteristics." In: *Information Technology & Tourism* 18.1 (2018), pp. 133–155. DOI: [10.1007/s40558-018-0104-0](https://doi.org/10.1007/s40558-018-0104-0).
- [233] Keith A. Markus. "Principles and Practice of Structural Equation Modeling by Rex B. Kline." In: *Structural Equation Modeling: A Multidisciplinary Journal* 19.3 (2012), pp. 509–512. DOI: [10.1080/10705511.2012.687667](https://doi.org/10.1080/10705511.2012.687667).
- [234] Thomas H. Massie, J. Kenneth Salisbury, et al. "The Phantom Haptic Interface: A Device for Probing Virtual Objects." In: *Proceedings of the ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. ASME, 1994.
- [235] Keigo Matsumoto, Yuki Ban, Takuji Narumi, Yohei Yanase, Tomohiro Tanikawa, and Michitaka Hirose. "Unlimited Corridor: Redirected Walking Techniques Using Visuo Haptic Interaction." In: *Proceedings of the ACM SIGGRAPH 2016 Emerging Technologies*. SIGGRAPH'16. ACM, 2016. DOI: [10.1145/2929464.2929482](https://doi.org/10.1145/2929464.2929482).
- [236] Brandon J. Matthews, Bruce H. Thomas, G. Stewart Von Itzstein, and Ross T. Smith. "Shape Aware Haptic Retargeting for Accurate Hand Interactions." In: *Proceedings of the 2022 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'22. IEEE, 2022. DOI: [10.1109/VR51125.2022.00083](https://doi.org/10.1109/VR51125.2022.00083).
- [237] Brandon J. Matthews, Bruce H. Thomas, G. Stewart Von Itzstein, and Ross T. Smith. "Towards Applied Remapped Physical-Virtual Interfaces: Synchronization Methods for Resolving Control State Conflicts." In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI'23. ACM, 2023. DOI: [10.1145/3544548.3580723](https://doi.org/10.1145/3544548.3580723).
- [238] Brandon J. Matthews, Bruce H. Thomas, Stewart Von Itzstein, and Ross T. Smith. "Remapped Physical-Virtual Interfaces with Bimanual Haptic Retargeting." In: *Proceedings of the 2019*

- IEEE Conference on Virtual Reality and 3D User Interfaces. VR'19. IEEE, 2019. DOI: [10.1109/VR.2019.8797974](https://doi.org/10.1109/VR.2019.8797974).*
- [239] John C. McClelland, Robert J. Teather, and Audrey Girouard. "Haptobend: Shape-Changing Passive Haptic Feedback in Virtual Reality." In: *Proceedings of the 5th Symposium on Spatial User Interaction. SUI'17. ACM, 2017. DOI: [10.1145/3131277.3132179](https://doi.org/10.1145/3131277.3132179).*
 - [240] William A. McNeely. "Robotic Graphics: A New Approach to Force Feedback for Virtual Reality." In: *Proceedings of IEEE Virtual Reality Annual International Symposium. IEEE, 1993. DOI: [10.1109/VRAIS.1993.380761](https://doi.org/10.1109/VRAIS.1993.380761).*
 - [241] Michael Meehan, Brent Insko, Mary Whitton, and Frederick P. Brooks. "Physiological Measures of Presence in Stressful Virtual Environments." In: *ACM Transactions on Graphics* 21.3 (2002), pp. 645–652. DOI: [10.1145/566654.566630](https://doi.org/10.1145/566654.566630).
 - [242] Paul Milgram and Fumio Kishino. "A Taxonomy of Mixed Reality Visual Displays." In: *IEICE Transactions on Information and Systems* E77-D.12 (1994), pp. 1321–1329.
 - [243] Kouta Minamizawa, Souichiro Fukamachi, Hiroyuki Kajimoto, Naoki Kawakami, and Susumu Tachi. "Gravity Grabber: Wearable Haptic Display to Present Virtual Mass Sensation." In: *Proceedings of the ACM SIGGRAPH 2007 Emerging Technologies. SIGGRAPH'07. ACM, 2007. DOI: [10.1145/1278280.1278289](https://doi.org/10.1145/1278280.1278289).*
 - [244] Reiji Miura, Shunichi Kasahara, Michiteru Kitazaki, Adrien Verhulst, Masahiko Inami, and Maki Sugimoto. "Multi-Soma: Distributed Embodiment with Synchronized Behavior and Perception." In: *Proceedings of the Augmented Humans International Conference 2021. AHs'21. ACM, 2021. DOI: [10.1145/3458709.3458878](https://doi.org/10.1145/3458709.3458878).*
 - [245] Roberto A. Montano Murillo, Sriram Subramanian, and Diego Martinez Plasencia. "Erg-O: Ergonomic Optimization of Immersive Virtual Environments." In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology. UIST'17. ACM, 2017. DOI: [10.1145/3126594.3126605](https://doi.org/10.1145/3126594.3126605).*
 - [246] Kosuke Morimoto, Kenta Hashiura, and Keita Watanabe. "Effect of Virtual Hand's Fingertip Deformation on the Stiffness Perceived Using Pseudo-Haptics." In: *Proceedings of the 29th ACM Symposium on Virtual Reality Software and Technology. VRST '23. ACM, 2023. DOI: [10.1145/3611659.3615689](https://doi.org/10.1145/3611659.3615689).*
 - [247] Stefanie Mueller, Tobias Mohr, Kerstin Guenther, Johannes Frohnhofen, and Patrick Baudisch. "FaBrickation: Fast 3D Printing of Functional Objects by Integrating Construction Kit Building Blocks." In: *Proceedings of the SIGCHI Conference on*

- Human Factors in Computing Systems*. CHI'14. ACM, 2014. DOI: [10.1145/2556288.2557005](https://doi.org/10.1145/2556288.2557005).
- [248] Thomas Muender, Anke V. Reinschluessel, Sean Drewes, Dirk Wenig, Tanja Döring, and Rainer Malaka. "Does It Feel Real? Using Tangibles with Different Fidelities to Build and Explore Scenes in Virtual Reality." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300903](https://doi.org/10.1145/3290605.3300903).
- [249] Ernst Mutschler, Hans-Georg Schaible, and Peter Vaupel. *Anatomie, Physiologie, Pathophysiologie des Menschen*. Vol. 6. Wissenschaftliche Verlagsgesellschaft mbH Stuttgart, 2007.
- [250] NO Myers. "Characterization of Surface Roughness." In: *Wear* 5:3 (1962), pp. 182–189. DOI: [https://doi.org/10.1016/0043-1648\(62\)90002-9](https://doi.org/10.1016/0043-1648(62)90002-9).
- [251] FC Müller-Lyer. "Optische Urteilstäuschungen." In: *Archiv für Anatomie und Physiologie, Physiologische Abteilung* 2 (1889), pp. 263–270.
- [252] Robin M. Neville, Fabrizio Scarpa, and Alberto Pirrera. "Shape Morphing Kirigami Mechanical Metamaterials." In: *Scientific Reports* 6:1 (2016), p. 31067. DOI: [10.1038/srep31067](https://doi.org/10.1038/srep31067).
- [253] Willy Nguyen, Klaus Gramann, and Lukas Gehrke. "Modeling the Intent to Interact With VR Using Physiological Features." In: *IEEE Transactions on Visualization and Computer Graphics* 30:8 (2024), pp. 5893–5900. DOI: [10.1109/TVCG.2023.3308787](https://doi.org/10.1109/TVCG.2023.3308787).
- [254] Karin Nieuwenhuizen, Lei Liu, Robert van Liere, and Jean-Bernard Martens. "Insights from Dividing 3D Goal-Directed Movements into Meaningful Phases." In: *IEEE Computer Graphics and Applications* 29:6 (2009), pp. 44–53. DOI: [10.1109/MCG.2009.121](https://doi.org/10.1109/MCG.2009.121).
- [255] Niels Christian Nilsson, Rolf Nordahl, and Stefania Serafin. "Immersion Revisited: A Review of Existing Definitions of Immersion and their Relation to Different Theories of Presence." In: *Human Technology* 12:2 (2016), pp. 108–134. DOI: [10.17011/HT/URN.201611174652](https://doi.org/10.17011/HT/URN.201611174652).
- [256] Niels Christian Nilsson, André Zenner, and Adalberto L. Simeone. "Propping Up Virtual Reality With Haptic Proxies." In: *IEEE Computer Graphics and Applications* 41:5 (2021), pp. 104–112. DOI: [10.1109/MCG.2021.3097671](https://doi.org/10.1109/MCG.2021.3097671).
- [257] Tommy Nilsson, Flavie Rometsch, Leonie Becker, Florian Dufresne, Paul Demedeiros, Enrico Guerra, Andrea Emanuele Maria Casini, Anna Vock, Florian Gaeremynck, and Aidan Cowley. "Using Virtual Reality to Shape Humanity's Return to the Moon: Key Takeaways from a Design Study." In: *Proceed-*

- ings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI'23. ACM, 2023. DOI: [10.1145/3544548.3580718](https://doi.org/10.1145/3544548.3580718).
- [258] Martin Nisser, Leon Cheng, Yashaswini Makaram, Ryo Suzuki, and Stefanie Mueller. "ElectroVoxel: Electromagnetically Actuated Pivoting for Scalable Modular Self-Reconfigurable Robots." In: *Proceedings of the 2022 IEEE International Conference on Robotics and Automation*. ICRA'22. ACM, 2022. DOI: [10.1109/ICRA46639.2022.9811746](https://doi.org/10.1109/ICRA46639.2022.9811746).
- [259] Kenneth H. Norwich. "On Theory of Weber Fractions." In: *Perception & Psychophysics* 42.3 (1987), pp. 286–298. DOI: [10.3758/BF03203081](https://doi.org/10.3758/BF03203081).
- [260] Torben Noto, Guangyu Zhou, Stephan Schuele, Jessica Templer, and Christina Zelano. "Automated Analysis of Breathing Waveforms using BreathMetrics: A Respiratory Signal Processing Toolbox." In: *Chemical Senses* 43.8 (2018), pp. 583–597. DOI: [10.1093/chemse/bjy045](https://doi.org/10.1093/chemse/bjy045).
- [261] Maki Ogawa, Keigo Matsumoto, Kazuma Aoyama, and Takuji Narumi. "Expansion of Detection Thresholds for Hand Redirection using Noisy Tendon Electrical Stimulation." In: *Proceedings of the 2023 IEEE International Symposium on Mixed and Augmented Reality*. ISMAR'23. IEEE, 2023. DOI: [10.1109/ISMAR59233.2023.00119](https://doi.org/10.1109/ISMAR59233.2023.00119).
- [262] Nami Ogawa, Takuji Narumi, and Michitaka Hirose. "Effect of Avatar Appearance on Detection Thresholds for Remapped Hand Movements." In: *IEEE Transactions on Visualization and Computer Graphics* 27.7 (2021), pp. 3182–3197. DOI: [10.1109/TVCG.2020.2964758](https://doi.org/10.1109/TVCG.2020.2964758).
- [263] Shogo Okamoto, Hikaru Nagano, and Yoji Yamada. "Psychophysical Dimensions of Tactile Perception of Textures." In: *IEEE Transactions on Haptics* 6.1 (2013), pp. 81–93. DOI: [10.1109/toh.2012.32](https://doi.org/10.1109/toh.2012.32).
- [264] Jason Orlosky, Yuta Itoh, Maud Ranchet, Kiyoshi Kiyokawa, John Morgan, and Hannes Devos. "Emulation of Physician Tasks in Eye-Track Virtual Reality for Remote Diagnosis of Neurodegenerative Disease." In: *IEEE Transactions on Visualization and Computer Graphics* 23.4 (2017), pp. 1302–1311. DOI: [10.1109/TVCG.2017.2657018](https://doi.org/10.1109/TVCG.2017.2657018).
- [265] Gonçalo Padrao, Mar Gonzalez-Franco, Maria V. Sanchez-Vives, Mel Slater, and Antoni Rodriguez-Fornells. "Violating Body Movement Semantics: Neural Signatures of Self-generated and External-generated Errors." In: *NeuroImage* 124.Pt A (2016), pp. 147–156. DOI: [10.1016/j.neuroimage.2015.08.022](https://doi.org/10.1016/j.neuroimage.2015.08.022).

- [266] Xueni Pan and Mel Slater. "A Preliminary Study of Shy Males Interacting with a Virtual Female." In: *Presence: The 10th Annual International Workshop on Presence*. ISPR, 2007.
- [267] Chaeyong Park, Jinhyuk Yoon, Seungjae Oh, and Seungmoon Choi. "Augmenting Physical Buttons with Vibrotactile Feedback for Programmable Feels." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST'20. ACM, 2020. DOI: [10.1145/3379337.3415837](https://doi.org/10.1145/3379337.3415837).
- [268] Courtnie J. Paschall, Jason S. Hauptman, Rajesh PN Rao, Jeffrey G. Ojemann, and Jeffrey Herron. "Designing Touch: Intracortical Neurohaptic Feedback in Virtual Reality." In: *Brain-Computer Interface Research: A State-of-the-Art Summary 11*. Springer, 2024, pp. 93–107. DOI: [10.1007/978-3-031-49457-4_10](https://doi.org/10.1007/978-3-031-49457-4_10).
- [269] Tam Pham, Zen Juen Lau, S. H. Annabel Chen, and Dominique Makowski. "Heart Rate Variability in Psychology: A Review of HRV Indices and an Analysis Tutorial." In: *Sensors* 21.12 (2021). 3998. DOI: [10.3390/s21123998](https://doi.org/10.3390/s21123998).
- [270] Lane Phillips, Brian Ries, Michael Kaeding, and Victoria Interrante. "Avatar Self-embodiment Enhances Distance Perception Accuracy in Non-photorealistic Immersive Virtual Environments." In: *Proceedings of the 2010 IEEE Virtual Reality Conference*. VR'10. IEEE, 2010. DOI: [10.1109/VR.2010.5444802](https://doi.org/10.1109/VR.2010.5444802).
- [271] Michal Piovarči, David I. W. Levin, Jason Rebello, Desai Chen, Roman Ďurikovič, Hanspeter Pfister, Wojciech Matusik, and Piotr Didyk. "An Interaction-Aware, Perceptual Model for Non-Linear Elastic Objects." In: *ACM Transactions on Graphics* 35.4 (2016). 55. DOI: [10.1145/2897824.2925885](https://doi.org/10.1145/2897824.2925885).
- [272] Ivan Poupyrev, Mark Billingham, Suzanne Weghorst, and Tadao Ichikawa. "The Go-Go Interaction Technique: Non-Linear Mapping for Direct Manipulation in VR." In: *Proceedings of the 9th annual ACM Symposium on User Interface Software and Technology*. UIST'96. ACM, 1996. DOI: [10.1145/237091.237102](https://doi.org/10.1145/237091.237102).
- [273] Kristopher J. Preacher, Derek D. Rucker, Robert C. MacCallum, and Alan W. Nicewander. "Use of the Extreme Groups Approach: A Critical Reexamination and New Recommendations." In: *Psychological Methods* 10.2 (2005), pp. 178–192. DOI: [10.1037/1082-989X.10.2.178](https://doi.org/10.1037/1082-989X.10.2.178).
- [274] Uwe Proske and Simon C. Gandevia. "The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force." In: *Physiological Reviews* 92.4 (2012), pp. 1651–1697. DOI: [10.1152/physrev.00048.2011](https://doi.org/10.1152/physrev.00048.2011).

- [275] Parinya Punpongsanon, Daisuke Iwai, and Kosuke Sato. "SoftAR: Visually Manipulating Haptic Softness Perception in Spatial Augmented Reality." In: *IEEE Transactions on Visualization and Computer Graphics* 21.11 (2015), pp. 1279–1288. DOI: [10.1109/TVCG.2015.2459792](https://doi.org/10.1109/TVCG.2015.2459792).
- [276] Susanne Putze, Dmitry Alexandrovsky, Felix Putze, Sebastian Höffner, Jan David Smeddinck, and Rainer Malaka. "Breaking The Experience: Effects of Questionnaires in VR User Studies." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376144](https://doi.org/10.1145/3313831.3376144).
- [277] Nimesha Ranasinghe, Ryohei Nakatsu, Hideaki Nii, and Pon-nampalam Gopalakrishnakone. "Tongue Mounted Interface for Digitally Actuating the Sense of Taste." In: *Proceedings of the 2012 16th International Symposium on Wearable Computers*. IEEE, 2012. DOI: [10.1109/ISWC.2012.16](https://doi.org/10.1109/ISWC.2012.16).
- [278] Nimesha Ranasinghe, Gajan Suthokumar, Kuan-Yi Lee, and Ellen Yi-Luen Do. "Digital Flavor: Towards Digitally Simulating Virtual Flavors." In: *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*. ICMI'15. ACM, 2015. DOI: [10.1145/2818346.2820761](https://doi.org/10.1145/2818346.2820761).
- [279] Sharif Razzaque. *Redirected Walking (PhD Thesis)*. The University of North Carolina at Chapel Hill, 2005.
- [280] Fernando Ribeiro and Jose Oliveira. "Factors Influencing Proprioception: What do They Reveal?" In: *Biomechanics in Applications*. IntechOpen, 2011. Chap. 14. DOI: [10.5772/20335](https://doi.org/10.5772/20335).
- [281] Brian Ries, Victoria Interrante, Michael Kaeding, and Lane Phillips. "Analyzing the Effect of a Virtual Avatar's Geometric and Motion Fidelity on Ego-centric Spatial Perception in Immersive Virtual Environments." In: *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology*. VRST'09. ACM, 2009. DOI: [10.1145/1643928.1643943](https://doi.org/10.1145/1643928.1643943).
- [282] Michael Rietzler, Florian Geiselhart, Jan Gugenheimer, and Enrico Rukzio. "Breaking the Tracking: Enabling Weight Perception Using Perceivable Tracking Offsets." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173702](https://doi.org/10.1145/3173574.3173702).
- [283] Gabriel Robles-De-La-Torre. "The Importance of the Sense of Touch in Virtual and Real Environments." In: *IEEE MultiMedia* 13.3 (2006), pp. 24–30. DOI: [10.1109/MMUL.2006.69](https://doi.org/10.1109/MMUL.2006.69).
- [284] Anne Roudaut, Diana Krusteva, Mike McCoy, Abhijit Karnik, Karthik Ramani, and Sriram Subramanian. "Cubimorph: Designing Modular Interactive Devices." In: *Proceedings of the*

- 2016 *IEEE International Conference on Robotics and Automation*. ICRA'16. IEEE, 2016. DOI: [10.1109/ICRA.2016.7487508](https://doi.org/10.1109/ICRA.2016.7487508).
- [285] Neung Ryu, Woojin Lee, Myung Jin Kim, and Andrea Bianchi. "ElaStick: A Handheld Variable Stiffness Display for Rendering Dynamic Haptic Response of Flexible Object." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST'20. ACM, 2020. DOI: [10.1145/3379337.3415862](https://doi.org/10.1145/3379337.3415862).
- [286] Mikel Sagardia, Katharina Hertkorn, Thomas Hulin, Simon Schätzle, Robin Wolff, Johannes Hummel, Janki Dodiya, and Andreas Gerndt. "VR-OOS: The DLR's Virtual Reality Simulator for Telerobotic On-orbit Servicing with Haptic Feedback." In: *Proceedings of the 2015 IEEE Aerospace Conference*. IEEE, 2015. DOI: [10.1109/AERO.2015.7119040](https://doi.org/10.1109/AERO.2015.7119040).
- [287] Shahabedin Sagheb, Frank Wencheng Liu, Alireza Bahremand, Assegid Kidane, and Robert LiKamWa. "SWISH: A Shifting-Weight Interface of Simulated Hydrodynamics for Haptic Perception of Virtual Fluid Vessels." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST'19. ACM, 2019. DOI: [10.1145/3332165.3347870](https://doi.org/10.1145/3332165.3347870).
- [288] Majed Samad, Elia Gatti, Anne Hermes, Hrvoje Benko, and Cesare Parise. "Pseudo-Haptic Weight: Changing the Perceived Weight of Virtual Objects By Manipulating Control-Display Ratio." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300550](https://doi.org/10.1145/3290605.3300550).
- [289] Marco Santello, Martha Flanders, and John F. Soechting. "Patterns of Hand Motion during Grasping and the Influence of Sensory Guidance." In: *Journal of Neuroscience* 22.4 (2002), pp. 1426–1435. DOI: [10.1523/JNEUROSCI.22-04-01426.2002](https://doi.org/10.1523/JNEUROSCI.22-04-01426.2002).
- [290] Yushi Sato, Takefumi Hiraki, Naruki Tanabe, Haruka Matsukura, Daisuke Iwai, and Kosuke Sato. "Modifying Texture Perception With Pseudo-Haptic Feedback for a Projected Virtual Hand Interface." In: *IEEE Access* 8 (2020), pp. 120473–120488. DOI: [10.1109/ACCESS.2020.3006440](https://doi.org/10.1109/ACCESS.2020.3006440).
- [291] Dominik Schmidt et al. "Kickables: Tangibles for Feet." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI'14. ACM, 2014. DOI: [10.1145/2556288.2557016](https://doi.org/10.1145/2556288.2557016).
- [292] Oliver Schneider, Bruno Fruchard, Dennis Wittchen, Bibhushan Raj Joshi, Georg Freitag, Donald Degraen, and Paul Strohmeier. "Sustainable Haptic Design: Improving Collaboration, Sharing, and Reuse in Haptic Design Research." In: *CHI Conference on Human Factors in Computing*

- Systems Extended Abstracts*. CHI EA'22. ACM, 2022. DOI: [10.1145/3491101.3503734](https://doi.org/10.1145/3491101.3503734).
- [293] Oliver Schneider, Karon MacLean, Colin Swindells, and Kellogg Booth. "Haptic Experience Design: What Hapticians Do and Where They Need Help." In: *International Journal of Human-Computer Studies* 107 (2017), pp. 5–21. DOI: [10.1016/j.ijhcs.2017.04.004](https://doi.org/10.1016/j.ijhcs.2017.04.004).
- [294] Philipp Schoessler, Daniel Windham, Daniel Leithinger, Sean Follmer, and Hiroshi Ishii. "Kinetic Blocks: Actuated Constructive Assembly for Interaction and Display." In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*. UIST'15. ACM, 2015. DOI: [10.1145/2807442.2807453](https://doi.org/10.1145/2807442.2807453).
- [295] Samuel B. Schorr and Allison M. Okamura. "Fingertip Tactile Devices for Virtual Object Manipulation and Exploration." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.3025744](https://doi.org/10.1145/3025453.3025744).
- [296] Thomas Schubert, Frank Friedmann, and Holger Regenbrecht. "The Experience of Presence: Factor Analytic Insights." In: *Presence: Teleoperators and Virtual Environments* 10.3 (2001), pp. 266–281. DOI: [10.1162/105474601300343603](https://doi.org/10.1162/105474601300343603).
- [297] Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. "Microstructures to Control Elasticity in 3D Printing." In: *ACM Transactions on Graphics* 34.4 (2015). 136. DOI: [10.1145/2766926](https://doi.org/10.1145/2766926).
- [298] Eric Schweikardt and Mark D. Gross. "The Robot is the Program: Interacting with RoBlocks." In: *Proceedings of the 2nd International Conference on Tangible and Embedded Interaction*. TEI'o8. ACM, 2008. DOI: [10.1145/1347390.1347427](https://doi.org/10.1145/1347390.1347427).
- [299] Valentin Schwind, Pascal Knierim, Nico Haas, and Niels Henze. "Using Presence Questionnaires in Virtual Reality." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300590](https://doi.org/10.1145/3290605.3300590).
- [300] Valentin Schwind, Pascal Knierim, Cagri Tasci, Patrick Franczak, Nico Haas, and Niels Henze. "“These are not my Hands!” Effect of Gender on the Perception of Avatar Hands in Virtual Reality." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.302560](https://doi.org/10.1145/3025453.302560).

- [301] Vivian Shen, Tucker Rae-Grant, Joe Mullenbach, Chris Harrison, and Craig Shultz. "Fluid Reality: High-Resolution, Untethered Haptic Gloves using Electroosmotic Pump Arrays." In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. UIST'23. ACM, 2023. DOI: [10.1145/3586183.3606771](https://doi.org/10.1145/3586183.3606771).
- [302] Jotaro Shigeyama, Takeru Hashimoto, Shigeo Yoshida, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. "Transcalibur: A Weight Shifting Virtual Reality Controller for 2D Shape Rendering Based on Computational Perception Model." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300241](https://doi.org/10.1145/3290605.3300241).
- [303] Hakim Si-Mohammed, Catarina Lopes-Dias, Maria Duarte, Ferran Argelaguet, Camille Jeunet, Géry Casiez, Gernot R. Müller-Putz, Anatole Lécuyer, and Reinhold Scherer. "Detecting System Errors in Virtual Reality Using EEG Through Error-Related Potentials." In: *Proceedings of the 2020 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'20. IEEE, 2020. DOI: [10.1109/VR46266.2020.00088](https://doi.org/10.1109/VR46266.2020.00088).
- [304] Adalberto L. Simeone, Eduardo Velloso, and Hans Gellersen. "Substitutional Reality: Using the Physical Environment to Design Virtual Reality Experiences." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI'15. ACM, 2015. DOI: [10.1145/2702123.2702389](https://doi.org/10.1145/2702123.2702389).
- [305] Mike Sinclair, Eyal Ofek, Mar Gonzalez-Franco, and Christian Holz. "CapstanCrunch: A Haptic VR Controller with User-supplied Force Feedback." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST'19. ACM, 2019. DOI: [10.1145/3332165.3347891](https://doi.org/10.1145/3332165.3347891).
- [306] Avinash K. Singh, Hsiang-Ting Chen, Yu-Feng Cheng, Jung-Tai King, Li-Wei Ko, Klaus Gramann, and Chin-Teng Lin. "Visual Appearance Modulates Prediction Error in Virtual Reality." In: *IEEE Access* 6 (2018), pp. 24617–24624. DOI: [10.1109/ACCESS.2018.2832089](https://doi.org/10.1109/ACCESS.2018.2832089).
- [307] Avinash K. Singh, Klaus Gramann, Hsiang-Ting Chen, and Chin-Teng Lin. "The Impact of Hand Movement Velocity on Cognitive Conflict Processing in a 3D Object Selection Task in Virtual Reality." In: *Neuroimage* 226 (2021), p. 117578. DOI: [10.1016/j.neuroimage.2020.117578](https://doi.org/10.1016/j.neuroimage.2020.117578).
- [308] Alexa F. Siu, Eric J. Gonzalez, Shenli Yuan, Jason B. Ginsberg, and Sean Follmer. "ShapeShift: 2D Spatial Manipulation and Self-Actuation of Tabletop Shape Displays for Tangible and Haptic Interaction." In: *Proceedings of the 2018 CHI Conference*

- on Human Factors in Computing Systems. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173865](https://doi.org/10.1145/3173574.3173865).
- [309] Antun Skuric, Hasan Sinan Bank, Richard Unger, Owen Williams, and David González-Reyes. "SimpleFOC: A Field Oriented Control (FOC) Library for Controlling Brushless Direct Current (BLDC) and Stepper Motors." In: *Journal of Open Source Software* 7.74 (2022), p. 4232. DOI: [10.21105/joss.04232](https://doi.org/10.21105/joss.04232).
- [310] Mel Slater. "Place Illusion and Plausibility can Lead to Realistic Behaviour in Immersive Virtual Environments." In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 364.1535 (2009), pp. 3549–3557. DOI: <http://doi.org/10.1098/rstb.2009.0138>.
- [311] Mel Slater. *Bayesian Methods in Statistics: From Concepts to Practice*. SAGE Publications Ltd, 2021.
- [312] Mel Slater, David-Paul Pertaub, Chris Barker, and David M Clark. "An Experimental Study on Fear of Public Speaking using a Virtual Environment." In: *CyberPsychology & Behavior* 9.5 (2006), pp. 627–633. DOI: [10.1089/cpb.2006.9.627](https://doi.org/10.1089/cpb.2006.9.627).
- [313] Mel Slater, Amela Sadagic, Martin Usoh, and Ralph Schroeder. "Small-Group Behavior in a Virtual and Real Environment: A Comparative Study." In: *Presence: Teleoperators and Virtual Environments* 9.1 (2000), pp. 37–51. DOI: [10.1162/105474600566600](https://doi.org/10.1162/105474600566600).
- [314] Mel Slater and Anthony Steed. "A Virtual Presence Counter." In: *Presence: Teleoperators and Virtual Environments* 9.5 (2000), pp. 413–434. DOI: [10.1162/105474600566925](https://doi.org/10.1162/105474600566925).
- [315] Mel Slater and Sylvia Wilbur. "A Framework for Immersive Virtual Environments (FIVE): Speculations on the Role of Presence in Virtual Environments." In: *Presence: Teleoperators and Virtual Environments* 6.6 (1997), pp. 603–616. DOI: [10.1162/pres.1997.6.6.603](https://doi.org/10.1162/pres.1997.6.6.603).
- [316] Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. "AIREAL: Interactive Tactile Experiences in Free Air." In: *ACM Transactions on Graphics* 32.4 (2013). 134. DOI: [10.1145/2461912.2462007](https://doi.org/10.1145/2461912.2462007).
- [317] Marco Speicher, Jan Ehrlich, Vito Gentile, Donald Degraen, Salvatore Sorce, and Antonio Krüger. "Pseudo-haptic Controls for Mid-air Finger-based Menu Interaction." In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI EA'19. ACM, 2019. DOI: [10.1145/3290607.3312927](https://doi.org/10.1145/3290607.3312927).

- [318] Misha Sra, Xuhai Xu, and Pattie Maes. "BreathVR: Leveraging Breathing as a Directly Controlled Interface for Virtual Reality Games." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173914](https://doi.org/10.1145/3173574.3173914).
- [319] Mandayam A. Srinivasan and Cagatay Basdogan. "Haptics in Virtual Environments: Taxonomy, Research status, and Challenges." In: *Computers & Graphics* 21.4 (1997), pp. 393–404. DOI: [10.1016/S0097-8493\(97\)00030-7](https://doi.org/10.1016/S0097-8493(97)00030-7).
- [320] Anthony Steed, Eyal Ofek, Mike Sinclair, and Mar Gonzalez-Franco. "A Mechatronic Shape Display Based on Auxetic Materials." In: *Nature Communications* 12 (2021). 4758. DOI: [10.1038/s41467-021-24974-0](https://doi.org/10.1038/s41467-021-24974-0).
- [321] Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. "Analyses of human sensitivity to redirected walking." In: *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology*. VRST'08. ACM, 2008. DOI: [10.1145/1450579.1450611](https://doi.org/10.1145/1450579.1450611).
- [322] Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. "Estimation of Detection Thresholds for Redirected Walking Techniques." In: *IEEE Transactions on Visualization and Computer Graphics* 16.1 (2010), pp. 17–27. DOI: [10.1109/TVCG.2009.62](https://doi.org/10.1109/TVCG.2009.62).
- [323] Carolin Stellmacher, André Zenner, Oscar Javier Ariza Nunez, Ernst Kruijff, and Johannes Schöning. "Continuous VR Weight Illusion by Combining Adaptive Trigger Resistance and Control-Display Ratio Manipulation." In: *Proceedings of the 2023 IEEE Conference Virtual Reality and 3D User Interfaces*. VR'23. IEEE, 2023. DOI: [10.1109/VR55154.2023.00040](https://doi.org/10.1109/VR55154.2023.00040).
- [324] Robert J. Sternberg and Karin Sternberg. *Cognitive Psychology*. Vol. 7th edition. Wadsworth Publishing, 2016.
- [325] Francesca Stival, Stefano Michieletto, Matteo Cognolato, Enrico Pagello, Henning Müller, and Manfredo Atzori. "A Quantitative Taxonomy of Human Hand Grasps." In: *Journal of NeuroEngineering and Rehabilitation* 16 (2019). 28. DOI: [10.1186/s12984-019-0488-x](https://doi.org/10.1186/s12984-019-0488-x).
- [326] Patrick L. Strandholt, Oana A. Dogaru, Niels C. Nilsson, Rolf Nordahl, and Stefania Serafin. "Knock on Wood: Combining Redirected Touching and Physical Props for Tool-Based Interaction in Virtual Reality." In: *Proceedings of the 2020 ACM Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376303](https://doi.org/10.1145/3313831.3376303).

- [327] Evan Strasnick, Christian Holz, Eyal Ofek, Mike Sinclair, and Hrvoje Benko. "Haptic Links: Bimanual Haptics for Virtual Reality Using Variable Stiffness Actuation." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3174218](https://doi.org/10.1145/3173574.3174218).
- [328] Paul Strohmeier and Kasper Hornbæk. "Generating Haptic Textures with a Vibrotactile Actuator." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI'17. ACM, 2017. DOI: [10.1145/3025453.3025812](https://doi.org/10.1145/3025453.3025812).
- [329] Evan A. Suma, Seth Clark, David Krum, Samantha Finkelstein, Mark Bolas, and Zachary Warte. "Leveraging Change Blindness for Redirection in Virtual Environments." In: *Proceedings of the 2011 IEEE Virtual Reality Conference*. IEEE, 2011. DOI: [10.1109/VR.2011.5759455](https://doi.org/10.1109/VR.2011.5759455).
- [330] Evan A. Suma, Zachary Lipps, Samantha Finkelstein, David M. Krum, and Mark Bolas. "Impossible Spaces: Maximizing Natural Walking in Virtual Environments with Self-Overlapping Architecture." In: *IEEE Transactions on Visualization and Computer Graphics* 18.4 (2012), pp. 555–564. DOI: [10.1109/TVCG.2012.47](https://doi.org/10.1109/TVCG.2012.47).
- [331] Ivan E. Sutherland. "A Head-Mounted Three Dimensional Display." In: *Proceedings of the AFIPS Fall Joint Computing Conference*. AFIPS'68. ACM, 1968. DOI: [10.1145/1476589.1476686](https://doi.org/10.1145/1476589.1476686).
- [332] Ryo Suzuki, Hooman Hedayati, Clement Zheng, James L. Bohn, Daniel Szafir, Ellen Yi-Luen Do, Mark D. Gross, and Daniel Leithinger. "RoomShift: Room-scale Dynamic Haptics for VR with Furniture-moving Swarm Robots." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376523](https://doi.org/10.1145/3313831.3376523).
- [333] Ryo Suzuki, Eyal Ofek, Mike Sinclair, Daniel Leithinger, and Mar Gonzalez-Franco. "HapticBots: Distributed Encountered-Type Haptics for VR with Multiple Shape-Changing Mobile Robots." In: *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology*. UIST'21. ACM, 2021. DOI: [10.1145/3472749.3474821](https://doi.org/10.1145/3472749.3474821).
- [334] Haruki Takahashi and Jeeun Kim. "Designing a Hairy Haptic Display Using 3D Printed Hairs and Perforated Plates." In: *Adjunct Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. UIST'22 Adjunct. ACM, 2022. DOI: [10.1145/3526114.3558655](https://doi.org/10.1145/3526114.3558655).
- [335] Yudai Tanaka, Arata Horie, and Xiang 'Anthony' Chen. "DualVib: Simulating Haptic Sensation of Dynamic Mass by Combining Pseudo-Force and Texture Feedback." In: *Proceedings of*

- the 26th ACM Symposium on Virtual Reality Software and Technology*. VRST'20. ACM, 2020. DOI: [10.1145/3385956.3418964](https://doi.org/10.1145/3385956.3418964).
- [336] Yudai Tanaka, Jacob Serfaty, and Pedro Lopes. "Haptic Source-Effector: Full-Body Haptics via Non-Invasive Brain Stimulation." In: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. CHI'24. ACM, 2024. DOI: [10.1145/3613904.3642483](https://doi.org/10.1145/3613904.3642483).
 - [337] Yudai Tanaka, Akifumi Takahashi, and Pedro Lopes. "Interactive Benefits from Switching Electrical to Magnetic Muscle Stimulation." In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. UIST'23. ACM, 2023. DOI: [10.1145/3586183.3606812](https://doi.org/10.1145/3586183.3606812).
 - [338] Yujie Tao, Cheng Yao Wang, Andrew D. Wilson, Eyal Ofek, and Mar Gonzalez-Franco. "Embodying Physics-Aware Avatars in Virtual Reality." In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI'23. ACM, 2023. DOI: [10.1145/3544548.3580979](https://doi.org/10.1145/3544548.3580979).
 - [339] Shan-Yuan Teng, Tzu-Sheng Kuo, Chi Wang, Chi-huan Chiang, Da-Yuan Huang, Liwei Chan, and Bing-Yu Chen. "PuPoP: Pop-up Prop on Palm for Virtual Reality." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST'18. ACM, 2018. DOI: [10.1145/3242587.3242628](https://doi.org/10.1145/3242587.3242628).
 - [340] Shan-Yuan Teng, K. D. Wu, Jacqueline Chen, and Pedro Lopes. "Prolonging VR Haptic Experiences by Harvesting Kinetic Energy from the User." In: *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. UIST'22. ACM, 2022. DOI: [10.1145/3526113.3545635](https://doi.org/10.1145/3526113.3545635).
 - [341] Jordan Tewell and Nimesha Ranasinghe. "A Review of Olfactory Display Designs for Virtual Reality Environments." In: *ACM Computing Surveys* 56.11 (2024). 276. DOI: [10.1145/3665243](https://doi.org/10.1145/3665243).
 - [342] Xavier de Tinguy, Claudio Pacchierotti, Mathieu Emily, Mathilde Chevalier, Aurélie Guignardat, Morgan Guillaudeux, Chloé Six, Anatole Lécuyer, and Maud Marchal. "How Different Tangible and Virtual Objects Can Be While Still Feeling the Same?" In: *Proceedings of the IEEE World Haptics Conference*. WHC'19. IEEE, 2019. DOI: [10.1109/WHC.2019.8816164](https://doi.org/10.1109/WHC.2019.8816164).
 - [343] Xavier de Tinguy, Claudio Pacchierotti, Anatole Lécuyer, and Maud Marchal. "Capacitive Sensing for Improving Contact Rendering With Tangible Objects in VR." In: *IEEE Transactions on Visualization and Computer Graphics* 27.4 (2021), pp. 2481–2487. DOI: [10.1109/TVCG.2020.3047689](https://doi.org/10.1109/TVCG.2020.3047689).

- [344] Cesar Torres, Tim Campbell, Neil Kumar, and Eric Paulos. "HapticPrint: Designing Feel Aesthetics for Digital Fabrication." In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*. UIST'15. ACM, 2015. DOI: [10.1145/2807442.2807492](https://doi.org/10.1145/2807442.2807492).
- [345] Tanh Quang Tran, HyunJu Shin, Wolfgang Stuerzlinger, and JungHyun Han. "Effects of Virtual Arm Representations on Interaction in Virtual Environments." In: *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*. VRST'17. ACM, 2017. DOI: [10.1145/3139131.3139149](https://doi.org/10.1145/3139131.3139149).
- [346] Hsin-Ruey Tsai, Ching-Wen Hung, Tzu-Chun Wu, and Bing-Yu Chen. "ElastOscillation: 3D Multilevel Force Feedback for Damped Oscillation on VR Controllers." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376408](https://doi.org/10.1145/3313831.3376408).
- [347] Hsin-Ruey Tsai, Jun Rekimoto, and Bing-Yu Chen. "ElasticVR: Providing Multilevel Continuously-Changing Resistive Force and Instant Impact Using Elasticity for VR." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300450](https://doi.org/10.1145/3290605.3300450).
- [348] Wen-Jie Tseng, Elise Bonnail, Mark McGill, Mohamed Khamis, Eric Lecolinet, Samuel Huron, and Jan Gugenheimer. "The Dark Side of Perceptual Manipulations in Virtual Reality." In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI'22. ACM, 2022. DOI: [10.1145/3491102.3517728](https://doi.org/10.1145/3491102.3517728).
- [349] Michael T. Turvey. "Dynamic Touch." In: *American Psychologist* 51.11 (1996), pp. 1134–1152. DOI: [10.1037/0003-066X.51.11.1134](https://doi.org/10.1037/0003-066X.51.11.1134).
- [350] Martin Usoh, Kevin Arthur, Mary C. Whitton, Rui Bastos, Anthony Steed, Mel Slater, and Frederick P. Brooks Jr. "Walking> Walking-in-Place> Flying, in Virtual Environments." In: *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH'99. ACM, 1999. DOI: [10.1145/311535.311589](https://doi.org/10.1145/311535.311589).
- [351] Martin Usoh, Ernest Catena, Sima Arman, and Mel Slater. "Using Presence Questionnaires in Reality." In: *Presence: Teleoperators and Virtual Environments* 9.5 (2000), pp. 497–503. DOI: [10.1162/105474600566989](https://doi.org/10.1162/105474600566989).
- [352] Ali Valipour, Mohammad H. Kargozarfard, Mina Rakhshi, Amin Yaghootian, and Hamid M. Sedighi. "Metamaterials and their Applications: An Overview." In: *Proceedings of the Institution of Mechanical Engineers, Part L: Journal of Materials*:

- Design and Applications* 236.11 (2022), pp. 2171–2210. DOI: [10.1177/1464420721995858](https://doi.org/10.1177/1464420721995858).
- [353] Eric-Jan Wagenmakers, Jonathon Love, Maarten Marsman, Tahira Jamil, Alexander Ly, Josine Verhagen, Ravi Selker, Quentin F Gronau, Damian Dropmann, Bruno Boutin, et al. “Bayesian Inference for Psychology. Part II: Example Applications with JASP.” In: *Psychonomic Bulletin & Review* 25 (2018), pp. 58–76. DOI: [10.3758/s13423-017-1323-7](https://doi.org/10.3758/s13423-017-1323-7).
 - [354] Yu-Wei Wang, Yu-Hsin Lin, Pin-Sung Ku, Yōko Miyatake, Yi-Hsuan Mao, Po Yu Chen, Chun-Miao Tseng, and Mike Y. Chen. “JetController: High-speed Ungrounded 3-DoF Force Feedback Controllers using Air Propulsion Jets.” In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI’21. ACM, 2021. DOI: [10.1145/3411764.3445549](https://doi.org/10.1145/3411764.3445549).
 - [355] Yuntao Wang et al. “MoveVR: Enabling Multiform Force Feedback in Virtual Reality Using Household Cleaning Robot.” In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI’20. ACM, 2020. DOI: [10.1145/3313831.3376286](https://doi.org/10.1145/3313831.3376286).
 - [356] Colin Ware, Kevin Arthur, and Kellogg S. Booth. “Fish Tank Virtual Reality.” In: *Proceedings of the INTERACT’93 and CHI’93 Conference on Human Factors in Computing Systems*. CHI’93. ACM, 1993. DOI: [10.1145/169059.169066](https://doi.org/10.1145/169059.169066).
 - [357] Ryoichi Watanabe, Yuichi Itoh, Masatsugu Asai, Yoshifumi Kitamura, Fumio Kishino, and Hideo Kikuchi. “The Soul of ActiveCube: Implementing a Flexible, Multimodal, Three-Dimensional Spatial Tangible Interface.” In: *Computers in Entertainment* 2.4 (2004), p. 15. DOI: [10.1145/1037851.1037874](https://doi.org/10.1145/1037851.1037874).
 - [358] Tzu-Yun Wei, Hsin-Ruey Tsai, Yu-So Liao, Chieh Tsai, Yi-Shan Chen, Chi Wang, and Bing-Yu Chen. “ElastiLinks: Force Feedback between VR Controllers with Dynamic Points of Application of Force.” In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST’20. ACM, 2020. DOI: [10.1145/3379337.3415836](https://doi.org/10.1145/3379337.3415836).
 - [359] Florian Weidner, Gerd Boettcher, Stephanie Arevalo Arboleda, Chenyao Diao, Luljeta Sinani, Christian Kunert, Christoph Gerhardt, Wolfgang Broll, and Alexander Raake. “A Systematic Review on the Visualization of Avatars and Agents in AR and VR displayed using Head-Mounted Displays.” In: *IEEE Transactions on Visualization and Computer Graphics* 29.5 (2023), pp. 2596–2606. DOI: [10.1109/TVCG.2023.3247072](https://doi.org/10.1109/TVCG.2023.3247072).
 - [360] Yannick Weiss, Steeven Villa, Albrecht Schmidt, Sven Mayer, and Florian Müller. “Using Pseudo-Stiffness to Enrich the Haptic Experience in Virtual Reality.” In: *Proceedings of the*

- 2023 *CHI Conference on Human Factors in Computing Systems*. CHI'23. ACM, 2023. DOI: [10.1145/3544548.3581223](https://doi.org/10.1145/3544548.3581223).
- [361] Johann Wentzel, Greg d'Eon, and Daniel Vogel. "Improving Virtual Reality Ergonomics Through Reach-Bounded Non-Linear Input Amplification." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376687](https://doi.org/10.1145/3313831.3376687).
- [362] Michael White, James Gain, Ulysse Vimont, and Daniel Lochner. "The Case for Haptic Props: Shape, Weight and Vibro-Tactile Feedback." In: *Proceedings of the 12th ACM SIGGRAPH Conference on Motion, Interaction and Games*. MIG'19. ACM, 2019. DOI: [10.1145/3359566.3360058](https://doi.org/10.1145/3359566.3360058).
- [363] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. "Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3173660](https://doi.org/10.1145/3173574.3173660).
- [364] Anusha Withana, Daniel Groeger, and Jürgen Steimle. "Tact-too: A Thin and Feel-Through Tattoo for On-Skin Tactile Output." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST'18. ACM, 2018. DOI: [10.1145/3242587.3242645](https://doi.org/10.1145/3242587.3242645).
- [365] 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](https://doi.org/10.1162/105474698565686).
- [366] Dennis Wittchen, Katta Spiel, Bruno Fruchard, Donald Degraen, Oliver Schneider, Georg Freitag, and Paul Strohmeier. "TactJam: An End-to-End Prototyping Suite for Collaborative Design of On-Body Vibrotactile Feedback." In: *Proceedings of the Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI'22. ACM, 2022. DOI: [10.1145/3490149.3501307](https://doi.org/10.1145/3490149.3501307).
- [367] Rui Xavier, José Luís Silva, Rodrigo Ventura, and Joaquim Armando Pires Jorge. "Pseudo-Haptics Survey: Human-Computer Interaction in Extended Reality and Teleoperation." In: *IEEE Access* 12 (2024), pp. 80442–80467. DOI: [10.1109/ACCESS.2024.3409449](https://doi.org/10.1109/ACCESS.2024.3409449).
- [368] Shuyuan Xiao, Tao Wang, Tingting Liu, Chaobiao Zhou, Xiaoyun Jiang, and Jianfa Zhang. "Active Metamaterials and Metadevices: A Review." In: *Journal of Physics D: Applied Physics* 53.50 (2020). 503002. DOI: [10.1088/1361-6463/abaced](https://doi.org/10.1088/1361-6463/abaced).

- [369] Biao Xie, Huimin Liu, Rawan Alghofaili, Yongqi Zhang, Yeling Jiang, Flavio Destri Lobo, Changyang Li, Wanwan Li, Haikun Huang, Mesut Akdere, et al. "A Review on Virtual Reality Skill Training Applications." In: *Frontiers in Virtual Reality* 2 (2021). 645153. DOI: [10.3389/frvir.2021.645153](https://doi.org/10.3389/frvir.2021.645153).
- [370] Yizhong Xin, Xiaojun Bi, and Xiangshi Ren. "Natural Use Profiles for the Pen: An Empirical Exploration of Pressure, Tilt, and Azimuth." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI'12. ACM, 2012. DOI: [10.1145/2207676.2208518](https://doi.org/10.1145/2207676.2208518).
- [371] Kotaro Yamaguchi, Ginga Kato, Yoshihiro Kuroda, Kiyoshi Kiyokawa, and Haruo Takemura. "A Non-Grounded and Encountered-Type Haptic Display Using a Drone." In: *Proceedings of the 2016 Symposium on Spatial User Interaction*. SUI'16. ACM, 2016. DOI: [10.1145/2983310.2985746](https://doi.org/10.1145/2983310.2985746).
- [372] Jackie Yang, Hiroshi Horii, Alexander Thayer, and Rafael Balagas. "VR Grabbers: Ungrounded Haptic Retargeting for Precision Grabbing Tools." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST'18. ACM, 2018. DOI: [10.1145/3242587.3242643](https://doi.org/10.1145/3242587.3242643).
- [373] Willa Yunqi Yang, Yumeng Zhuang, Luke Andre Darcy, Grace Liu, and Alexandra Ion. "Reconfigurable Elastic Metamaterials." In: *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. UIST'22. ACM, 2022. DOI: [10.1145/3526113.3545649](https://doi.org/10.1145/3526113.3545649).
- [374] Boris Yazmir, Miriam Reiner, Hillel Pratt, and Miriam Zacksenhouse. "Brain Responses to Errors During 3D Motion in a Hapto-Visual VR." In: *Proceedings of Haptics: Perception, Devices, Mobility, and Communication*. EuroHaptics'16. Springer, 2016. DOI: [10.1007/978-3-319-42324-1_12](https://doi.org/10.1007/978-3-319-42324-1_12).
- [375] Yan Yixian, Kazuki Takashima, Anthony Tang, Takayuki Tanno, Kazuyuki Fujita, and Yoshifumi Kitamura. "ZoomWalls: Dynamic Walls that Simulate Haptic Infrastructure for Room-scale VR World." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST'20. ACM, 2020. DOI: [10.1145/3379337.3415859](https://doi.org/10.1145/3379337.3415859).
- [376] Shigeo Yoshida, Yuqian Sun, and Hideaki Kuzuoka. "PoCoPo: Handheld Pin-Based Shape Display for Haptic Rendering in Virtual Reality." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI'20. ACM, 2020. DOI: [10.1145/3313831.3376358](https://doi.org/10.1145/3313831.3376358).
- [377] Run Yu and Doug A. Bowman. "Pseudo-Haptic Display of Mass and Mass Distribution During Object Rotation in Virtual Reality." In: *IEEE Transactions on Visualization and Com-*

- puter Graphics* 26.5 (2020), pp. 2094–2103. DOI: [10.1109/TVCG.2020.2973056](https://doi.org/10.1109/TVCG.2020.2973056).
- [378] André Zenner, Chiara Karr, Martin Feick, Oscar Ariza, and Antonio Krüger. “The Detectability of Saccadic Hand Offset in Virtual Reality.” In: *Proceedings of the 29th ACM Symposium on Virtual Reality Software and Technology*. VRST ’23. ACM, 2023. DOI: [10.1145/3611659.3617223](https://doi.org/10.1145/3611659.3617223).
 - [379] André Zenner, Kristin Ullmann, Chiara Karr, Oscar Ariza, and Antonio Krüger. “The Staircase Procedure Toolkit: Psychophysical Detection Threshold Experiments Made Easy.” In: *Proceedings of the 29th ACM Symposium on Virtual Reality Software and Technology*. VRST’23. ACM, 2023. DOI: [10.1145/3611659.3617218](https://doi.org/10.1145/3611659.3617218).
 - [380] André Zenner. *Advancing Proxy-Based Haptic Feedback in Virtual Reality (PhD Thesis)*. Saarland University, 2022. DOI: [10.22028/D291-37879](https://doi.org/10.22028/D291-37879).
 - [381] André Zenner, Chiara Karr, Martin Feick, Oscar Ariza, and Antonio Krüger. “Beyond the Blink: Investigating Combined Saccadic & Blink-Suppressed Hand Redirection in Virtual Reality.” In: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. CHI’24. ACM, 2024. DOI: [10.1145/3613904.3642073](https://doi.org/10.1145/3613904.3642073).
 - [382] André Zenner and Antonio Krüger. “Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality.” In: *IEEE Transactions on Visualization and Computer Graphics* 23.4 (2017), pp. 1285–1294. DOI: [10.1109/TVCG.2017.2656978](https://doi.org/10.1109/TVCG.2017.2656978).
 - [383] André Zenner and Antonio Krüger. “Drag:on – A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift.” In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI’19. ACM, 2019. DOI: [10.1145/3290605.3300441](https://doi.org/10.1145/3290605.3300441).
 - [384] André Zenner and Antonio Krüger. “Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality.” In: *Proceedings of the 2019 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR’19. IEEE, 2019. DOI: [10.1109/VR.2019.8798143](https://doi.org/10.1109/VR.2019.8798143).
 - [385] André Zenner, Kora Persephone Regitz, and Antonio Krüger. “Blink-Suppressed Hand Redirection.” In: *Proceedings of the 2021 IEEE Virtual Reality and 3D User Interfaces*. VR’21. IEEE, 2021. DOI: [10.1109/VR50410.2021.00028](https://doi.org/10.1109/VR50410.2021.00028).

- [386] André Zenner, Kristin Ullmann, and Antonio Krüger. "Combining Dynamic Passive Haptics and Haptic Retargeting for Enhanced Haptic Feedback in Virtual Reality." In: *IEEE Transactions on Visualization and Computer Graphics* 27.5 (2021), pp. 2627–2637. DOI: [10.1109/TVCG.2021.3067777](https://doi.org/10.1109/TVCG.2021.3067777).
- [387] Yiwei Zhao and Sean Follmer. "A Functional Optimization Based Approach for Continuous 3D Retargeted Touch of Arbitrary, Complex Boundaries in Haptic Virtual Reality." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI'18. ACM, 2018. DOI: [10.1145/3173574.3174118](https://doi.org/10.1145/3173574.3174118).
- [388] Yiwei Zhao, Lawrence H. Kim, Ye Wang, Mathieu Le Goc, and Sean Follmer. "Robotic Assembly of Haptic Proxy Objects for Tangible Interaction and Virtual Reality." In: *Proceedings of the ACM Conference on Interactive Surfaces and Spaces*. ISS'17. ACM, 2017. DOI: [10.1145/3132272.3134143](https://doi.org/10.1145/3132272.3134143).
- [389] Joshua Z. Zheng, Sara De La Rosa, and Aaron M. Dollar. "An Investigation of Grasp Type and Frequency in Daily Household and Machine Shop Tasks." In: *Proceedings of the 2011 IEEE International Conference on Robotics and Automation*. ICRA'11. IEEE, 2011. DOI: [10.1109/ICRA.2011.5980366](https://doi.org/10.1109/ICRA.2011.5980366).
- [390] Yuqi Zhou and Voicu Popescu. "Tapping with a Handheld Stick in VR: Redirection Detection Thresholds for Passive Haptic Feedback." In: *Proceedings of the 2022 IEEE Conference on Virtual Reality and 3D User Interfaces*. VR'22. IEEE, 2022. DOI: [10.1109/VR51125.2022.00026](https://doi.org/10.1109/VR51125.2022.00026).
- [391] Kening Zhu, Taizhou Chen, Feng Han, and Yi-Shiun Wu. "HapTwist: Creating Interactive Haptic Proxies in Virtual Reality Using Low-Cost Twistable Artefacts." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI'19. ACM, 2019. DOI: [10.1145/3290605.3300923](https://doi.org/10.1145/3290605.3300923).

COLOPHON

This document was typeset using the typographical look-and-feel classicthesis developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". classicthesis is available for both L^AT_EX and L^YX:

<https://bitbucket.org/amiede/classicthesis/>

Happy users of classicthesis usually send a real postcard to the author, a collection of postcards received so far is featured here:

<http://postcards.miede.de/>

Final Version as of March 30, 2025 (classicthesis).