
SAARLAND UNIVERSITY

Faculty of Mathematics and Computer Science
Department of Computer Science
Dissertation



Advancing Proxy-Based Haptic Feedback 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
André Zenner (M.Sc.)
Saarbrücken
2022

Date of the Colloquium: October 28, 2022

Dean: Professor Dr. Jürgen Steimle

Members of the Examination Board:

Chair: Professor Dr. Christian Theobalt

Reporter: Professor Dr. Antonio Krüger

Dr. Anatole Lécuyer

Scientific Assistant: Dr. Maximilian Altmeyer

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; all URLs were last retrieved on June 17, 2022). 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.

Acknowledgements

With more than half a decade in the making, this thesis could be realized only thanks to the continued support and encouragement of many great people, to which I want to say “thank you!”

First, I thank **Antonio Krüger** for the opportunity to pursue my PhD in his research group. To me, Antonio has always been a great mentor and ever since I first talked to Antonio as an undergraduate student, I was impressed by his optimism. I learned a lot from Antonio (not only about research, but also about topics like sailing, pinball machines, movies, juggling, climbing, and biking) and I thank him for always encouraging me, supporting me, and for providing valuable feedback – be it about prototypes, experiments, or, say, on how I can improve the wiring of my do-it-yourself e-bike (which I would not have built if it wasn’t for the supportive and sporty atmosphere in Antonio’s group). Also, what I thank Antonio for especially is the enormous amount of freedom and trust that I had the pleasure to experience throughout my entire PhD. I was very happy that I could always approach the topics that interested me in the ways that I wanted to – a circumstance that I know cannot be taken for granted.

With my thesis now completed, I also thank **Anatole Lécuyer** for agreeing to review it. My conversations with Anatole at my very first scientific conference (IEEE VR 2017 in Los Angeles) really motivated me to continue this endeavor and shaped my view of the VR research community. Hence, I feel honored to have you in my committee and I hope you enjoy reading this thesis.

On my exciting PhD journey, I had the pleasure to spend time with a lot of great people – far too many to name them all here – and I want to thank everyone at the **Ubiquitous Media Technology Lab**. In particular, I thank **Denise Kahl** and **Nico Herbig** for creating the super friendly and uncomplicated atmosphere in our office, and for always supporting me when needed. I especially want to thank Nico for being a great friend since our very first days at the university. The thought that we jointly worked our ways through, from Bachelor’s over Master’s to the end of our PhDs, makes me smile every time I think about it. I know that I can always count on you, be it in work-related matters or, say, on a sailing boat at sea, and I am always very happy and optimistic when I know you in my team.

My appreciation goes also to **Donald Degraen**, **Felix Kosmalla**, and **Martin Feick**, who I always enjoy spending time and co-authoring papers with. Through you guys, I learned a lot about fabrication, prototyping, climbing, Scotland, whisky, e-bikes, brewing beer, and smoking salmons, and I want to say “thank you!” for always offering pragmatic feedback and a helping hand. Let’s make more cool papers together in the future!

A big “thank you!” goes also to **Pascal Lessel**, **Florian Daiber**, and **Frederic Kerber**, who always shared their years of experience with me, and from whom I learned many important things, such as how to approach project proposals, how to organize teaching activities, how to (not) play the Werewolves card game, and

how much (and fast) human beings can move without using a motor (swimming excluded). Furthermore, even though we rarely worked together on a specific project, I always enjoy it when I can spend time with **Maximilian Altmeyer** and **Marcel Köster**, because I know that I will leave the conversation in a great mood and more often than not with new ideas and helpful advice. Also, without the company of Nico and Max, the Software Campus programme would not have been the same (not at all). Finally, I also want to thank **Gerrit Kahl**, who was my advisor back when I was still a student research assistant, for his patience and support when working with the early version of me.

I am very grateful also to each of my co-authors and project partners. Besides the ones already mentioned, my thank goes in particular to **Oscar Ariza** and **Frank Steinicke**, with whom I have the honor to continue advancing VR haptics in our joint VHIVE project. Moreover, my appreciation goes to **Adalberto Simeone** and **Niels Nilsson** for the pleasant collaborations, which I hope we can continue in the future. Speaking of collaborators, I am also very thankful to the student research assistants that supported me during my PhD, namely (in alphabetical order), **Chiara Karr**, **Sören Klingner**, **David Liebemann**, **Akhmajon Makhsadov**, **Fajar Nugroho**, **Marc Ruble**, and **Kristin Ullmann**. My thanks goes also to **Aline** and **Dominik Barré** as well as **Eva Gehring** for their help in the production of illustrations and videos. Moreover, I thank everyone who chose me as their Bachelor's or Master's thesis advisor, especially **Hannah Kriegler** and **Kora Regitz**, whose hard work eventually became part of this thesis. Also, I thank **Tim Düwel**, whose Master's research connects to the topics in this dissertation. Finally, my appreciation goes to all study participants.

A huge "thanks for all your support!" further goes to **Iris Lambrecht**, **Gundula Kleiner**, **Lisa-Marie Merziger**, and **Christine Pyttlik** for helping me solve the countless organizational issues along the way, and to **Margaret De Lap** for proof-reading many of my publications and teaching me a lot about English writing.

My greatest "thanks!" eventually goes to two groups of people. Firstly, to my parents **Jutta** and **Markus Zenner** for *always* supporting me, *no matter* what I am up to, and for sponsoring ≥ 24 semesters of computer science studies. *You are the best!* In addition, the second *You are the best!* award goes to my friends, with whom I always have a great time recharging my batteries.

Last but not least, I thank several institutions for their support. First of all, I especially thank the **German Research Center for Artificial Intelligence** (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI) and **Saarland University** for funding me throughout my PhD. In this context, my thank goes also to the **German Federal Ministry of Education and Research** (Bundesministerium für Bildung und Forschung, BMBF) and the **German Research Foundation** (Deutsche Forschungsgemeinschaft, DFG) for partly funding the research in this thesis. Furthermore, I thank **Scheer Holding** and the team at the **EIT ICT Labs Germany GmbH** for their support during my participation in the **Software Campus**.

Abstract

This thesis advances haptic feedback for Virtual Reality (VR). Our work is guided by Sutherland's 1965 vision of the ultimate display, which calls for VR systems to control the existence of matter. To push towards this vision, we build upon proxy-based haptic feedback, a technique characterized by the use of passive tangible props. The goal of this thesis is to tackle the central drawback of this approach, namely, its inflexibility, which yet hinders it to fulfill the vision of the ultimate display. Guided by four research questions, we first showcase the applicability of proxy-based VR haptics by employing the technique for data exploration. We then extend the VR system's control over users' haptic impressions in three steps. First, we contribute the class of Dynamic Passive Haptic Feedback (DPHF) alongside two novel concepts for conveying kinesthetic properties, like virtual weight and shape, through weight-shifting and drag-changing proxies. Conceptually orthogonal to this, we study how visual-haptic illusions can be leveraged to unnoticeably redirect the user's hand when reaching towards props. Here, we contribute a novel perception-inspired algorithm for Body Warping-based Hand Redirection (HR), an open-source framework for HR, and psychophysical insights. The thesis concludes by proving that the combination of DPHF and HR can outperform the individual techniques in terms of the achievable flexibility of the proxy-based haptic feedback.

Zusammenfassung

Diese Arbeit widmet sich haptischem Feedback für Virtual Reality (VR) und ist inspiriert von Sutherlands Vision des ultimativen Displays, welche VR-Systemen die Fähigkeit zuschreibt, Materie kontrollieren zu können. Um dieser Vision näher zu kommen, baut die Arbeit auf dem Konzept proxy-basierter Haptik auf, bei der haptische Eindrücke durch anfassbare Requisiten vermittelt werden. Ziel ist es, diesem Ansatz die für die Realisierung eines ultimativen Displays nötige Flexibilität zu verleihen. Dazu bearbeiten wir vier Forschungsfragen und zeigen zunächst die Anwendbarkeit proxy-basierter Haptik durch den Einsatz der Technik zur Datenexploration. Anschließend untersuchen wir in drei Schritten, wie VR-Systeme mehr Kontrolle über haptische Eindrücke von Nutzern erhalten können. Hierzu stellen wir Dynamic Passive Haptic Feedback (DPHF) vor, sowie zwei Verfahren, die kinästhetische Eindrücke wie virtuelles Gewicht und Form durch Gewichtsverlagerung und Veränderung des Luftwiderstandes von Requisiten vermitteln. Zusätzlich untersuchen wir, wie visuell-haptische Illusionen die Hand des Nutzers beim Greifen nach Requisiten unbemerkt umlenken können. Dabei stellen wir einen neuen Algorithmus zur Body Warping-based Hand Redirection (HR), ein Open-Source-Framework, sowie psychophysische Erkenntnisse vor. Abschließend zeigen wir, dass die Kombination von DPHF und HR proxy-basierte Haptik noch flexibler machen kann, als es die einzelnen Techniken alleine können.

List of Publications

Parts of this dissertation, including text passages, figures, and tables, as well as ideas, concepts, implementations, applications, studies, results, discussions, and conclusions presented in this work have previously been published. The following list outlines the publications that cover central contributions of this thesis and have been incorporated into this dissertation.

Journal Articles

- **Zenner, A.** and Krüger, A. (2017). Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1285–1294.
(appears in chapters 4 and 5)
- **Zenner, A.**, Makhsoodov, A., Klingner, S., Liebemann, D., and Krüger, A. (2020c). Immersive Process Model Exploration in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(5):2104–2114.
(appears in chapter 3)
- **Zenner, A.**, Ullmann, K., and Krüger, A. (2021c). Combining Dynamic Passive Haptics and Haptic Retargeting for Enhanced Haptic Feedback in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2627–2637.
(appears in chapter 12)

Full Conference Papers

- **Zenner, A.** and Krüger, A. (2019a). Drag:on – A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’19, pages 1–12. ACM.
(appears in chapter 6)
- **Zenner, A.** and Krüger, A. (2019b). Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’19, pages 47–55. IEEE.
(appears in chapter 8)
- **Zenner, A.**, Regitz, K. P., and Krüger, A. (2021b). Blink-Suppressed Hand Redirection. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’21, pages 75–84. IEEE.
(appears in chapter 9)

Magazine Article

- Nilsson, N. C., **Zenner, A.**, and Simeone, A. L. (2021a). Propping Up Virtual Reality With Haptic Proxies. *IEEE Computer Graphics and Applications*, 41(05):104–112.
(appears in chapter 2)

Late-Breaking Works and Posters

- **Zenner, A.**, Klingner, S., Liebemann, D., Makhsadov, A., and Krüger, A. (2019b). Immersive Process Models. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'19, pages 1–6. ACM.
(appears in chapter 3)
- **Zenner, A.** and Krüger, A. (2020). Shifting & Warping: A Case for the Combined Use of Dynamic Passive Haptics and Haptic Retargeting in VR. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST'20 Adjunct, pages 1–3. ACM.
(appears in chapter 12)
- **Zenner, A.**, Kriegler, H. M., and Krüger, A. (2021a). HaRT – The Virtual Reality Hand Redirection Toolkit. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'21, pages 1–7. ACM.
(appears in chapter 10)

Demonstration Paper

- **Zenner, A.**, Degraen, D., Daiber, F., and Krüger, A. (2020a). Demonstration of Drag:on – A VR Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'20, pages 1–4. ACM.
(appears in chapter 6)

Position Papers

- **Zenner, A.** (2020). Enhancing Proxy-Based Haptics in Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, VRW'20, pages 549–550. IEEE.
(appears in chapter 1)
- Nilsson, N. C., **Zenner, A.**, Simeone, A. L., Degraen, D., and Daiber, F. (2021b). Haptic Proxies for Virtual Reality: Success Criteria and Taxonomy. In *Proceedings of the Workshop on Everyday Proxy Objects for Virtual Reality at the ACM Conference on Human Factors in Computing Systems*, EPO4VR'21, pages 1–5.
(appears in chapter 2)

Patent

- **Zenner, A.** and Krüger, A. (2021). Hand-Virtual Reality/Augmented Reality-Steuergerät, Virtual Reality/Augmented Reality-System mit demselben sowie Verfahren zur Simulation der Haptik. *German Patent 10 2019 105 854*. June 05, 2021.
(appears in chapter 6)

Supervised Bachelor's Theses

- Kriegler, H. M. (2020). A Toolkit for Hand Redirection in Virtual Reality. Bachelor's thesis, Saarland University. Advisor: **Zenner, A.**
(presented in [Zenner et al., 2021a]; appears in chapter 10)
- Regitz, K. P. (2020). Leveraging Blink Suppression for Hand Redirection in Virtual Reality. Bachelor's thesis, Saarland University. Advisor: **Zenner, A.**
(presented in [Zenner et al., 2021b]; appears in chapter 9)

In total, the author of this dissertation has supervised four completed Bachelor's theses (with two of them contributing to this dissertation as outlined above), and four completed Master's theses. At the time of submission of this dissertation, the supervision of four additional Bachelor's theses and one additional Master's thesis is still ongoing.

In addition to the publications mentioned above, the author of this dissertation has contributed to several further publications in the fields of Virtual Reality (VR) and Haptics. Those related to the topics of this dissertation are mentioned in the text where applicable and can be found in the Bibliography.

List of Contributions

In the following, we provide an overview of the parts of this thesis that were contributed by the thesis author (column **Author**). We further list which contributions were the joint work of the author and his collaborators (column **Author and Collaborator(s)**), and which parts of the work presented here were contributed by the collaborators (students and research assistants) while being supervised by the author (column **Collaborator(s) under Guidance of Author**). To complete the overview, we also list which contributions were primarily made by collaborators (column **Collaborator(s) only**).

Note:

Collaborators include the co-authors of the publications indicated in the table. Where applicable, we also indicate the names of additional collaborators that are no co-authors (typically students and research assistants) and mark the contributions they were involved in with a “(+)”. All the chapters not indicated in this list were contributed by the author for this thesis.

Contributions of			
Author	Author and Collaborator(s)	Collaborator(s) under Guidance of Author	Collaborator(s) only
<i>Research Approach</i>			
paper writing	idea	–	–
PhD thesis writing			
<i>Taxonomy of Techniques for Achieving Similarity and Colocation</i>			
PhD thesis writing	taxonomy structure classification criteria paper writing	–	idea
<i>Immersive Process Model Exploration</i>			
paper writing	idea	implementation	–
PhD thesis writing	concept development study design analysis of results	study execution	–
<i>Dynamic Passive Haptic Feedback</i>			
paper writing	idea	–	–
PhD thesis writing	classification criteria		

Contributions of			
Author	Author and Collaborator(s)	Collaborator(s) under Guidance of Author	Collaborator(s) only
<i>Shifty -- Dynamic Passive Haptic Feedback Based on Weight Shift</i>			
Chapter 5 – References: [Zenner and Krüger, 2017], GitHub			
Collaborator(s): co-author, Fajar Nugroho (+), Marc Ruble (+)			
implementation	idea	demo application (+)	–
study execution	concept development		
analysis of results	study design		
paper writing			
open-source repository			
PhD thesis writing			
<i>Drag:on -- Dynamic Passive Haptic Feedback Based on Drag and Weight Shift</i>			
Chapter 6 – References: [Zenner and Krüger, 2019a; Zenner et al., 2020a], GitHub			
Collaborator(s): co-authors, Akhmadjon Makhsadov (+), Sören Klingner (+)			
study design	idea	–	–
study execution	concept development		
analysis of results	implementation (+)		
paper writing			
open-source repository			
PhD thesis writing			
<i>Continuous Hand Redirection</i>			
Chapter 8 – References: [Zenner and Krüger, 2019b]			
Collaborator(s): co-author			
implementation	idea	–	–
study design	concept development		
study execution			
analysis of results			
paper writing			
PhD thesis writing			
<i>Blink-Suppressed Hand Redirection</i>			
Chapter 9 – References: [Regitz, 2020; Zenner et al., 2021b]			
Collaborator(s): co-authors			
paper writing	idea	implementation	–
PhD thesis writing	concept development	study execution	
	study design	BSc thesis writing	
	analysis of results		
<i>Unity Staircase Procedure Toolkit</i>			
Chapter 9 – References: GitHub			
Collaborator(s): Kristin Ullmann			
idea	concept development	implementation	–
	open-source repository		

Contributions of			
Author	Author and Collaborator(s)	Collaborator(s) under Guidance of Author	Collaborator(s) only
<i>Hand Redirection Toolkit</i>			
Chapter 10 – References: [Kriegler, 2020; Zenner et al., 2021a], GitHub			
Collaborator(s): co-authors, Kristin Ullmann (+), Chiara Karr (+)			
idea	concept development	implementation	–
paper writing	study design	study execution	
PhD thesis writing	open-source repository (+)	analysis of results	
		BSc thesis writing	
<i>Combining Dynamic Passive Haptic Feedback and Hand Redirection</i>			
Chapter 12 – References: [Zenner and Krüger, 2020; Zenner et al., 2021c]			
Collaborator(s): co-authors, David Liebemann (+), Sören Klingner (+), Marc Ruble (+)			
concept development	idea	–	–
study design	implementation (+)		
paper writing	study execution		
PhD thesis writing	analysis of results		

Contents

I	Introduction and Related Work	1
1	Introduction	3
1.1	Virtual Reality	3
1.2	Experiencing Haptics in Virtual Reality	5
1.2.1	The Role of Haptic Feedback	6
1.2.2	Approaches to Haptic Feedback in Virtual Reality	6
1.3	Motivation and Problem Statement	8
1.4	Research Questions	10
1.5	Methods and Approach	12
1.6	Contributions	13
1.7	Thesis Outline	15
2	Background and Related Work	19
2.1	Virtual Reality	19
2.1.1	Historical Roots of Virtual Reality	19
2.1.2	What is Virtual Reality?	21
2.1.3	Immersion and Presence	23
2.1.4	Technical Realization and Application Areas	26
2.1.5	Conclusion	30
2.2	Human Perception	30
2.2.1	Basics of Human Perception	31
2.2.2	Visual Perception	32
2.2.3	Haptic Perception	35
2.2.4	Multisensory Perception	40
2.2.5	Conclusion	42
2.3	Haptic Feedback for Virtual Reality	43
2.3.1	Classification	43
2.3.2	Active Haptic Feedback	46
2.3.3	Passive Haptic Feedback	53
2.3.4	Conclusion	60

2.4	Central Challenges for Proxy-Based Haptic Feedback	61
2.4.1	The Challenge of <i>Similarity</i>	62
2.4.2	The Challenge of <i>Colocation</i>	62
2.4.3	Orthogonality & Classification of Solutions	63
2.5	Physical Approaches to Similarity and Colocation	64
2.5.1	Offline Physical Strategies	65
2.5.2	Real-Time Physical Strategies	67
2.5.3	Conclusion	76
2.6	Virtual Approaches to Similarity and Colocation	76
2.6.1	Offline Virtual Strategies	77
2.6.2	Real-Time Virtual Strategies	78
2.6.3	Conclusion	98
2.7	Combined Approaches to Similarity and Colocation	98
2.7.1	Encountered-Type Haptics & Haptic Retargeting	99
2.7.2	Dynamic Passive Haptics & Haptic Retargeting	99
2.7.3	Conclusion	100
2.8	Research Methodology	100
2.8.1	Theoretical Aspects	100
2.8.2	Technical Aspects	101
2.8.3	Virtual Reality-Related Aspects	101
2.8.4	Task-Related Aspects	103
2.8.5	Perception-Related Aspects	103
2.8.6	Other Aspects	105
II	Proxy-Based Haptics for VR – A Novel Application	107
3	Immersive Process Model Exploration	109
3.1	Introduction	110
3.1.1	Immersive Data Analysis	110
3.1.2	Process Models	111
3.1.3	Motivation	112
3.2	The <i>Immersive Process Model Exploration</i> System	113
3.2.1	2D to 3D Mapping	114
3.2.2	Logical Walkthrough	115
3.2.3	Haptic Interactions	117
3.3	Evaluation of <i>Immersive Process Model Exploration</i>	121

3.3.1	Participants	123
3.3.2	Apparatus	123
3.3.3	Procedure	124
3.3.4	Design	124
3.3.5	Results	125
3.4	Discussion of <i>Immersive Process Model Exploration</i>	127
3.4.1	User Opinions: Benefits and Drawbacks of 2D and VR . .	127
3.4.2	When to Choose VR and When to Choose 2D?	128
3.4.3	The Impact of Passive Haptic Proxies	129
3.5	Conclusion & Contribution to the Research Questions	130
III	Enhancing Proxy-Based Haptics – The <i>Physical</i> Approach	133
4	Dynamic Passive Haptic Feedback	135
4.1	Introduction	136
4.1.1	The Gap in the Active-Passive Haptics Continuum	136
4.1.2	The Lack of Varying Kinesthetic Proxy Sensations	137
4.1.3	Motivation	137
4.2	Definition	138
4.3	Discussion	138
4.4	Conclusion & Contribution to the Research Questions	140
5	Shifty – Dynamic Passive Haptics Based on Weight Shift	141
5.1	Introduction	142
5.2	Concept of <i>Shifty</i>	143
5.3	Implementation of <i>Shifty</i>	144
5.3.1	Hardware	145
5.3.2	Software	146
5.4	Evaluation of <i>Shifty</i>	147
5.4.1	Experiment 1: Change in Object Length and Thickness .	147
5.4.2	Experiment 2: Picking Up Virtual Objects	152
5.5	Discussion of <i>Shifty</i>	156
5.5.1	Discussion of Experiment 1	156
5.5.2	Discussion of Experiment 2	157
5.5.3	From DPHF to AHF	159
5.5.4	Limitations	160
5.6	Applying <i>Shifty</i> in Different Application Domains	161

5.7	Conclusion & Contribution to the Research Questions	162
6	Drag:on – Dynamic Passive Haptics Based on Drag and Weight Shift	165
6.1	Introduction	166
6.2	Concept of <i>Drag:on</i>	167
6.3	Implementation of <i>Drag:on</i>	169
6.3.1	Hardware	170
6.3.2	Software	171
6.4	Evaluation of <i>Drag:on</i>	172
6.4.1	Participants	172
6.4.2	Apparatus	172
6.4.3	Procedure	173
6.4.4	Design	176
6.4.5	Results	177
6.5	Discussion of <i>Drag:on</i>	179
6.5.1	Discussion of Part 1	179
6.5.2	Discussion of Part 2	180
6.5.3	Recommendations	180
6.5.4	Application Areas	181
6.5.5	Limitations	181
6.6	Conclusion & Contribution to the Research Questions	182
7	Concluding Remarks	185
IV	Enhancing Proxy-Based Haptics – The <i>Virtual</i> Approach	187
8	Continuous Hand Redirection	189
8.1	Introduction	191
8.2	<i>Continuous Hand Redirection</i>	192
8.2.1	Horizontal Warping	193
8.2.2	Vertical Warping	194
8.2.3	Gain Warping	194
8.3	Evaluation of <i>Continuous Hand Redirection</i>	194
8.3.1	Participants	196
8.3.2	Apparatus	196
8.3.3	Procedure	197
8.3.4	Design	199

8.3.5	Results	200
8.4	Discussion of <i>Continuous Hand Redirection</i>	203
8.4.1	Detection Thresholds	203
8.4.2	Distraction & Subjective Impressions	206
8.4.3	Limitations	206
8.5	Conclusion & Contribution to the Research Questions	207
9	Blink-Suppressed Hand Redirection	209
9.1	Introduction	210
9.2	<i>Blink-Suppressed Hand Redirection</i>	211
9.2.1	Concept	211
9.2.2	Assumptions	213
9.2.3	Algorithm	214
9.3	Evaluation of <i>Blink-Suppressed Hand Redirection</i>	218
9.3.1	Hypotheses	218
9.3.2	Participants	219
9.3.3	Apparatus	219
9.3.4	Procedure	220
9.3.5	Design	221
9.3.6	Results	222
9.4	Discussion of <i>Blink-Suppressed Hand Redirection</i>	224
9.4.1	The Detectability of $BSHR_{+0\%}$ (<i>Mode 1</i>) and <i>Cheng</i>	224
9.4.2	The Detectability of Combined Redirection (<i>Mode 2</i>)	225
9.4.3	Limitations	226
9.5	Conclusion & Contribution to the Research Questions	226
10	Hand Redirection Toolkit	229
10.1	Introduction	230
10.2	<i>The Hand Redirection Toolkit</i>	231
10.2.1	Overview	232
10.2.2	Hand Redirection Component	232
10.2.3	User Movement Component	233
10.2.4	Analysis Component	234
10.3	Evaluation of the <i>Hand Redirection Toolkit</i>	235
10.3.1	Methodology	235
10.3.2	Participants	236
10.3.3	Procedure	236

10.3.4	Results	237
10.4	Discussion of the <i>Hand Redirection Toolkit</i>	238
10.4.1	Benefits	238
10.4.2	Limitations	239
10.5	Conclusion & Contribution to the Research Questions	239
11	Concluding Remarks	241
V	Enhancing Proxy-Based Haptics – The <i>Combined</i> Approach	243
12	Combining Dynamic Passive Haptics and Hand Redirection	245
12.1	Introduction	246
12.2	<i>Combined DPHF and HR</i>	247
12.2.1	Proof-of-Concept Scenario	247
12.2.2	Haptic Rendering Techniques	249
12.3	Evaluation of <i>Combined DPHF and HR</i>	251
12.3.1	Metrics	252
12.3.2	Thought Experiments	253
12.3.3	Hypotheses	256
12.3.4	User Experiments	256
12.4	Discussion of <i>Combined DPHF and HR</i>	266
12.4.1	Enhancing <i>Similarity</i>	266
12.4.2	Enhancing <i>Colocation</i>	267
12.4.3	Beyond the Proof-of-Concept Scenario	268
12.4.4	Limitations	269
12.5	Conclusion & Contribution to the Research Questions	269
VI	General Conclusions	273
13	Summary and Contributions	275
13.1	Summary	275
13.1.1	Applying Proxy-Based Haptic Feedback	276
13.1.2	Improving Proxy-Based Haptic Feedback	277
13.2	Contributions	281
14	Future Work and Closing Remarks	287
14.1	Future Work	287

14.1.1 Advancing Systems and Applications	287
14.1.2 Advancing the Physical Approach	288
14.1.3 Advancing the Virtual Approach	291
14.1.4 Advancing the Combined Approach	294
14.2 Closing Remarks	294
List of Figures	297
List of Tables	301
List of Algorithms	303
Abbreviations	305
Bibliography	307

Part I

Introduction and Related Work



Chapter 1

Introduction

Virtual Reality (VR) is currently on its way to becoming a ubiquitous technology that bears the potential to revolutionize how we experience digital data and the way we interact with it. This thesis advances the field of VR in manifold ways with a dedication to the perceptual dimension of haptics. In this chapter, we introduce the fundamental motivation for the research conducted in this thesis and the problems we address throughout this work. We summarize the principal research questions guiding our efforts, briefly outline the contributions we make, and provide an overview of the structure of this thesis.

1.1 Virtual Reality

As humans, we have always been able to experience the physical world around us with all our senses. We can see and hear other people in the environment, can touch and feel objects we interact with, and can smell and taste our food. However, with the arrival of the Information Age in the last century, our everyday world changed. Digital content has become a prominent part of our lives and started to accompany our physical world. Since we cannot see, hear, or feel digital data the way we can perceive physical artifacts, there is a fundamental need for appropriate interfaces to access and interact with digital data.

Virtual Reality (VR) represents one of the highest evolutionary levels of human-computer interfaces as it is tailored to the human senses to a truly unique extent. VR *immerses* users in digital content by catering multiple senses in a way that allows them to experience a simulated, virtual world [Butz and Krüger, 2017; p. 232 f.]. By this, the technology enables novel ways to experience digital content and to interact with it. VR takes on the challenging task of transporting the user into alternative realities that, in the optimal case, are indistinguishable from



Figure 1.1: A typical VR system that conveys the look and sound of an Immersive Virtual Environment (IVE) through a Head-Mounted Display (HMD) and headphones. Interactions with virtual objects are performed by means of handheld VR controllers, which usually provide basic haptic feedback through vibration.

reality in terms of their perceptual quality and convincibility while empowering users to experience things they could not experience in the real world.

While the concept of VR attracts significant attention these days, its historical roots go back as far as to the 19th century [Jerald, 2015; pp. 15 ff.]. The advances in computing and hardware manufacturing achieved in the last decades have made VR technology affordable and widely available on the consumer market [Butz and Krüger, 2017; p. 223]. By investing only as much as, for example, in a new smartphone, users can now buy commercial VR equipment as shown in Figure 1.1 and immerse themselves in interactive, virtual worlds at home.

With the entertainment sector being the current driving force for the commercial VR industry, VR's growing success can be observed, for example, in Valve's review of the year 2020. Their report on SteamVR indicates growing VR game sales, growing VR playtimes, "*over 104 million PC VR sessions*" that took place on the gaming platform in 2020, and concludes that the "*[d]emand for PC VR is strong and growing as more than 1.7 million users on Steam experienced VR games for the first time in 2020*".¹ PricewaterhouseCoopers even projects that the growth of the VR sector will be the fastest among the entertainment and media segments in their forecast period from 2021 to 2025². But entertainment is not the only application for VR technology. Fields that take advantage of the unique capabilities of VR are varied, spanning, for example, domains like simulation, training, retail, creative work, communication, and the healthcare sector [Jerald, 2015; pp. 12 f.].

So the question arises: Are we there yet? Have we finally achieved the vision of the "*ultimate display*" brought forward by VR pioneer Ivan Sutherland [1965], who framed VR as a "*looking glass into a mathematical wonderland*" that simulates virtual worlds to all the human senses?

Taking a closer look, one can note that the graphical and auditory rendering for VR has indeed reached impressive quality levels. As the example on the left in

¹Steam - 2020 Year in Review blog post by Valve Corporation. <https://bit.ly/3BkJiBn>

²PwC Global Entertainment & Media Outlook 2021–2025, www.pwc.com/outlook. <https://pwc.to/3m1UiK>



Figure 1.2: Left: Screenshot of the VR game *The Climb* by Crytek (image taken from official press kit³; © 2016 Crytek). Right: No matter which virtual pan the user will interact with in this VR kitchen store, all the pans will feel the same and unrealistically lightweight when picked up with today's VR controllers.

Figure 1.2 demonstrates, details of virtual geometry and materials, as well as the propagation of light and sound through virtual scenes can be precisely simulated and convincingly conveyed thanks to powerful rendering algorithms, tracking systems, and modern Head-Mounted Displays (HMDs). So, are the VR systems we can buy in electronic stores today the realization of the ultimate display?

A look at the third human sense reveals that this is not the case. Instead, with regard to the ability to touch and feel virtual content inside simulated worlds, VR is still found to be far from fulfilling Sutherland's vision. Rather, the haptic feedback provided by today's VR technology is crucially limited [Wang et al., 2019b]. Users cannot reach out and feel a virtual wall, neither can they perceive properties of virtual objects such as their weight or material [Butz and Krüger, 2017; p. 232]. In most commercial VR systems, the rich and varied haptic impressions we are used to from reality are replaced by simple vibration conveyed through handheld controllers, such as those carried by the user in Figure 1.1. This substitution lacks realism and is often not convincing. As a result, users are deprived of their ability to perceive things with the breadth of haptic channels usually available to them and are constantly reminded that the virtual worlds presented to them are actually not real – defeating the whole purpose of VR.

It is this aspect – the haptic fidelity of virtual worlds – that bears great potential to unlock unparalleled levels of immersion [National Research Council, 1995; p. 162] and User Experience (UX), and that we aim to advance with this thesis.

1.2 Experiencing Haptics in Virtual Reality

In the real world, haptic information that tells us about the composition and qualities of objects, is omnipresent [MacLean, 2000]. Our sense of touch helps us to gain information and to interact in effective, efficient, and safe ways. Tactile cues tell us about the world around us through receptors in our skin and

³*The Climb* - Press Kit by Crytek. <https://bit.ly/3i7QF8h>

kinesthetic sensations empower us to perceive movements of our own limbs and external forces [Lederman and Klatzky, 2009; Proske and Gandevia, 2012]. Being so used to haptic cues from the real world, we consequently expect haptic sensations to also accompany the simulated worlds we immerse ourselves in with VR technology.

1.2.1 The Role of Haptic Feedback

Consider the example of buying a new pan in a local kitchen store. To find the most suitable pan, we inspect available models visually, but also grasp them at their handle to feel the quality of the material and to test how ergonomically it is shaped. Most likely, we would also wield the pan around to assess its weight and balance, evaluating how well it fits our needs. Without consciously thinking about it, this short multisensory exploration of the pans provides us with a vast amount of information and impressions, all of which are fused into our perception of the pan as a whole – ultimately leading to the decision of buying it or continuing to search for a more suitable model.

Transferring this example scenario to VR, one can easily imagine a virtual kitchen store, similar to our example illustrated on the right in Figure 1.2, that allows users to buy pans via their personal VR equipment from the comfort of their homes. However, even if implemented with modern commercial VR equipment, users would not be able to experience the products the way they could experience them in a real store. While modern graphics could produce photo-realistic visual renderings of the pans, today's VR systems would fail to provide users with realistic impressions of the material, shape, weight, and balance of the different pans. Instead, the virtual pans would all feel the same. Considering, for example, their mass properties, the perceived weight of the pans would equal that of the controller, rendering them unrealistically lightweight. Apart from abstract vibration patterns, users could hardly feel any differences among the pans and information critical to their purchase decision would be missing or misleading.

The kitchen example makes apparent the pivotal role of haptics in the perception of our environment – be it real or virtual – and the fundamental lack of support for realistic haptic feedback in today's commercial VR systems. This lack, however, is not the result of a lack in research on VR haptics, but rather the consequence of the complexity and limitations inherent to most previously proposed solutions for haptic feedback in VR.

1.2.2 Approaches to Haptic Feedback in Virtual Reality

Research on VR haptics started in the 1960s [Brooks et al., 1990] and solutions today are classified along a conceptual continuum spanning from Active Haptic Feedback (AHF) to Passive Haptic Feedback (PHF) sketched in Figure 1.3.

Having been strongly connected to the field of robotics, early research on VR

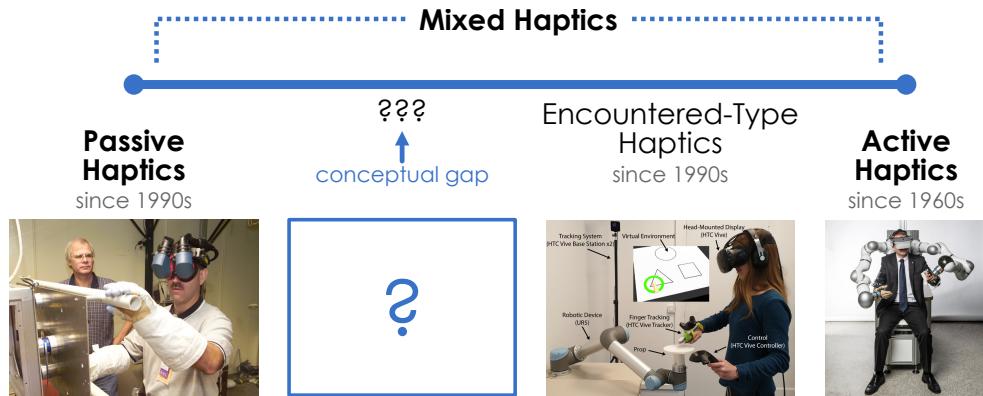


Figure 1.3: The Active-Passive Haptics Continuum. Left: An astronaut training with PHF provided by physical proxies (Credit: NASA; Source: Wikimedia Commons⁴). Center: A user interacting with a proxy presented by an encountered-type haptic display (extracted from [Mercado et al., 2020]; © 2020 IEEE). Right: The active haptic interface HUG (Credit & Source: DLR⁵). Most previous research focused on the active and passive poles, and on the active half of the continuum – leaving the passive half a conceptual gap.

haptics focused on solutions close to the active end of this continuum. Solutions based on Active Haptic Feedback (AHF) track the user's movements, simulate the user's interactions (e.g., collisions) with virtual objects inside the IVE, and render resulting output forces through actuators onto the user's body [Srinivasan and Basdogan, 1997]. This approach to haptics in VR can provide flexible feedback, but cannot easily convey multiple haptic modalities (e.g., forces, textures, temperatures) at once. Moreover, AHF is characterized by a heavy use of computer-controlled actuation (e.g., robotic arms as shown on the right in Figure 1.3), high mechanical complexity, and high computational complexity due to the involved physical simulations. Hence, most active approaches are expensive and only offer severely limited workspaces. As a consequence, active solutions to VR haptics found their way into highly specialized professional domains like teleoperation and research [Burdea, 1999b], but – apart from basic vibration feedback – failed to reach the VR consumer domain.

It was in 1994 that an alternative approach to haptic feedback for VR was introduced. Hinckley et al. [1994] proposed the use of “*passive interface props*” to realize a novel User Interface (UI) for neurosurgical planning. The idea was picked up for the domain of immersive VR and researchers continued developing and studying the concept of Passive Haptic Feedback (PHF) [Lindeman, 1999], which, in contrast to AHF, does not require any actuators nor complex mechanics or computations. PHF takes advantage of passive physical props, so called *haptic proxies*, that embody virtual objects – instead of using active machines that only

⁴Image by NASA from Wikimedia Commons showing astronaut John M. Grunsfeld using VR hardware. Unchanged. License: CC BY 2.0. <https://w.wiki/48VE>

⁵Image of the HUG by the DLR. Cropped and scaled. License: CC BY 3.0. <https://bit.ly/2Y2uKsg>

simulate their presence. These proxy objects serve as tangible counterparts that users can grasp and interact with naturally using their hands. Passive haptic proxies can deliver rich, multimodal haptic impressions of shape, material, temperature, weight, and other properties that enhance the sense of presence and perceived realism inside the IVE [Hoffman, 1998; Insko, 2001]. Even low-fidelity props were found to yield compelling illusions allowing for high-quality haptic feedback compatible with minimal budgets [Insko, 2001; Simeone et al., 2015].

Despite these many advantages accompanying the use of passive haptic proxies, however, PHF also did not establish as the default haptic feedback technique in consumer VR systems. Since PHF exclusively leverages *passive* proxies, the technique suffers from an inherent inflexibility and a lack of generality. Each change inside the IVE needs to be reflected manually in the physical environment and props need to be exchanged or modified. Moreover, as IVEs become larger, also the number of employed proxy objects needs to be increased as a single proxy can only serve as a counterpart for objects that have similar haptic properties and are located in the same place. All these factors rendered PHF a fit for professional domains like specialist training (e.g., of astronauts as shown on the left in Figure 1.3) or out-of-home VR entertainment, but prevented a wide adoption of the technique in other settings where more flexibility is required, such as in domestic VR setups. Yet, the concept of haptic proxies promises great potential also for these domains if only the problem of its inflexibility could be solved – leading us to the motivation of this thesis.

1.3 Motivation and Problem Statement

In his seminal paper, Sutherland [1965] described the “*ultimate display*”, i.e., the most complete implementation of VR, as “[...] 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 [...].”

VR systems that employ Passive Haptic Feedback (PHF) come close to this vision of the ultimate display with proxies providing realistic, multisensory haptic feedback. As the alternative approach of Active Haptic Feedback (AHF) suffers from its high complexity while struggling to provide holistic impressions of virtual objects, we stick to the concept of proxy-based haptics when pushing the edge of research towards the vision of the ultimate display.

With the naïve use of passive proxies, however, much of the *control* aspect described by Sutherland is left unrealized. As a consequence, the central motivation of this thesis is provided by the observation that:

Central Motivation

On the way towards an ultimate display, techniques are required that grant VR systems control over the proxy-based feedback and how it is perceived.

As the user's perception is a multisensory process that combines vision and haptics, we argue that such control can be achieved [Nilsson et al., 2021a]:

1. through *physical techniques* operating in the real environment (e.g., by leveraging a combination of proxies and actuation)
2. through *virtual techniques* that operate inside the IVE (e.g., by taking advantage of visual-haptic illusions and other perceptual phenomena)

Moreover, to reach a feedback quality as high as that of an ultimate display, any application employing proxy-based haptics and any technique enhancing the concept needs to solve two fundamental challenges, as we will review in chapter 2 [Lohse et al., 2019; Strandholt et al., 2020; Nilsson et al., 2021a,b]:

1. The Challenge of *Similarity*, which dictates that for the user, a proxy should *feel like* the virtual object it represents.
2. The Challenge of *Colocation*, which dictates that for the user, a proxy should be *spatially aligned* with the virtual object it represents.

Guided by these two central challenges and the two approaches to increase the *control* over proxy-based haptic feedback for VR, we address the following two overarching problems in this thesis:

Applicability

How can proxy-based VR haptics be applied?

While proxy-based haptic feedback for VR has been applied in several use cases already (as reviewed in chapter 2), many application domains are still unexplored. The expectations of users concerning the haptic feedback differs between domains, as does the design freedom granted to the developers of IVEs and haptic feedback. A demonstration of the successful utilization of proxy-based haptics in a novel application domain would further emphasize the concept's applicability and its potential to lead towards an ultimate display.

Improvement

How can proxy-based VR haptics be improved?

The naïve use of passive proxy objects for haptic feedback in VR is severely limited in terms of its flexibility. Once the proxy setup is in place, the VR system has no control over the user's haptic perceptions at runtime. This inflexibility hinders the adoption and utility of proxy-based haptics. The development of novel techniques, and the investigation and improvement of existing approaches that grant the VR system control over the user's haptic perception of proxies are thus central on the path towards an ultimate display. Specifically, techniques enhancing the reusability of props and the flexibility of the feedback provided by them appear promising.

1.4 Research Questions

To organize our research efforts, we break down the two problems above into four research questions – one guiding our research on the **Applicability**, and three leading our attempts regarding the **Improvement** of proxy-based VR haptics. This structure is based on the research approach proposed previously in:

Zenner, A. (2020). Enhancing Proxy-Based Haptics in Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, VRW'20, pages 549–550. IEEE. © 2020 IEEE. Final published version available in the IEEE Xplore® Digital Library. DOI: [10.1109/VRW50115.2020.00126](https://doi.org/10.1109/VRW50115.2020.00126)

Novel Application (Topic of Part II) Haptic proxies have been applied in several VR application domains as we will review in chapter 2. Yet, for many domains, it remains unclear if and how proxy-based haptics can be of support. In this context, we identified a novel application domain for proxy-based VR haptics: the domain of business process model exploration. Since abstract data, such as a graph-based description of a business process, is usually perceived only with the visual sense, users have no preconception of how interactions with this data should feel like. This renders the domain of process model exploration an exciting opportunity for proxy-based haptics, offering great haptic design freedom. By developing and evaluating a novel VR system for this domain, we investigate the **Applicability** of proxy-based haptics in the context of the first research question:

Research Question 1	Applicability
How can proxy-based haptic feedback support the domain of process model exploration?	

Physical Approach (Topic of Part III) Previous research on physical techniques that improve the flexibility of proxy-based haptics concentrated on a hybrid concept known as Encountered-Type Haptic Feedback (ETHF) [McNeely, 1993; Tachi et al., 1994], which combines ideas of PHF and AHF. As we will review in chapter 2, ETHF leverages robotic actuation (e.g., robotic arms or drones) to relocate proxy elements so that they can be haptically encountered by the user in a just-in-time manner. ETHF systems, however, suffer from similar limitations as conventional AHF as they also heavily rely on computer-controlled actuation and active force rendering to convey specific object properties (e.g., compliance or weight). Hence, ETHF is located within the active half of the haptics continuum – leaving the passive half a conceptual gap as indicated in Figure 1.3. Although individual approaches of previous research can be located in the passive half of the continuum, an overarching concept for this part of the continuum is missing. In addition, research is lacking techniques that can deliver varying kinesthetic sensations (e.g., different impressions of weight) through haptic proxies without

the need for active force feedback. It is these gaps that we fill with our first **Improvement** of proxy-based haptics in Part III. Specifically, we introduce the concept of Dynamic Passive Haptic Feedback (DPHF) and two DPHF-proxies when addressing the second research question:

Research Question 2	Improvement
How can the gap in the haptics continuum be filled with a concept that enhances the flexibility of proxy-based haptic feedback and enables improved kinesthetic perceptions in VR with only minimal actuation?	

Virtual Approach (Topic of Part IV) Besides physical techniques to enhance proxy-based haptics, previous research also considered purely software-based approaches inspired by the dominant impact of vision on our perception [Gibson, 1933]. Consequently, a second field of research was established, focusing on how visual-haptic illusions can empower VR systems to control the user's perception of haptic proxies. As our review in chapter 2 will show, particularities of human perception, such as visual dominance, change blindness, and visual suppression, offer exciting opportunities to "*play with senses*" in VR [Lécuyer, 2017]. The flexibility of proxy-based systems can be enhanced, for example, with techniques that only manipulate the user's visual impression of the IVE, such as pseudo-haptics, Redirected Walking (RDW), change blindness remapping, redirected touching, or haptic retargeting. Several of these approaches rely on Body Warping-based Hand Redirection (HR) – a fundamental technique that allows the VR system to control the user's real hand movement. In contrast to established domains like RDW, however, the domain of HR is still a comparably new research area. Even though HR represents the foundation of several virtual techniques that can enhance proxy-based haptics, crucial aspects of HR are still understudied. It has been found that HR can go unnoticed, but for common scenarios, the degree to which the technique can be applied without users noticing it is unknown. Moreover, while the related concept of RDW has successfully made use of change blindness and visual suppression, approaches to HR that take advantage of such perceptual phenomena are mostly unexplored. Finally, accessible and unified software frameworks that lower the barriers for researchers and developers to work with HR are missing. We address these points with our second **Improvement** of proxy-based haptics by studying the virtual approach of HR in Part IV. Specifically, we investigate the detectability of HR, introduce a novel HR technique that takes advantage of human eye blinks, and propose an open-source software toolkit for HR as we attend to the third research question:

Research Question 3	Improvement
What limitations and potentials does human perception imply for the technique of Body Warping-based Hand Redirection (HR)?	

Combined Approach (Topic of Part V) Research on the combination of physical and virtual techniques that enhance proxy-based haptics is scarce. Although many concepts are theoretically compatible, most combinations have not ever been scientifically investigated. With our third **Improvement** for proxy-based haptic feedback in Part V, we push the scientific knowledge forward and investigate a combination of the concept of DPHF proposed in Part III with the concept of HR studied in Part IV. Our results provide further grounds for the novel area of research on hybrid techniques, which, eventually, might lead VR closer to the realization of an ultimate display. Specifically, we consider both the challenges of *Similarity* and *Colocation* when we explore the fourth research question:

Research Question 4	Improvement
How can the physical approach of DPHF and the virtual approach of HR be combined to improve the flexibility of proxy-based haptic feedback?	

1.5 Methods and Approach

To answer these research questions, we leverage methods from the fields of Human-Computer Interaction (HCI), VR, and psychophysics. Throughout this thesis, we conduct several individual research projects, each following an approach consisting of three central phases:

1. **Concept Development:** concepts for novel VR systems, devices, and/or algorithms are developed based on new ideas, identified requirements, and findings of related work.
2. **Implementation:** high-fidelity prototypes of these conceptualized VR systems, devices, and/or algorithms are realized, encompassing the creation of novel hard- and/or software.
3. **Evaluation:** user studies and experiments are conducted with the developed prototypes to evaluate them and their underlying concepts, or to gain psychophysical insights.

Specifically, we develop one VR system that takes advantage of proxy-based haptic feedback in Part II to answer Research Question 1 (**RQ 1**), targeting the domain of immersive data exploration. To address **RQ 2** in Part III, we further propose the novel theoretical concept of Dynamic Passive Haptic Feedback (DPHF) and showcase its potential by developing two novel DPHF-proxies each in the form factor of a handheld VR controller. Similar to the development of the system in Part II, also the development of these novel haptic VR controllers covers an extensive technology stack. Implementation efforts span from hardware development (e.g., mechanics, 3D-printing, electrical circuits), over low-level software engineering (e.g., micro-controller programming), to high-level software development (e.g., programming with 3D-engines, integration of VR interfaces, development of

interactive VR applications and software for conducting experiments, collecting data, and data analysis). In Part IV, **RQ 3** is addressed by developing novel HR algorithms (e.g., utilizing eye tracking), psychophysical experiments, and two open-source software frameworks. To answer **RQ 4** in Part V, we ultimately build on the physical prototypes developed in Part III and the results and software of our perceptual studies conducted in Part IV.

The majority of our evaluations is realized as controlled user studies in the laboratory, complemented by a remote expert user study in Part IV and two thought experiments in Part V. We gather both qualitative and quantitative data, and analyze it using established methods (e.g., thematic analysis [Braun and Clarke, 2012] or statistical testing [Dix et al., 2003; pp. 332 ff.]). Our methodological approach is described in chapter 2.

1.6 Contributions

The goal of this thesis is to showcase the applicability and versatility of proxy-based haptic feedback for VR (**RQ 1**), and to improve it (**RQ 2**, **RQ 3**, and **RQ 4**). To achieve this goal, we first apply proxy-based haptics to a novel and promising application domain before we advance the field by proposing novel and investigating existing approaches that increase the flexibility of proxy-based haptics. In this way, the research in this thesis advances the fields of VR and HCI, haptics, and psychophysics. The principal contributions of this thesis can be summarized as follows:

Theoretical Contributions This thesis contributes in manifold ways to the theories of haptic feedback and VR.

Firstly, the following chapter contributes an in-depth literature review of VR haptics. The sections about physical, virtual, and combined techniques for improving proxy-based haptics constitute, to the best of our knowledge, the most comprehensive overview of this field so far. Following our literature review, in Part II, we then apply proxy-based haptics for the first time to the domain of process model exploration. By this, we contribute to the theoretical understanding of the potential that proxy-based haptic feedback bears for this and related domains – underlining the versatility of haptic proxies.

Further, with our work in Part III, we fundamentally contribute to the domain of VR haptics by complementing the Active-Passive Haptics continuum, and by proposing novel concepts for kinesthetic haptic feedback in VR. We introduce the theoretical concept of Dynamic Passive Haptic Feedback (DPHF), which represents a more passive combination of PHF and AHF than ETHF and primarily tackles the challenge of proxy *Similarity*. Based on DPHF, we then propose two novel approaches for conveying kinesthetic haptic impressions in VR: (1) a concept based on varying a proxy's inertial configuration (through varying its

weight distribution), and (2) a concept based on varying a proxy's air resistance (through varying its surface area). The first of those concepts has already given rise to several follow-up works by other researchers as our literature review in chapter 2 indicates. Furthermore, the second concept was recently patented in Germany [Zenner and Krüger, 2021].

In Part IV then, we make theoretical contributions to the field of visual-haptic illusions when proposing a novel algorithm for HR that takes advantage of human eye blinks and resulting change blindness. By this, we point out a new way to take advantage of the limitations of human perception for hand-based interactions in VR that establish proxy *Colocation*. Moreover, we contribute to the field of psychophysics by studying the degree to which different HR algorithms can be applied without users noticing the involved visual-proprioceptive manipulation.

Finally, the combination of DPHF and HR proposed and investigated in Part V contributes a theoretical understanding of how hybrid strategies can further enhance the flexibility and rendering quality of proxy-based haptics. Specifically, our results indicate the combination of our weight-shifting proxy concept (contribution of Part III) with unnoticeable HR (contribution of Part IV) to be superior in solving the *Similarity* and *Colocation* challenges compared to the individual techniques when rendering virtual weight shift.

Technical Contributions Beside advancing theories, this thesis also makes several contributions of technical nature.

The contribution of Part II encompasses a VR system that enables users to interactively familiarize themselves with business processes through immersive VR experiences supported by proxy-based haptic feedback.

Part III then contributes two haptic VR controller prototypes that implement the theoretical concept of DPHF. Both are specialized on rendering kinesthetic impressions of virtual objects and interactions while utilizing only minimal actuation. The first device, *Shifty*, is a weight-shifting controller that can convey impressions of different virtual weights and shapes following the first theoretical concept of adaptive inertia. The second device, *Drag:on*, is a shape-changing controller that can convey impressions of different virtual resistances, scales, materials, and fill states, following both the first concept of inertial adjustments and the second concept of adaptive air resistance. In both cases, our technical contributions are supplemented by open-source repositories featuring hardware building plans and software for recreating the devices.

Furthermore, the technical contributions of Part III are complemented by those of Part IV, which encompass open-source software toolkits for hand redirection and psychophysical experiments. Our *Hand Redirection Toolkit (HaRT)* features reference implementations of hand redirection techniques previously published by ourselves and other authors, as well as our proposed HR algorithm that leverages blinks and eye tracking. In addition, we contribute an open-source toolkit that supports the implementation of psychophysical experiments based

on the staircase procedure [Kingdom and Prins, 2016c; p. 53].

Finally, Part V contributes an additional VR system suitable for the empirical validation of our theoretical considerations about the combination of DPHF and HR. The system combines *Shifty* with unnoticeable HR to render different weight distributions inside a virtual stick.

Design Contributions Thirdly, this thesis also makes contributions related to the design of VR systems, devices, and applications.

Part II shows how IVEs, haptic proxies, and interactions can be designed that convey abstract graph-based data in VR.

Part III proposes designs for novel handheld VR controllers that can convey different kinesthetic impressions. These designs are realizable with low-cost materials, simple mechanics, low-power actuators, and low-complexity algorithms, and are thus suitable for integration in consumer VR in- and output devices. Moreover, we demonstrate how virtual objects and interactions can be visualized so as to give rise to convincing visual-haptic illusions in conjunction with the haptic feedback provided by our proposed controllers.

Furthermore, in Part IV, we reveal how much a user's hand can be redirected in VR without noticing it. These results can inform the design of IVEs and interactions that covertly employ visual-haptic illusions in VR.

Finally, we demonstrate in Part V how immersive interactions can be designed that not only involve a physical or a virtual real-time technique to enhance proxy-based haptics, but combine both strategies in order to render kinesthetic impressions of weight shift and to compensate for misaligned proxies.

1.7 Thesis Outline

This thesis is split into six parts as illustrated in Figure 1.4. We conclude the first part in the following chapter 2 with an in-depth literature review. We will summarize the historical roots of VR and its definition, introduce the concepts of immersion and presence, followed by a review of basic aspects of human perception that are of importance to the later parts of this thesis. After our revision of these foundations, we continue with a comparison of the two main approaches to haptics in VR (AHF and PHF) before we highlight the two central challenges for proxy-based haptics: the challenges of *Similarity* and *Colocation*. Once these challenges are introduced, we continue with a comprehensive review of previous approaches to solve them, following our taxonomy in [Nilsson et al., 2021a,b]. We conclude chapter 2 with a summary of our research methods.

As indicated in Figure 1.4, the following four parts of this thesis are each dedicated to one of the four research questions outlined above. Part II presents our work on **RQ 1**. Here, in chapter 3, we investigate the **Applicability** of proxy-based haptic

feedback in a domain with much design freedom and present our concept of *Immersive Process Model Exploration* alongside a novel VR system that implements it. We further report on a user evaluation in which we evaluated the effects of our proposed system on UX and on the understandability of process models, while comparing it to a traditional 2D interface on a tablet device.

Next, in Part III, our efforts towards an **Improvement** of proxy-based haptics are reported – starting with our work on physical techniques and **RQ 2**. In chapter 4, we propose the theoretical concept of *Dynamic Passive Haptic Feedback (DPHF)* before we introduce two implementations of this concept with *Shifty* in chapter 5 and *Drag:on* in chapter 6. For both devices, we first introduce their underlying, novel feedback concepts, describe their implementation, and report on user studies evaluating the feedback provided by them.

In Part IV, we then turn towards the investigation of the virtual technique of Body Warping-based Hand Redirection (HR) and **RQ 3**. We first study the detectability of HR in chapter 8 with a psychophysical experiment that reveals the extent to which *Continuous Hand Redirection* can be applied without being noticed by users. Following this investigation, in chapter 9, we introduce *Blink-Suppressed Hand Redirection (BSHR)*, a novel HR algorithm that takes advantage of human eye blinks to hide manipulations of the virtual hand when redirecting a user’s reach trajectory. A second psychophysical study here reveals both the Detection Thresholds (DTs) of a common state-of-the-art HR technique proposed by Cheng et al. [2017b] and our proposed technique – complementing our results of chapter 8. Finally, in chapter 10, we propose the *Hand Redirection Toolkit (HaRT)*, an open-source software framework targeted at researchers and developers that provides reference implementations of several hand redirection techniques. The toolkit also incorporates the results of chapter 8 and chapter 9.

Our last research question **RQ 4** is then addressed in Part V. Here, in chapter 12, we investigate how a combination of our physical approach of DPHF proposed in Part III and HR applied below the DTs found in Part IV can better solve the challenges of *Similarity* and *Colocation* than the individual techniques. Our investigation is concerned with a proof-of-concept scenario that focuses on perceiving the weight distribution of a virtual stick. The results of our experiments verify our theoretical predictions in that they showcase the combined technique to outperform the individual techniques with regard to both tackling proxy *Similarity* and *Colocation*.

We conclude this thesis in Part VI with a summary of our work and a review of our contributions in chapter 13, followed by ideas for future work in chapter 14.

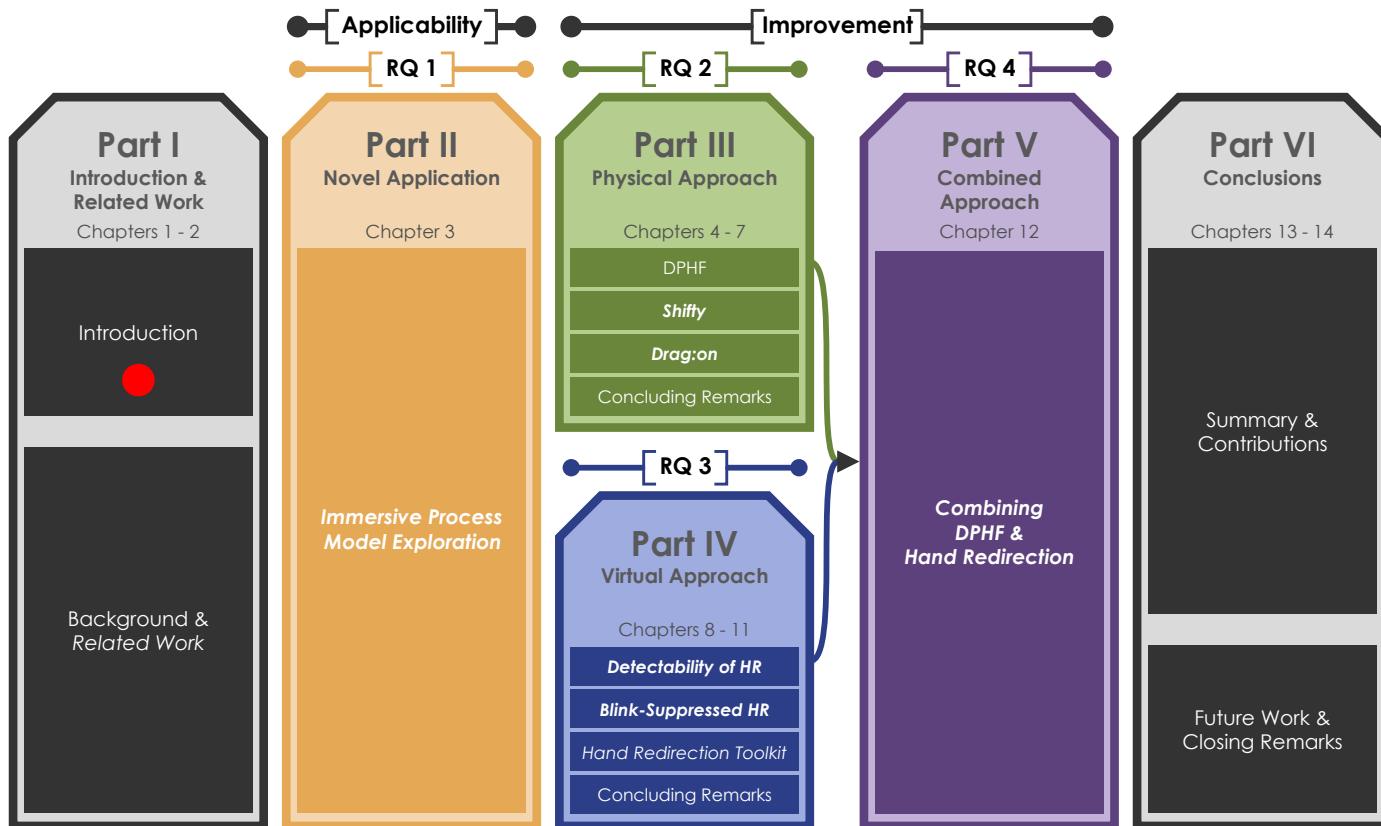


Figure 1.4: Thesis Outline. Contents previously published are marked in *italic*. Projects published as full papers in journals or conferences are highlighted in **bold**. Each box within a part represents a chapter. Red dot = "You are here!". The color-coding of the research questions is used throughout the thesis.

Chapter 2

Background and Related Work

This thesis grounds itself on scientific and technological achievements of the last decades. With the fields of VR and haptics research being inherently multidisciplinary, this chapter establishes the background for this thesis. We discuss related work that will cover topics from philosophical considerations, over computer science, neuroscience, psychology, physics, mechanical engineering, and the broad field of Human-Computer Interaction (HCI). Our survey starts with a reflection on the concept of Virtual Reality (VR) and continues with a brief review of how humans perceive a reality – highlighting the role of the sense of touch. Following this, the main part of this chapter provides a review of various approaches to realizing haptic feedback for VR – first, introducing active and passive haptics, followed by an in-depth review of the challenges of proxy-based haptics and corresponding solutions proposed in the past. The chapter concludes with an overview of the methodologies employed in this thesis.

2.1 Virtual Reality

We start by revisiting the historical roots and the definition of VR. To lay the grounds for this thesis, we also introduce immersion and presence, and end our background section on VR with a brief review of common application areas.

2.1.1 Historical Roots of Virtual Reality

The precursors of what we know as Virtual Reality (VR) technology today date back many decades – long before computers were invented. With head-worn displays being one of the most iconic hallmarks of VR, the roots of VR trace well



Figure 2.1: Left: A stereoscope from the early 1930's (Source: Wikimedia Commons; public domain⁷). Right: A modern cardboard VR device. A smartphone can be inserted to display a stereoscopic view of rendered virtual scenes.

back to the invention of the stereoscope at the beginning of the 19th century [Jerald, 2015; pp. 15–18]. These early viewing devices allowed to view Three-Dimensional (3D) images of a scene using pairs of separate pictures for the left and right eye. The development of stereoscopic viewing devices continued throughout the 19th and 20th century and evolved from static images to 3D film. On the path to even more realistic illusions, inventions of that time already started to incorporate additional senses. One of the most famous examples of such a multisensory immersive apparatus is the Sensorama by Heilig [1962]. Users of the Sensorama could sit inside a machine and view a stereoscopic 3D film that was accompanied by synchronized sound and haptic feedback (vibrations, tilting of the user's seat, wind), and smell⁶. These early concepts on the path to modern VR still reverberate in today's research communities and the VR industry. Figure 2.1 illustrates the similarities between a viewing device from the times of stereoscopes and a modern cardboard VR Head-Mounted Display (HMD).

A major leap towards modern VR technology is marked by the work of Sutherland [1968]. With the ceiling-mounted system shown on the left in Figure 2.2, known as the Sword of Damocles [Jerald, 2015; p. 22], Sutherland was one of the first to present an HMD apparatus that could track the user's head and display stereoscopic images rendered by a computer to the user. With this setup it was possible, for example, to render the wire-frame of a virtual room surrounding the user [Sutherland, 1968] while letting the user freely look around in it. Following these achievements, in the 1980s, the first commercial HMDs became available. The technical specifications of these devices and the necessary investments, however, were still far from the mass-compatible consumer devices available today. The HMDs back then were mainly used in certain well-funded research institutions (e.g., at specific universities, in laboratories associated with governmental institutions or the military, or specific companies). In the 1990s, the VR industry grew further and much basic research on VR was conducted. This development

⁶Sensorama on Wikipedia. <https://w.wiki/3xFt>

⁷Image of a 1930's stereoscope from Wikimedia Commons. Public Domain. <https://w.wiki/3xFu>



Figure 2.2: Left: The *Sword of Damocles* HMD of Sutherland (extracted from [Sutherland, 1968]; © 1968 ACM). Right: The HTC Vive – an example of a modern HMD.

was accompanied by the interest of the public and prominent appearances of VR technology in popular culture such as science-fiction literature and movies. At the same time, first-hand VR experiences became accessible through location-based entertainment, for example in theme parks and arcades.

At the end of the 1990s and the beginning of the 2000s, however, it became apparent that the state of the technology back then was not yet matured enough for VR to become an affordable everyday technology. As a result, the interest of the public in VR decreased and many VR companies closed their doors. Nonetheless, in the following ten years, which are known as the “VR winter” [Jerald, 2015; p. 27], basic research on VR continued. It was with the general advances of the computing field and the increased popularity of smartphones that technically more advanced and significantly cheaper VR hardware (especially HMDs) could be produced. The release of the first consumer-grade VR systems of this new generation in 2016, such as the Oculus Rift CV1⁸ and the HTC Vive⁹, marked the point where VR finally arrived on the mass market and the public’s interest in VR was reignited.

Still today, VR is a comparably young scientific field and as such develops rapidly with new devices, products, problems, and solutions arriving on a daily basis. As the humorous adage goes, it is difficult to make predictions, especially about the future¹⁰, and especially concerning the evolution of the technological landscape. However, as the VR user base at home and in professional contexts grows, chances seem good that VR will not vanish from our everyday lives anymore. Instead, it might now establish as a ubiquitous technology.

2.1.2 What is Virtual Reality?

The central concept of Virtual Reality (VR) is to allow users to experience simulated worlds that can be different from what we usually refer to as *the real world*,

⁸Oculus Rift CV1 on Wikipedia. https://w.wiki/_z5n2

⁹HTC Vive on Wikipedia. <https://w.wiki/3xG2>

¹⁰For a trace back to the origins of this adage see: <https://bit.ly/3gsTLTC>.

or *reality*. Trying to define Virtual Reality, however, is difficult as the concept can be approached from different perspectives (e.g., philosophically, technically, or from the standpoint of perception and presence [Steuer, 1992]).

Concept and Definition of Virtual Reality

The Merriam-Webster.com Dictionary frames Virtual Reality in a way that suits the perspective this thesis takes on VR very well. It defines *Virtual Reality* as “*an artificial environment which is experienced through sensory stimuli (such as sights and sounds) provided by a computer and in which one’s actions partially determine what happens in the environment*” and adds “*also : the technology used to create or access a virtual reality*”¹¹. Jerald frames VR similarly writing “*virtual reality is defined to be a computer-generated digital environment that can be experienced and interacted with as if that environment were real*” [Jerald, 2015; p. 9].

But what does *real* actually mean? When looking at it from a perceptual perspective, we perceive the world through our sensory organs by looking at it, listening to it, touching things in it, smelling, and tasting it. When experiencing the real world, our biological sensors are constantly transforming light that entered our eyes, sonic waves picked up by our ears, and physical forces acting on our body into signals sent to our brain. It is inside our brain that the reality we experience manifests as these signals are filtered, combined and interpreted. Philosophically speaking, we can thus only experience the *objective* reality that surrounds us in a *subjective* way as our perception of reality is formed in our minds [Jerald, 2015; pp. 59 ff.].

It is precisely at this point that the concept of VR picks up. By generating appropriate signals and providing them to our biological sensors, a VR system allows us to subjectively experience alternative realities – simulating a virtual reality to us, also referred to as the *Immersive Virtual Environment* (IVE). If the simulation is done *well*, our brain is able to interpret the artificial inputs as if they came from the reality we are used to interpret, constructing a coherent perception [Gonzalez-Franco and Lanier, 2017]. If done *perfectly*, the world we would perceive based on these synthetic inputs would appear as real to us as the objective reality. If a VR system stimulates our senses only *insufficiently*, however, users are likely to not feel *present* inside the simulated reality or might even suffer from sickness symptoms [Kennedy et al., 1993; Jerald, 2015; pp. 159 ff.].

Virtuality Continuum

The *Virtuality Continuum* by Milgram and Kishino [1994] depicted in Figure 2.3 allows to classify systems that simulate different realities. Specifically, systems are categorized along a continuum spanning from the *Real Environment* to the *Virtual Environment*. A system’s location within the continuum is determined

¹¹Merriam-Webster.com online dictionary entry for *virtual reality*. <https://bit.ly/3rLpLYV>



Figure 2.3: Virtuality Continuum (adapted from [Milgram and Kishino, 1994]).

based on the extent to which real-world and artificially created stimuli contribute to the user's perception of the experienced reality. This extent is captured in what is called the *Extent of World Knowledge*, i.e., the degree to which the world that is conveyed is modeled and known to the computer system.

While for experiencing the *Real Environment*, nothing needs to be computer-modeled, experiencing a fully *Virtual Environment* (the most extreme form of VR) requires the computer system to know everything about the world that is to be experienced in order to provide the corresponding stimuli. When stimuli from the real and virtual world are mixed, for example by blending virtual visual objects into the user's view of the real environment, realities in between the two extreme poles are created. These are classified as *Mixed Reality (MR)* and further categorized as either *Augmented Reality (AR)* or *Augmented Virtuality (AV)* based on the degree of virtual content involved, i.e., whether real-world or computer-generated stimuli dominate perception.

Milgram and Kishino originally focused primarily on the visual domain to categorize systems. This original view of the continuum can be extended to a 2D visual-haptic reality-virtuality continuum as presented in subsection 2.3.1 [Jeon and Choi, 2009]. For most of this thesis, however, we generalize Milgram and Kishino's classification by projecting all sensory dimensions onto the continuum in Figure 2.3 and classify systems taking into account also the origins of sounds, touch, smell, and taste [Borst and Volz, 2003].

2.1.3 Immersion and Presence

To make a virtual world become the subjective reality for a user, VR systems aim to make the user experience presence. Users feeling present in a simulated environment feel as if they "*are there (PI)*" and the events of the virtual world are "*really happening (Psi)*" to them [Slater, 2009]. These two illusions, known as the "*place illusion (PI)*" and "*plausibility illusion (Psi)*", lead to users behaving realistically inside virtual environments [Slater, 2009] – a key aspect for the success of many VR systems (e.g., for training, simulation, or exposure therapy).

In this context, two central terms exist in the field of VR research: *immersion* and *presence* [Slater and Wilbur, 1997]. While both terms are often taken as synonyms, they describe two distinct concepts in the scientific language as outlined below.

Immersion

Immersion describes an objective property of a VR system, a technical quantity, which is independent of the user [Slater and Wilbur, 1997; Jerald, 2015; pp. 45 f.]. The immersion of a VR system is defined by its technical setup; specifically, by how well it stimulates the user's senses from a technical point of view. Determining factors are for example the display resolution, latency, quality of pixel colors and headphones, range of sounds, colors, or vibrations that can be conveyed [Jerald, 2015; p. 45]. Moreover, also aspects such as how the IVE responds to user actions and input, how virtual characters inside an IVE behave, and how the user's virtual body is realized play an important role for immersion. In summary, research has established an understanding of several aspects that contribute to immersion [Steuer, 1992; Slater and Wilbur, 1997; Jerald, 2015; p. 45], such as:

1. The *inclusiveness* of a VR system specifying how well stimuli originating from the real environment are locked out and made unnoticeable.
2. The *extensiveness* of multimodal stimulation, i.e., how many different perceptual channels are being stimulated.
3. The *matching* of artificial stimuli to inputs like movements or interactions.
4. The *surroundedness* describing how complete, or "*panoramic*" [Jerald, 2015; p. 45], the stimuli are presented.
5. The *vividness* of the stimulation, i.e., the quality of the stimuli provided.
6. The *interactivity* of the VR system and application.
7. The *plot* deciding how convincing and coherent the virtual scenario, story, or experience of the artificial reality is conveyed to the user.

An alternative conception of immersion is brought forward by Slater [2009] who regards immersion as defined by the set of "*valid actions*" that can be carried out within the IVE. Slater [2009] argues that the set of valid actions encompasses all such actions that allow a user to either perceive (through "*valid sensorimotor actions*", such as looking around or touching something in the IVE), or change the virtual world (through "*valid effectual actions*", such as moving a virtual object).

Whichever of the two conceptions of immersion are considered, the work in this thesis aims to *increase* immersion. Our approach to achieve this is by introducing and studying techniques that improve central dimensions of VR systems, such as extensiveness, matching, surroundedness, vividness, and interactivity. Following Slater's view on immersion, the ideas brought forward in this thesis enlarge the set of "*valid sensorimotor actions*" that allow users to explore IVEs haptically.

Presence

In contrast to immersion, presence describes a "*state of consciousness*" [Slater and Wilbur, 1997] of the user that can ensue when the user is exposed to the

artificial stimuli of an IVE. The most prominent aspect of presence is a user's perception of "*being there*" in the IVE [Slater and Wilbur, 1997; Jerald, 2015; p. 46]. While immersion can be considered an objective, technical quantity of the VR system, presence is a function of the user and immersion [Steuer, 1992]. As such, presence can vary with time and across people¹². Ultimately, however, presence is facilitated by immersive systems and increased immersion can increase the probability of users experiencing presence – but immersion alone is not a guarantee for presence [Jerald, 2015; p. 46].

The International Society for Presence Research defines the concept of presence in an explication statement through 12 paragraphs. Presence is described therein as "*a psychological state or subjective perception in which even though part or all of an individual's current experience is generated by and/or filtered through human-made technology, part or all of the individual's perception fails to accurately acknowledge the role of the technology in the experience.*"¹². This central description highlights the importance of the involved technology, i.e., the VR system, to be itself unnoticeable to the user during VR experiences. By providing stimuli that closer match those from the real world, discrepancies between the artificial and natural stimulation are eliminated and the VR system as a consequence becomes increasingly transparent to the user.

Presence research distinguishes various central dimensions of presence. Jerald [2015; pp. 47 ff.] categorizes them into the following four "*core components*":

1. *The Illusion of Being in a Stable Spatial Place* describes the sense of "*being there*" inside the IVE, also referred to as the place illusion (PI) [Slater, 2009], and is flagged by Jerald [2015; p. 47] as the most important factor.
2. *The Illusion of Self-Embodiment* relates to the extent to which a virtual body is conveyed to the user in a compelling way so that the user accepts the virtual body as their own during the simulation [Slater, 2009].
3. *The Illusion of Physical Interaction* refers to the ability to touch, feel, and interact with virtual objects portrayed in the IVE in a way that conveys their physicality. This encompasses also the use of haptic feedback. Referring to the International Society for Presence Research¹², our work aims to strengthen this aspect by empowering the user to "*[perceive] that the objects, events, and/or people s/he encounters look, sound, smell, feel, etc. as they do or would in the physical world*".
4. *The Illusion of Social Communication* conveying the feeling of communication with entities such as virtual characters in the IVE.

As presence is a volatile state, specific events can cause the user to fall back in a state of not feeling present anymore, even while inside a virtual scene that the user felt present in just moments before. These can be events, for example, which

¹²*The Concept of Presence: Explication Statement* by the International Society for Presence Research [2000]. <https://bit.ly/3JG80FJ>

surface limitations in a VR system's inclusiveness. Such events might abruptly remind the user of the physical world, such as when feeling the tension of a cable, colliding with physical objects not incorporated in the IVE, hearing people in the real environment talking, or reaching through virtual objects without perceiving an appropriate haptic response. Also bugs or glitches that make apparent the limitations of the simulation can transport the user's state of consciousness back to reality. The consequence of such events is referred to as a "*break in presence*" (BIP) [Slater and Steed, 2000]. To maintain presence, VR systems should be designed so as to reduce the probability of BIPs to occur.

When studying the quality of a VR system, it is valuable to measure whether users experience presence. To do so, research has come up with a variety of methods to assess presence ranging from subjective measures (e.g., self-assessment questionnaires [Van Baren and IJsselsteijn, 2004]), physiological measures (e.g., recording the response of the user's body to (potentially stressful) events happening inside the IVE) [Insko, 2001; and references therein], behavioral measures (e.g., observing the extent to which users behave realistically in response to virtual events) [Usoh et al., 1999], to counting BIPs [Slater and Steed, 2000].

2.1.4 Technical Realization and Application Areas

With the theoretical definitions of VR, immersion, and presence in place, we complete our introduction to VR by briefly outlining how VR systems are technically realized and by summarizing in which domains VR is used.

Technical Realization of Virtual Reality Systems

The term *VR system* describes the technical setup that simulates the IVE and presents the user with appropriate stimuli. As such, VR systems consist of components that stimulate the senses through multimodal feedback.

Visual Feedback The simplest approach to immerse users visually in a 3D scene is *Desktop VR*, i.e., by rendering the virtual scene multiple times per second using computer graphics algorithms and displaying the rendered images monoscopically or stereoscopically (e.g., with shutter glasses) on a desktop monitor. If combined with head tracking, the rendered view can be adapted to the head position in front of the monitor – a concept known as *Fish Tank VR* [Ware et al., 1993] (shown on the left in Figure 2.4). *Cave Automatic Virtual Environments* (CAVEs) (shown in the center of Figure 2.4) are more immersive alternatives based on large-scale stereoscopic projection systems that provide a larger Field of View (FOV) [Kuhlen and Hentschel, 2014]. CAVEs track the user's head and project the corresponding perspective rendering of the IVE onto projection walls surrounding the user. Costs and space requirements, however, limit the application domains of CAVEs severely [Kuhlen and Hentschel, 2014].



Figure 2.4: Left: Early implementation of Fish Tank VR (extracted from [Ware et al., 1993]; © 1993 ACM). Center: A user in a CAVE system experiencing projection-based VR (Source: Wikimedia Commons; public domain¹³). Right: The author of this thesis wearing a wireless version of the HTC Vive Pro, a modern HMD-based VR system.

A better trade-off between costs and immersion is provided by spatially tracked *Head-Mounted Displays (HMDs)* (shown on the right in Figure 2.4). While not yet covering the full FOV of humans (200° horizontal and 135° vertical [Jerald, 2015; pp. 89 f., and references therein]), modern HMDs still enable surrounding and stereoscopic experiences. The *HTC Vive* and *HTC Vive Pro* headsets used in several of our experiments, for example, provide a FOV¹⁴ of 110°. To prevent visual-proprioceptive discrepancies that can result in sickness symptoms, HMDs employ high update rates (e.g., 90Hz in the *HTC Vive* and *HTC Vive Pro*¹⁴). While low-end HMDs like *Google Cardboard* (shown in Figure 2.1) only track head rotation with a plugged-in smartphone (Three Degrees of Freedom (3-DoF) tracking), more sophisticated VR systems use rotation and translation tracking (6-DoF tracking)¹⁵.

Auditory Feedback Most VR systems also immerse users aurally in virtual scenes by means of headphones, or, more rarely, speakers in the real environment. Similar to how the images displayed in the HMD adapt with high update rates and low latency to the tracking of the user, also the audio feedback provided to the left and right ears adapt to the user's location and orientation in the IVE based on spatial sound simulation¹⁶.

Haptic Feedback Haptic feedback, the central topic of this thesis, plays a crucial role in how interactions with the IVE are perceived. In contrast to visual and auditory feedback, however, the quality of the haptic feedback provided by today's commercial VR systems is still severely minimal. In VR systems that employ handheld VR controllers, the controllers serve both as input devices (reading built-in buttons, joysticks, and tracking) and, if supported at all, as

¹³Image of a user in a CAVE system from Wikimedia Commons. Public Domain. https://commons.wikimedia.org/wiki/File:User_in_a_CAVE_system.jpg

¹⁴Technical specifications of *HTC Vive* and *HTC Vive Pro*. <https://bit.ly/3Bhfu8X> and <https://bit.ly/3kuUhkZ>

¹⁵VR positional tracking on Wikipedia. <https://w.wiki/3yKn>

¹⁶VRWorks Audio by NVIDIA. <https://bit.ly/2WBf4va>

haptic output devices. Their haptic rendering capabilities are usually limited to controller “*rumble*” [Jerald, 2015; p. 306], i.e., simple vibration implemented by vibration actuators attached to the controller casing. Vibration patterns typically vary in vibration amplitude, frequency, and duration, and are acutely limited in terms of vividness and surroundedness [Wang et al., 2019b]. Moreover, vibration is frequently used as a substitute for other haptic perceptions like collision forces, tension, or virtual weight [Cooper et al., 2018]. With such an inflationary use of vibration, only abstract feedback mappings and limited stimulus matching is achieved. Moreover, some of our biological sensors that contribute to haptic perception are simply not addressed by vibration inside handheld controllers, as apparent from section 2.2 on Human Perception. Consequently, novel solutions for haptic feedback in VR bear great potential to boost immersion. Today’s VR controllers still fail to convey most of the haptic impressions we are used to from the real world – a circumstance constituting a central motivation for this thesis.

Olfactory and Gustatory Feedback Still the most exotic types of feedback are those of smell and taste. Olfactory feedback is realized by stimulating olfactory receptors in the nose [Barfield and Danas, 1996]. Sensations of smell can be achieved, for example, by mixing odors and transporting them to the user’s nose either through tubes or with an air cannon [Yanagida et al., 2004; and references therein]. Similarly, gustatory feedback aims to induce perceptions of taste, such as the sensation of sweetness. This can be achieved, for example, by relying on thermal actuation of the tongue [Karunananayaka et al., 2018; and references therein]. Approaches to olfactory and gustatory feedback in VR are still at an early stage of development and mainly found in specialized research laboratories.

Software VR systems usually employ 3D engines to define the look and behavior of simulated worlds. 3D engines are comprehensive software frameworks that handle the composition of virtual scenes, the execution of scripts that define the behavior of virtual objects therein, visual and auditory rendering, and the integration of tracking systems and haptic devices. The VR systems for this thesis were realized with the *Unity* 3D engine¹⁷ by Unity Technologies.

Application Areas of Virtual Reality

The unparalleled capabilities of inducing a sense of presence in simulated worlds render VR a human-computer interface that suits uniquely to specific application areas. Analyzing what VR has been successfully used for, we identified three central types of use cases enabled by VR. According to these types, VR can allow users to experience scenarios in an immersive way that in reality would either be:

1. *impossible*, e.g., because it would contradict the laws of time, physics, or other constraints

¹⁷Unity by Unity Technologies. <https://bit.ly/3ztuIHe>

2. *non-economic*, e.g., because it would be too expensive, too much effort, too time-consuming, or due to other constraints
3. *too risky*, e.g., because of financial, mental, or physical dangers involved

Offering a solution to these types of use cases, VR has found its way into a variety of application fields. Table 2.1 lists several examples of where VR is used nowadays. While far from complete, it showcases the diversity of the areas that can benefit from VR technology.

A prominent domain for VR is that of simulation and training. Here, VR draws from its capabilities to simulate real-world situations in a way that is more memorable, less abstract, more realistic, and often more entertaining than most conventional approaches, while still being comparably risk-free and economic compared to training in the real world. VR also enables novel ways of visualization and creative work. Especially the field of design and production can benefit from the possibilities of remote collaboration and fast iteration cycles that come with VR. Moreover, the field of healthcare explores the use of VR already for

<i>Field</i>	<i>Area</i>	<i>Example</i>	<i>Reference</i>	<i>Type</i>
Simulation/ Training	<i>specialist training</i>	firefighters	web link	
		astronauts	web link	
		offshore drilling	web link	non-economic
	<i>sports</i>	fencing	web link	too risky
		hockey	web link	
		sailing	web link	
Visualization/ Creative Work	<i>education</i>	museums	web link	impossible
	<i>data visualization</i>	immersive analysis	[Zielasko et al., 2016]	
			[Laha et al., 2012]	impossible
			[Sousa et al., 2017]	
	<i>retail</i>	cars	web link	
		airplanes	web link	
		furniture	[Zenner et al., 2020b]	
Healthcare	<i>architecture</i>	previewing spaces	web link	non-economic
	<i>film making</i>	set design	web link	
	<i>product design</i>	aviation industry	web link	
	<i>therapy</i>	exposure therapy	[Hoffman, 2004]	too risky
Communication	<i>diagnosis</i>	Parkinson diagnosis	[Orlosky et al., 2017]	non-economic
	<i>collaboration</i>	multiuser teaching virtual conferences	web link web link	non-economic
Entertainment	<i>gaming</i>	PC VR gaming	web link	impossible
	<i>location-based</i>	VR arcades	web link	non-economic
	<i>events</i>	film festivals	web link	too risky

Table 2.1: Several examples of areas that successfully employ VR. The rightmost column classifies the examples according to our three identified types of use cases.

decades. Sensors build inside VR hardware can be taken advantage of for (remote) diagnosis and the psychological consequences of experiencing controlled, immersive scenarios enable novel ways of therapy. Furthermore, multiuser VR bears great potential to revolutionize remote collaboration and communication. Finally, still the most influencing domain that drives the current growth of consumer VR is the field of entertainment where VR enables novel gaming and event experiences. The dark blue entry in Table 2.1 highlights the domain this thesis contributes to with the system presented in Part II. The contributions of Parts III, IV, and V are not bound to a specific application domain but are universally applicable and can benefit VR systems in arbitrary fields.

2.1.5 Conclusion

The preceding sections provide basic background information on VR which allows to assess the work presented in this thesis in the bigger picture. We reviewed the origins of VR and introduced immersion as a technical property of the VR system. Moreover, we reviewed that when the user's cognitive state allows for it, users can feel present in a simulated world, with presence being a function of the user's psychological state and the immersion of the VR system.

Based on this central relationship, the overarching aim of this thesis is to facilitate the sense of presence by increasing the immersion of VR systems. Specifically, we aim to enhance immersion both qualitatively and quantitatively through the introduction of novel haptic feedback concepts and visual-haptic illusion techniques that increase their effectiveness.

2.2 Human Perception

"The good news is that VR does not need to perfectly replicate reality. Instead we just need to present the most important stimuli well and the mind will fill in the gaps."

– Jerald [2015; p. 61]

Following the recommendations of Jerald [2015; pp. 55 ff.], this thesis develops novel concepts for VR with the workings of human perception in mind. Leveraging the particularities and limitations of our perceptual system allows us to make VR technology as "*transparent*" as possible. The following sections will introduce background information and related research about human perception that the different concepts, devices and techniques presented in this thesis are based on. Our review will make apparent that in order to achieve certain perceptions, it usually suffices to stimulate the sensory system in a way that is *sufficiently good* rather than a perfect replication of the real world. Since human perception is a vast topic area, our review focuses on aspects relevant for this thesis, such as visual suppression and change blindness, as well as on how different haptic

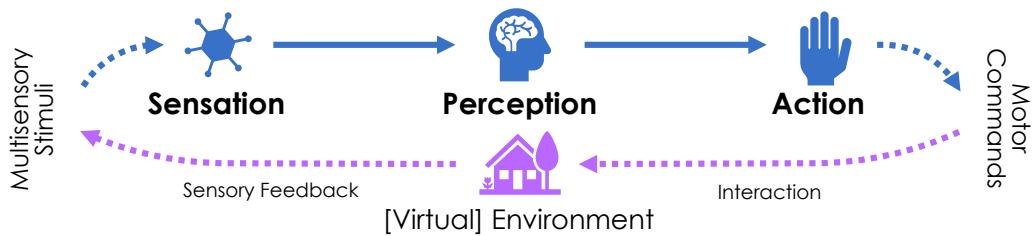


Figure 2.5: The *Perception-Action loop* (adapted from [Ernst and Bülthoff, 2004]).

qualities are perceived. At the end, we conclude our review with a summary of important findings in the domain of multisensory integration.

2.2.1 Basics of Human Perception

The perception of the world around us is a complex process that cannot be explained by a single model and still is only partially understood by science. In this process, one distinguishes between *sensations* and *perceptions* [Jerald, 2015; p. 72]. Sensations refer to the lower level of this process and describe the reactions of sensory organs to stimuli (e.g., physical or chemical) triggering them. Perception describes the higher level cognitive process that filters and interprets sensations.

In a simplified form, the process of perception is commonly represented as a three-staged *Perception-Action loop* [Ernst and Bülthoff, 2004] (depicted in Figure 2.5). In the first stage of this loop, the environment excites our sensory organs, which create sensations. Next, the sensations are interpreted by our brain to create perceptions. These perceptions, in turn, lead to a response, i.e., actions such as movements, carried out by our body's effectors, i.e., our muscles. These actions then affect the environment, which consequently leads to new sensations and a re-iteration of the loop. VR systems are directly coupled to the Perception-Action loop when simulating an environment. Their central task is to provide coherent sensations at the appropriate times as a response to user interactions with the IVE so as to keep the user re-iterating the perceptual loop.

Classically, five different senses are distinguished in human perception: seeing, hearing, touching, smelling, and tasting. Besides these, modern physiology distinguishes a set of further human senses such as the senses to perceive temperature, pain, balance (i.e., the vestibular sense), and the position and movement of our own body (i.e., the proprioceptive sense)¹⁸. Most relevant for this thesis are the domains of visual perception, and the proprioceptive and tactile senses (that both are attributed to the domain of haptics [Lederman and Klatzky, 2009]), as well as the topic of multisensory perception [Ernst and Bülthoff, 2004].

¹⁸German article about the senses ("Sinn (Wahrnehmung)") on Wikipedia. <https://w.wiki/3ykr>

2.2.2 Visual Perception

Visual perception, together with auditory perception, as well as the sense of smell, is considered a *distant-perceiving sense*¹⁸. The visual sense empowers us to see our environment from a distance and can be considered the most important perceptual channel for VR. When our eye lids are opened, light that is reflected from the surrounding environment enters our eyes through the cornea, is focused by the lens inside our eye and hits the photoreceptors located on the retina at the back of our eyes [Fairchild, 2005; pp. 2 ff.]. The signals generated when photons fall onto the photoreceptors are transferred through the optic nerve into our brain, which processes the image data, for example to detect patterns, movements, or to recognize objects in our vision.

One can further distinguish central and peripheral vision [Jerald, 2015; p. 88]. While central vision is optimized to perceive visual details with high visual acuity and in color, peripheral vision is not sensitive to color but enables vision in dark conditions [Fairchild, 2005; pp. 8 ff.]. Peripheral vision is not optimized to perceive fine details, but is sensitive to detect motion in the FOV [Jerald, 2015; p. 88]. The most accurate vision is achieved within foveal vision, which covers only 2° to 5° in our central vision [Marwecki et al., 2019; Fairchild, 2005; p. 5].

In HMD-based VR systems, visual stimuli from the real world are almost completely blocked out and substituted by artificial stimuli originating from the pixels of the built-in display. The color that is projected onto the user's retina effectively mimics the light reflected from the IVE. By constantly adapting the pixel colors on the screen in response to the user's head motions, the illusion of *looking at* the virtual scene can be conveyed. In addition to that, by providing a separate image for the right and left eye, each showing the IVE from slightly offset points of view resembling the separation between our eyes, users can view the IVE stereoscopically. Such binocular cues can provide an additional sense of depth [Jerald, 2015; pp. 120 f.]. Any perceivable latency in the visual display loop can tremendously deteriorate the quality of a VR experience and quickly lead to users feeling motion sick [Stauffert et al., 2021; Jerald, 2015; pp. 183 ff.].

The following subsections discuss the physiology of human eye blinks, the phenomenon of visual suppression, and change blindness – particularities of visual perception that the technique we propose in chapter 9 is taking advantage of to hide visual manipulations of the IVE.

Blink-Induced Visual Suppression

In order to protect our eyes and to lubricate the cornea¹⁹, humans typically blink 10 to 20 times per minute [Doughty, 2002; Leigh and Zee, 2015; p. 241]. During an eye blink, the eyelids are quickly moved down and lifted up again, occluding

¹⁹Information about *Eyelid Movements* by Dr. Janet Fitzakerley from the University of Minnesota Medical School Duluth [2015]. <https://bit.ly/3LuTmxA>

the pupil for 100ms – 150ms [Volkmann, 1986]. This closing of the eyelids hinders light from entering the eye and consequently reduces retinal illumination. While blinks effectively blind us for a short time, they usually go completely unnoticed. Potential reasons for this have been subject of research in the past, as have the general neural processes linked to blinking. An important aspect related to the unnoticeability of blinks is the phenomenon of visual suppression. Bristow et al. [2005] found indication of a neural suppression mechanism that is triggered during blinks. This mechanism affects specific parietal and prefrontal brain regions, and is potentially responsible for the reduction of visual input during blinks usually going unnoticed. Thus, during blinks, our vision is blinded by both the mechanical interruption as a result of the closing of the eyelids, accompanied by additional active visual suppression in the brain. *Blink-induced visual suppression* begins shortly before blink onset and lasts for 100ms – 200ms [Volkmann, 1986]. Similarly, human vision is also suppressed during fast ballistic eye movements known as saccades [Volkmann, 1986]. In contrast to eye blinks, however, saccades on average only last 50ms.

Blinks can be either voluntary (e.g., in the context of social interaction), spontaneous (i.e., periodically occurring without external stimuli¹⁹), a reflex (e.g., initiated by bright light or objects that approach rapidly¹⁹), or externally triggered [Crnovrsanin et al., 2014; VanderWerf et al., 2003]. In order to make use of eye blinks in the context of VR, it is important to consider the frequency of blinking. Previous research found the blink frequency to drop when using computer monitors [Patel et al., 1991]. However, recent results by Dennison et al. [2016] indicate that the use of HMDs increases blink frequencies compared to computer monitors. This motivates our research in chapter 9 on techniques that utilize blink-induced phenomena like visual suppression and change blindness (introduced in the following).

From a technical perspective, eye blinks can be tracked robustly with eye trackers built into modern HMDs such as the *HTC Vive Pro Eye*²⁰. Moreover, blinks last for a relatively long time span with regard to the update rates of interactive VR systems. A blink of 150ms, for example, spans 13 frames when considering a VR system that renders at 90Hz update rates (i.e., with a rendering time of 11ms per frame) – leaving enough time to manipulate the IVE.

Change Blindness

In addition to inducing visual suppression, blinks and saccades also give rise to a perceptual phenomenon known as *change blindness* [O'Regan et al., 2000]. Change blindness as defined by Simons and Levin [1997] is “*the inability to detect changes to an object or scene*”. It is a particularity of visual perception that has been observed since the 19th century and has been subject to research in psychology for

²⁰Technical specifications of HTC Vive Pro Eye with integrated eye tracking. <https://bit.ly/3zmWYvf>

decades²¹. Experimental results provide evidence that people are likely to miss changes, even dramatic ones, introduced to a visual image, when these changes occur during a saccade, when people blink, when a blank visual stimulus is provided for a short time, or when a visual masking pattern (also referred to as a “mudsplash”) is added to the visual image [Simons and Levin, 1997].

Simons and Levin [1997] provide a review of experimental results and corresponding conclusions related to change blindness. It is found that the apparent shortcoming of being susceptible to not noticing changes to a visual scene can be regarded rather a positive feature of human visual perception than a flaw. Scientific evidence suggests that when viewing a scene, our perception extracts and stores only what appears to us as the most important information to understand the visual scene we see. Additional details in the scene that do not affect what is of central importance to us are not retained to the same extent. As a result, our perceptual system can more easily create a stable perception of our environment since it is not distracted by details, which are irrelevant for our understanding of the scene²². Simons and Levin [1997] summarize that when repeatedly processing views of a scene, “*if the gist [of a scene] is the same, our perceptual system assumes the details are the same.*”. Even if we pay attention to a specific object in the visual image, we might fail to detect changes related to the object. This could be, for example, if the change affects a property of the object that is not central to our understanding of the object – or in other words, that is not considered a part of the “gist of the scene” by our perceptual system²³. Simons and Levin [1997] conclude that “*attention is necessary, but not sufficient, for change detection*”.

Viewing the phenomenon of change blindness from the perspective of lower-level perception, findings indicate that when a change in the scene does not effect a direct “*motion on our retina that attracts attention*” [Simons and Levin, 1997], it is likely to go unnoticed. Blinks, saccades, flashing blank stimuli or patterns all suddenly change our visual input, which makes the changes of visual input originating from changes in the scene less noticeable.

In 2010, Steinicke et al. investigated whether change blindness also occurs in stereoscopic viewing conditions as encountered in VR systems. Their experiments with stereoscopic projection systems revealed that change blindness effects occur in VR systems to the same extent as they do in monoscopic viewing conditions (e.g., 2D desktop monitors) [Steinicke et al., 2010a]. As a result of these findings and due to its unique characteristics, the phenomenon has seen much attention by VR research in recent years. Various techniques have been proposed that take advantage of change blindness to covertly introduce changes to the IVE. Especially the domains of redirected walking and haptic retargeting have

²¹Change blindness on Wikipedia. <https://w.wiki/3ynK>

²²The following video on YouTube by the Transport for London nicely demonstrates change blindness. <https://youtu.be/ubNF9QNEQLA>

²³The following video on YouTube by Daniel Simons documents a change blindness experiment outside the laboratory, in which participants fail to notice that the person they talk to is exchanged during the conversation. <https://youtu.be/FWSxSQsspiQ>

been actively exploring the use of change blindness, as we will outline in subsection 2.6.2. Inspired by the success of these approaches in unnoticeably adapting the scene to solve interaction problems, we propose the first technique to take advantage of blink-induced change blindness for hand redirection in chapter 9.

2.2.3 Haptic Perception

Haptic perception, in contrast to visual perception, is considered a *direct contact sense*¹⁸. The term *haptic* originates from the Greek term *haptikós*, which means “suitable for touch”²⁴. Haptic perception spans multiple individual sensory systems and covers skin-related sensations such as surface information, pressure, vibration, and stretch, as well as sensations related to the position and movement of our own body parts, sensations of pain, and temperature²⁵. As such, haptic perception both informs us about external stimuli, e.g., properties of objects we are in contact with, as well as the internal state of our body. The breadth of information provided by our haptic perception informs us about our environment and enables us to act therein in an effective, efficient, and safe way. While we are often not consciously aware of these sensations, we are used to them through life-long experience. In order to serve expectations from real-world experiences, VR systems should aim to complement visual and auditory experiences with haptic sensations when users explore virtual worlds. Important in this context is the classification of haptics into *tactile* (or *cutaneous*) and *kinesthetic* (or *proprioceptive*) perception [Srinivasan and Basdogan, 1997; Lederman and Klatzky, 2009].

Tactile Perception

The term *tactile* originates from the Latin word for “tangible”, *tactilis*²⁶. Tactile perception is also referred to as cutaneous perception, which stems from the Latin word for “skin”, *cutis*²⁷. Tactile perception is formed through the sensations signaled by the four central tactile receptors in our skin, which each are responsible for sensing different aspects of a touch [Lederman and Klatzky, 2009]:

- Meissner’s corpuscles are responsible for low frequency vibration detection.
- Merkel cells are responsible for the detection of pressure and very low frequency vibration, as well as for the perception of coarse texture.
- Ruffini endings enable the perception of skin stretch.
- Pacinian corpuscles are responsible for high frequency vibration detection and the perception of fine texture details.

²⁴Haptic perception on Wikipedia. <https://w.wiki/3zMv>

²⁵German article about haptic perception (“*Haptische Wahrnehmung*”) on Wikipedia. <https://w.wiki/3zMw>

²⁶Origin of the word *tactile*. <https://bit.ly/3Dt0hUh>

²⁷Origin of the word *cutaneous*. <https://bit.ly/38qRTpN>

In addition, Klatzky et al. [1987] also attribute thermal sensation to the domain of tactile perception, which is signaled through additional thermoreceptors that sense changes in skin temperature [Lederman and Klatzky, 2009]. These receptors operate within a temperature range of 5°C to 45°C with C fibers responding to warmth and A delta fibers responding to cold [Wang et al., 2020a].

The information signaled by all of these receptors encodes information that allows us to perceive fine surface details such as the shape, edges, or material properties of objects we touch. Moreover, to gather information about an object, we naturally perform so called exploratory procedures [Klatzky et al., 1987] – i.e., we actively explore the properties of objects. For example, we might follow the edges of an object’s surface with the fingers to assess its shape, might squeeze it to assess its hardness, or might slide our finger across it to sense its texture.

Kinesthetic Perception

The term *kinesthetic* originates from the Greek terms for “to move”, *kinein*, and “sensation”, *aisthēsis*²⁸, and describes the perception of movements. In this context, the notion of *proprioception*, stemming from the Latin words for “own”, *proprius*, and “a receiving”, *receptio*²⁹, is largely used as a synonym and relates to the perception of one’s own body movement. In the context of this thesis we refer to kinesthetic perception when arguing about the perception of forces and inertia that arise during interaction with (virtual) objects, as well as the perception of related object properties such as weight, size, or shape. Complementarily, and based on the original meaning of the term, we refer to proprioception when arguing primarily about perceiving one’s own body rather than objects.

While fine surface details are sensed by our tactile system, the kinesthetic system is responsible for the perception of the “*position and movement of our limbs and trunk, the sense of effort, the sense of force, and the sense of heaviness*” [Proske and Gandevia, 2012]. Central for perceiving these dimensions are mechanoreceptors in our muscles, tendons, and joints [Jones, 2000]:

- Primary spindle receptors sense changes in length of the muscles, and are responsive to the dynamics of the change, such as the velocity and acceleration of the length-change.
- Secondary spindle receptors also sense changes in length of the muscles but are less responsive to the dynamics of the change compared to primary spindle receptors. Instead, they signal static muscle length and positional information of the limb.
- Golgi tendon organs are located at the junction between tendon and muscle and sense the muscle forces and tension.

In addition to the receptors mentioned above, also signals from sensors in the

²⁸Origin of the word *kinesthetic*. <https://bit.ly/3gNmJh0>

²⁹Origin of the word *proprioceptor*. <https://bit.ly/3gLWBTR>

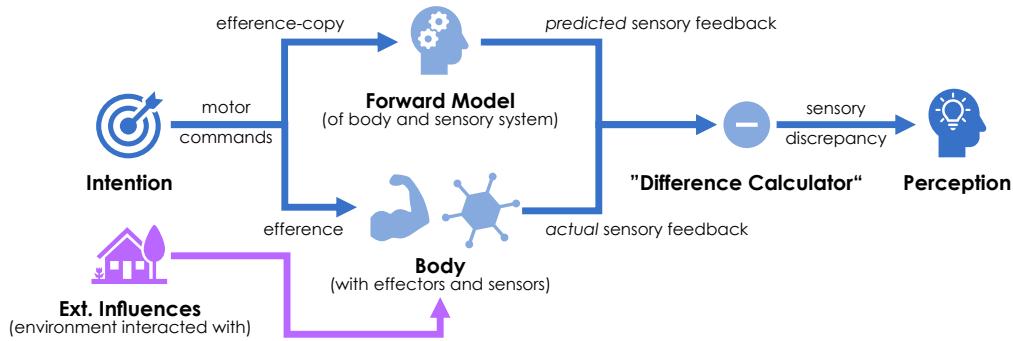


Figure 2.6: Perception Model (adapted from [Proske and Gandevia, 2012]).

skin can be involved in kinesthetic perception [Proske and Gandevia, 2012; van Beers et al., 2002]. Such skin-related information is especially important for the perception of the position and movement of our fingers [Jones, 2000]. Given all the data about the length of our muscles, the dynamics of the changes in length, the state of our skin, as well as the generated tension and forces, our brain infers the position and movement of our limbs, as well as the forces acting on our body. This, in turn, informs us also about the kinesthetic properties of objects.

Researchers proposed a four-staged model illustrated in Figure 2.6 of how such perceptions are formed [Proske and Gandevia, 2012; Jerald, 2015; pp. 73 f.]:

1. The process starts with an intended movement and the construction of motor commands in our brain to execute the movement.
2. These motor commands are then sent out as efferent signals to the motors (i.e., our muscles). In addition, a copy of the efferent signals (the efference-copy) is sent to a *"Forward Model"* [Proske and Gandevia, 2012]. This model estimates, based on the efference-copy, the input (i.e., the feedback) to expect from our sensors as a consequence of the executed action.
3. Upon reception of the motor commands, the muscles execute the action. The sensory input registered during execution depends on the executed commands themselves and external influences (e.g., the properties of an object interacted with).
4. The actual sensory feedback (the afferent signal) registered by the receptors is finally compared to the predicted sensory feedback. This step can be thought of as a difference calculation. The resulting discrepancy between expected and actual sensory feedback gives rise to a perception and informs about external stimuli, such as the properties of objects interacted with.

Perception of Object Properties

When assessing the properties of an object, one can distinguish between the properties related to its material, its geometry, and hybrid properties [Lederman

and Klatzky, 2009]. Material properties encompass hardness (hard or soft), thermal aspects (warm or cold), macro roughness (uneven or flat), fine roughness (rough or smooth), and friction (moist or dry, sticky or slippery) [Wang et al., 2019b]. Geometric information is given through an object's shape and size [Lederman and Klatzky, 2009]. Hybrid properties are aspects of an object that relate both to material and geometry. An example for a hybrid property is the weight and the weight distribution of an object, which is defined by the density of the material and its spatial distribution inside the object's volume.

Exploratory Procedures To learn about the haptic qualities of an object, we intuitively use "*exploratory procedures*" [Klatzky et al., 1987]. This means that we perform active movements and probing actions while in contact with the object to sample data about it. Such procedures help us leveraging the full potential of our tactile and kinesthetic senses.

Exploratory actions are executed in an economic way so as to "*produce the highest return for the least cost*" in terms of gained information, temporal efficiency and information accuracy [Klatzky et al., 1987]. Especially substance-related dimensions such as hardness and texture were found to be efficiently assessed through haptic exploration [Klatzky et al., 1987]. Typical examples for exploratory actions include, among others [Lederman and Klatzky, 2009]:

- moving the fingers laterally across an object's surface to sense its texture
- applying pressure to an object to sense its hardness
- following the contour of an object with the fingers to sense its shape
- enclosing an object with the hand to sense its size and volume
- holding an object without support to sense its weight
- wielding an object to assess properties like heaviness and shape-related attributes such as length or thickness [Turvey, 1996; Kingma et al., 2004]

Dynamic Touch Of particular importance for this thesis is the lattermost procedure of holding and moving an object to assess kinesthetic properties related to weight, size, shape, and inertia. This type of perception is known as "*dynamic touch*" [Turvey, 1996] and infers object characteristics from the physics of rotation, taking into account the muscular effort required to hold and rotate objects [Chan, 1995; Turvey, 1996; Kingma et al., 2004; and references therein].

The techniques proposed in Parts III and V build upon the findings of psychophysical research conducted by Chan, Turvey, Kingma et al. and colleagues, who conducted a series of experiments to investigate "*what mechanical quantities are present during the wielding of an object that constrain [...] one's perceptions of the object*" [Turvey, 1996]. Their results surfaced that our perception relies on physical invariants of the object that relate to rotation and the distribution of mass inside the object, i.e., mechanical quantities that do not vary with the speed or torque

with which we wield an object. Specifically, it was found that the haptically perceived length, width, and heaviness of an object wielded with the hand is a function of the object's invariant *moment of inertia* [Chan, 1995; Turvey, 1996], its invariant *static moment* [Kingma et al., 2002, 2004], or even both.

The moment of inertia (I) of an object is a physical quantity that “*determines the torque needed for a desired angular acceleration about a rotational axis, akin to how mass determines the force needed for a desired acceleration*”³⁰. It describes what can be thought of as the rotational resistance and is determined by the distribution of the object's mass relative to the rotation axis. The farther the mass elements of an object are from a rotation axis, the greater is I and the more torque and effort is required to rotate it about this axis. This, in turn, affects the force structure sensed by the kinesthetic system in the wrist when interacting with the object.

In the general case, the moment of inertia is captured by the inertia tensor, which describes the object's resistance to rotations about a point for different rotation axes. Any inertia tensor can be described by its eigenvectors and eigenvalues. The experiments of Turvey and colleagues revealed the eigenvectors and eigenvalues of the inertia tensor to be centrally important to our perception of many object properties. Specifically, for short rod-shaped, handheld objects (such as VR controllers), the haptic perception of length, width, and even more generally “*a crude type of shape perception*” [Turvey, 1996] is achieved as a function of the eigenvalues of the object's inertia tensor about the wrist. Moreover, Chan [1995] also found the felt and the known diameter of an object to influence our perception of its length. Here, however, the perception is inversely correlated to diameter with larger diameters reducing perceived length and vice versa. This effect is assumed to be linked to variations in wrist stiffness and muscle tension and the way these affect the perception of the moment of inertia [Chan, 1995].

Besides shape perception, initially, also the perception of heaviness was thought to be primarily related to the inertia tensor. However, Kingma et al. later found the static moment (M , given by $M = m \cdot d$, with m being the mass of an object and d being the distance between the object's Center of Mass (CoM) and the point of rotation) to be the most relevant mechanical parameter for the perception of heaviness when wielding an object, as well as for the perception of length when statically holding an object. The static moment was also identified to be important for length perception when wielding an object. For short rod-shaped objects as considered in this thesis, however, the results of Kingma et al. [2004] indicate the relevance of the static moment to be reduced and the importance of the inertia tensor to be increased (see also [Park et al., 2021]).

These results on dynamic touch show that our muscles and kinesthetic sensors drive our perception of object heaviness and shape by acting as “*smart instruments [...] registering invariants of rotational dynamics that connect physically to the spatial dimensions and other properties of the body and its attachments*” [Turvey, 1996]. In Parts III and V, we transfer these results from perceptual research to VR haptics.

³⁰Moment of inertia on Wikipedia. <https://w.wiki/3oZU>

2.2.4 Multisensory Perception

“If the task of the display is to serve as a looking-glass into the mathematical wonderland constructed in computer memory, it should serve as many senses as possible.”

– Sutherland [1965]

We usually perceive our environment with multiple senses at once and the sensory signals originating from our visual, auditory, haptic, gustatory, and olfactory systems are joined, giving rise to *multisensory perception*. Ernst and Bülthoff [2004] argue that *“the key to robust perception is the combination and integration of multiple sources of sensory information”*. Researchers distinguish two ways of how different sensory inputs are fused to arrive at a percept [Ernst and Bülthoff, 2004]:

- *sensory combination* occurs when different sensory channels complement each other (i.e., do not provide redundant information). Here, the inputs from different sensory modalities are combined so as to maximize information – arriving at a more robust perception.
- *sensory integration* occurs when different sensory channels provide estimates about the same aspect to be perceived (i.e., provide redundant information). As here, estimates from different modalities can be conflicting, an integration process combines the individual estimates into an overall percept.

Ernst and Banks [2002] found that the process of sensory integration, more specifically the integration of visual and haptic information, follows a principle that can be well described by a Maximum-Likelihood Estimation (MLE). During perception, the visual and haptic senses might arrive at different estimates of an environmental property that is to be perceived. At the same time, both the visual and haptic estimates produced by our nervous system are subject to different amounts of variance due to noise – or, in other words, differ in their reliability. The findings of Ernst and Banks [2002] indicate that the finally perceived property after sensory integration can be well described by a weighted sum that follows the MLE rule, which states that *“the optimal means of estimation (in the sense of producing the lowest-variance estimate) is to add the sensor estimates weighted by their normalized reciprocal variances”* [Ernst and Banks, 2002]. In other words, the visual and haptic information is combined in a way that weighs each individual modality according to their reliability. As such, the combined overall percept achieves a lower variance than the visual or haptic estimators [Ernst and Banks, 2002], increasing the reliability of our overall perception. As reviewed by Ernst and Bülthoff [2004], multiple studies have already confirmed that *“humans integrate information both within and across sensory modalities in just such an efficient way”*.

According to this model of multisensory integration, the reliability of the sensory modalities governs which of the modalities *dominates* our perception. In a VR system, discrepant sensory inputs, such as conflicting visual and haptic cues, can be a result of technical limitations or intentional effects. In this context, two special cases can arise: *visual dominance* and *haptic dominance*.

Visual Dominance

When our vision is more reliable than our haptic senses, perception is dominated by what we see rather than what we feel – a phenomenon known as visual dominance (or *visual capture*) [Ernst and Banks, 2002; Ernst and Bülthoff, 2004].

The effect of visual dominance can be remarkably strong, as already found in early experiments by Gibson [1933], in which participants wore glasses that distorted their vision. As a result, physically straight lines visually appeared to the participants as curved. When haptically exploring the lines while looking at them through the distortion glasses, participants even felt the lines as curved due to vision dominating their perception. Only when the visual perception of the lines was cut off (e.g., by closing the eyes or looking away), physically straight lines were also felt as straight. Gibson [1933] reported that “*this dominance of the visual over the kinæsthetic perception was so complete that when subjects were instructed to make a strong effort to dissociate the two, i.e. to ‘feel it straight and see it curved,’ it was reported either difficult or impossible to do so*”. Lederman and Klatzky [2009] further state that when assessing the properties of objects we interact with, “*the relative weighting of vision in relation to touch is greater when geometric properties are being judged than when material properties are tested*”.

Visual dominance also plays an important role in proprioception. In the perception of hand position, which is of central importance to our work in Part IV, for example, visual information dominates over proprioception for specific directions (e.g., in azimuth/horizontal direction) [van Beers et al., 2002].

The existence of the visual dominance phenomenon has dramatic consequences for the design of VR systems as it allows discrepancies between the visual and haptic senses to go unnoticed. Moreover, visual dominance allows for crossmodal effects, i.e., for example, to influence the perception of haptic properties by means of visual feedback. The combination of visual dominance with controlled visual-haptic conflicts opens a design space for creators of VR systems [Lécuyer, 2017]. Many techniques presented in this thesis rely on visual dominance to either convey intended perceptions of virtual object properties (e.g., as in Part III), or to facilitate interaction inside the IVE (e.g., as in Part IV).

Haptic Dominance

Complementarily to visual dominance, haptic dominance (or *haptic capture*) occurs when the variance of the haptic sensory channel is lower than that of the visual channel [Ernst and Banks, 2002]. In this case, haptic information is weighted more strongly than visual information.

Haptic dominance can be observed during the integration of specific visual and proprioceptive information, such as during the perception of hand position [van Beers et al., 2002]. Here, the weight of the haptic sense is increased, for example, as less visual information is available to subjects, or when subjects actively move

their hand themselves (in contrast to passively being moved by an experimenter). In extreme cases, proprioceptive weights can even exceed those of visual input [van Beers et al., 2002; and references therein].

Noteworthy in the context of this thesis is the experiment by van Beers et al. [2002]. In their study it was found that the reliability of vision and proprioception for estimating hand position varies also with direction. While for the horizontal direction, for example, vision dominates multisensory integration, in the depth direction, proprioception is weighted more strongly and dominates the perception of hand location [van Beers et al., 2002].

2.2.5 Conclusion

In the preceding sections, we introduced the basic principles of how humans perceive their environment with the Perception-Action loop [Ernst and Bülthoff, 2004], and reviewed how visual and haptic perception is formed.

Regarding visual perception, the phenomenon of change blindness and the way it is connected to eye blinks is of central importance to Part IV of this thesis, as is the phenomenon of visual dominance. It becomes apparent from our brief review of the physiology of blinks that the moments in which we blink are great opportunities for covertly manipulating the IVE. Being (1) a cause for change blindness, (2) usually unnoticed due to visual suppression, (3) frequently occurring for natural reasons, (4) relatively long (compared to saccades), (5) reliably and easily tracked with off-the-shelf eye trackers, and (6) triggerable through external stimuli, blinks lend themselves to being used for VR interaction techniques. Consequently, we take advantage of blinks in our work in chapter 9.

Our review of the haptic senses introduced the two classes of tactile and kinesthetic perceptions. While tactile perceptions arise from stimulation of our cutaneous receptors and inform about fine surface properties, kinesthetic perceptions arise from the sensory organs located in our tendons, joints, and muscles. Directly linked to proprioception, i.e., the perception of the movement and location of our own limbs, kinesthetic impressions inform about interaction forces and related object properties such as size or weight. Analyzing the haptic feedback conveyed by today's commercial VR systems makes apparent that only very basic forms of tactile sensations are supported. Today's VR controllers primarily rely on simple vibration and lack any means to stimulate the kinesthetic sensors. Kinesthetic effects are either not supported at all, simulated only through visual feedback, or mapped to abstract vibrotactile stimuli – an approach called *sensory substitution* [Cooper et al., 2018; Jerald, 2015; pp. 304 ff.]. It is this very lack of kinesthetic impressions in VR systems that motivates the research on proxy-based haptic feedback and related illusion techniques in this thesis. The need for novel feedback mechanisms compatible with handheld form factors further represents the key motivation for the development of the VR controllers in Part III.

Concluding our background review of human perception, we refer back to the

quote by Jerald [2015] at the beginning of this section. The way human perception works and combines inputs from different sources highlights that for a convincing simulation, perfect stimuli that are congruent across modalities are not always necessary. Perceptual phenomena like change blindness and visual or haptic dominance can be taken advantage of in VR [Steinicke et al., 2010a; Lécuyer, 2017] – relying on the human brain to “*fill in the gaps*” [Jerald, 2015; p. 61].

2.3 Haptic Feedback for Virtual Reality

“Being able to touch, feel, and manipulate objects in an environment, in addition to seeing (and hearing) them, provides a sense of immersion [...] that is otherwise not possible. It is quite likely that much greater immersion in a VE can be achieved by the synchronous operation of even a simple haptic interface with a visual and auditory display, than by large improvements in, say, the fidelity of the visual display alone.”

– National Research Council [1995; p. 162]

Including haptic feedback in computing systems has a long tradition, with applications spanning from everyday devices like smartphones to highly-specialized systems for research or industrial use [Burdea, 1999a]. The potential of haptics to enhance immersion has been recognized already in the early days of VR as this section’s introductory quote highlights. Augmenting IVEs with haptic sensations can yield positive effects that, depending on the use case, intensify presence [Insko, 2001], enhance interactions and interfaces [Hinckley et al., 1994; Lindeman et al., 1999], or improve task performance [Jiang et al., 2005; Burdea, 1999b].

Research on VR haptics roots back to as early as the 1960s [Brooks et al., 1990; Burdea, 1999b]. While most work in the early decades focused on what is in the following called active haptics, making use of world- and body-grounded actuation [Burdea, 1999a,b], the breadth of the field has increased significantly since then. As apparent from our review, many approaches are now also targeting the use at home, at work, or in other everyday contexts outside VR laboratories.

We start our overview with a summary of common taxonomies for VR haptics, followed by an in-depth review of different feedback concepts, devices, and systems that belong to the two key classes of active and passive haptics. Our review highlights the diversity of techniques considered in the past and allows to track the work in this thesis within the landscape of VR haptics.

2.3.1 Classification

Most commonly, haptic feedback systems are classified based on:

- the *perceptual category* of the feedback:
tactile or kinesthetic
[Srinivasan and Basdogan, 1997; Wang et al., 2020a; Jerald, 2015; pp. 37 ff.]

- the necessity of *contact* between user and device:
contact-based or contactless
[Rakkolainen et al., 2021]
- the realization in terms of *groundedness*:
ungrounded or body-grounded or world-grounded
[Srinivasan and Basdogan, 1997; Wang et al., 2020a; Jerald, 2015; pp. 38 f.]
- the realization in terms of the *actuation method*:
electrical motors or pneumatics or voice-coil actuators etc.
Wang et al. [2020a] provide a list of common actuation methods.

Moreover, most central for this thesis is the classification of haptic solutions for VR according to the *Haptic Reality-Virtuality Continuum* by Jeon and Choi [2009], also referred to as the *Active-Passive Haptics Continuum*. This continuum is depicted on the x-axis of Figure 2.7 and represents an adaptation of the Reality-Virtuality Continuum by Milgram and Kishino [1994] to the haptic modality. In line with Milgram and Kishino's original taxonomy, the haptic continuum spans from *Haptic Reality* to *Haptic Virtuality* covering *Haptic Mixed Reality*. Feedback solutions can be located within the continuum based on the extent of world knowledge that needs to be managed by the system to provide haptic feedback.

Haptic Reality refers to *Passive Haptics* [Lindeman et al., 1999], i.e., solutions that only rely on haptic sensations originating from the real world, such as from the interaction with tangible objects (called proxies). As passive haptic solutions do not need to compute the haptic stimuli conveyed to the user, they do not need to manage information for haptic rendering. In contrast to that, *Haptic Virtuality* refers to *Active Haptics* [Lindeman et al., 1999]), i.e., approaches that synthetically generate haptic stimuli by means of software and convey them with computer-controlled actuators. For this, active systems need to manage exhaustive knowledge about the objects in the IVE.

The class of *Haptic Mixed Reality* (also called *Mixed Haptics*) encompasses techniques that combine haptic sensations from the real world and artificially generated stimuli. Such solutions only need to manage partial information about the IVE to provide haptic feedback. As for the visual taxonomy, the extent of artificial stimuli and managed information thereby determines whether a technique belongs to the class of *Haptic Augmented Reality* or *Haptic Augmented Virtuality*.

Since VR systems usually involve multimodal feedback, Jeon and Choi [2009] combined Milgram and Kishino's taxonomy with the haptic continuum. The result is the 2D *Composite Visual-Haptic Reality-Virtuality Continuum* depicted in Figure 2.7. Here, the x-axis classifies a system's haptic feedback, and the y-axis classifies the visual feedback. The 2D continuum distinguishes nine classes of visual-haptic experiences. While *visual reality - haptic reality* (lower left corner) represents the real world, *visual virtuality - haptic virtuality* (upper right corner) represents the most literal form of VR with all-artificial stimulation.

The majority of related work presented in the following belongs to the uppermost

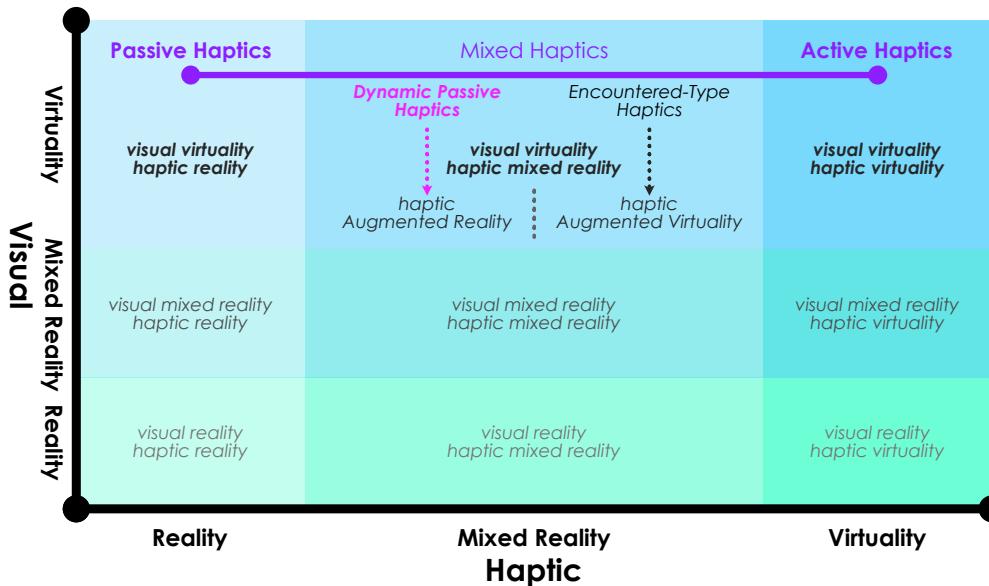


Figure 2.7: The 2D Visual-Haptic Reality-Virtuality Continuum (adapted from [Jeon and Choi, 2009]). The top row classifies the concept of Dynamic Passive Haptic Feedback proposed in Part III in the Active-Passive Haptics Continuum.

row in Figure 2.7. Moreover, the system presented in Part II of this thesis belongs to the upper left class of *visual virtuality - haptic reality*. The work in Part III and Part V can be classified as *visual virtuality - haptic mixed reality*, and the techniques studied in Part IV can be employed in any system that provides *visual virtuality*.

Besides this classification by Jeon and Choi [2009], haptic feedback systems for VR can also be categorized according to the *Haptic Fidelity Framework* recently proposed by Münder et al. [2022]. Here, haptic solutions are broadly categorized along the dimensions of *Haptic Fidelity* and *Versatility*. While *Haptic Fidelity* describes the degree to which a technique can provide realistic sensory feedback (a dimension ranging from *abstract* to *realistic*), the dimension of *Versatility* describes how specific a technique is tailored to an individual use case or whether it can be applied in a more general way to multiple application scenarios (a dimension ranging from *generic* to *specific*). The authors conclude that “[...] there are [...] no systems yet that fall into the category of very realistic but also very generic feedback. Such systems which create realistic feedback for any use case present the ultimate goal of haptic feedback for VR [...]” [Münder et al., 2022] – thus, to realize the ultimate display, VR haptics research needs to come up with solutions that provide flexible and highly realistic haptic impressions.

The following sections review the two poles of the Active-Passive Haptics Continuum. For each class, we outline the underlying concept, examples, as well as advantages and disadvantages. Each class's review is structured according to the classification that provides the most interesting view on the concept.

2.3.2 Active Haptic Feedback

Active Haptic Feedback (AHF) represents one extreme end of the Active-Passive Haptics Continuum and refers to *haptic virtuality* [Jeon and Choi, 2009].

General Concept

When Active Haptic Feedback (AHF) is employed, all haptic stimuli experienced by the user are artificially created by the VR system. Similar to how computer graphics algorithms visually render a scene by simulating light propagation and computing pixel colors, AHF employs *haptic rendering* algorithms to simulate and compute the kinesthetic forces and tactile stimuli that a user would experience during interaction with an IVE [Srinivasan and Basdogan, 1997]. AHF systems then *haptically display* the computed stimuli to the user. For this, AHF leverages computer-controlled actuation to exert the computed forces on the user's body, physically stimulating the user's haptic receptors [Burdea, 1999b].

To achieve convincing haptic feedback, AHF requires suitable actuators and end-effectors, i.e., actuated parts that are in contact with the user and through which stimuli are conveyed. Moreover, systems must precisely track user movements and manage information about virtual object properties like shapes and material. Software-wise, AHF rendering algorithms must provide a realistic simulation of interaction forces and tactile stimuli while achieving high update rates of around 1kHz [Srinivasan and Basdogan, 1997; Wang et al., 2019b]. As these requirements closely align with the field of robotics, a large synergy between active VR haptics and the field of robotics exists [Burdea, 1999a].

Most AHF techniques are specialized on either the kinesthetic or tactile modality (e.g., only vibration [Strohmeier and Hornbæk, 2017] or only forces [Van der Linde et al., 2002]). Others can render multimodal feedback by integrating tactile actuation inside the end-effector of a kinesthetic device [Choi et al., 2018; Wang et al., 2020a]. Further multimodal approaches integrate multiple tactile modalities such as skin-stretch, vibration and thermal feedback [Murakami et al., 2017]. In general, however, multimodal AHF interfaces are still rare [Wang et al., 2020a]. Complex multimodal combinations simulating, for example, "*softness, texture, thermal, 3-D shape, and weight*" [Wang et al., 2020a], do not exist today since actuators responsible for different modalities often spatially interfere, increasing the complexity of multimodal AHF interfaces.

Active Kinesthetic Feedback

Active interfaces providing kinesthetic sensations have received significant research attention in the past and have been the most common types of haptic devices for VR for a long time [Burdea, 1999b]. Since kinesthetic AHF devices exert forces and torques on the user's body, the physical action-reaction principle is of central importance for their design, as forces can only be applied relative to



Figure 2.8: Three examples of world-grounded kinesthetic AHF. Left: *PHANToM* force feedback device [Massie and Salisbury, 1994] (Source: Wikimedia Commons; public domain³¹). Center: *HapticMaster* (extracted from [Van der Linde et al., 2002]; © 2002 Van der Linde and co-authors). Right: *SPIDAR* system (extracted from [Sato, 2002]; © 2002 Sato).

a reference object. Thus, kinesthetic AHF is commonly categorized according to groundedness [Wang et al., 2020a; Jerald, 2015; pp. 38 f.].

World-Grounded Active Kinesthetic Feedback World-grounded devices are physically fixed to a static object (such as a table, or a wall) and provide forces that restrict the user's motion relative to the world. Prominent examples depicted in Figure 2.8 are motor-driven devices such as the *PHANToM* [Massie and Salisbury, 1994] and the *HapticMaster* [Van der Linde et al., 2002]. Most of these devices possess an actuated robotic arm with multiple segments and a stylus at the end (the end-effector). Alternatives to robotic arms are string-based force displays such as the *SPIDAR* [Sato, 2002]. Users interact with the IVE mediated through this end-effector, i.e., by moving it within the workspace of the device.

Haptic rendering is achieved by tracking the movement of the end-effector and constantly checking for collisions with virtual objects inside the IVE at high update rates. If the end-effector collides with virtual geometry, corresponding output forces are computed and applied to the end-effector, providing feedback and closing the perceptual feedback loop [Srinivasan and Basdogan, 1997].

Body-Grounded Active Kinesthetic Feedback Body-grounded solutions are worn by the user (e.g., exoskeletons or gloves) or held in the user's hand during interaction. Such devices can only apply forces or restrict movement relative to the parts of the body where the device is grounded. Examples are haptic gloves like the *CyberGrasp*³² and the *Rutgers Master II-ND* [Bouzit et al., 2002], which pneumatically actuates the fingers relative to the palm. *CLAW* by Choi et al. [2018], shown in Figure 2.9, is a body-grounded controller conveying forces through a built-in servo motor. Another example is the foot-worn *STROE*, which

³¹Image of a PHANTOM Premium A 1.5 haptic device made by SensAble from Wikimedia Commons. Public Domain. https://w.wiki/_z2kD

³²*CyberGrasp* system by CyberGlove Systems Inc.. <https://bit.ly/3mYW18Z>

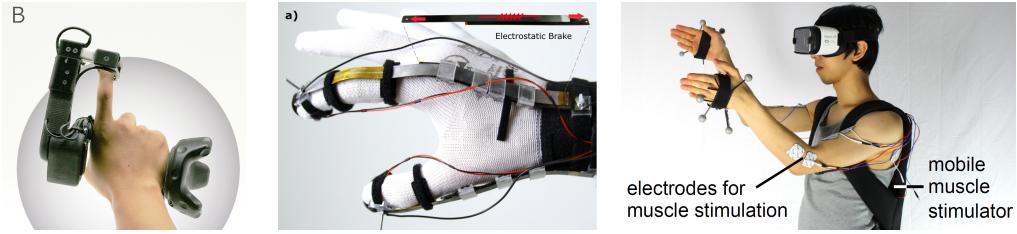


Figure 2.9: Three examples of body-grounded kinesthetic AHF. Left: *CLAW* (extracted from [Choi et al., 2018]; © 2018 Choi and co-authors). Center: *DextrES* providing haptic feedback based on electrostatic brakes (extracted from [Hinchet et al., 2018]; © 2018 Hinchet and co-authors). Right: Rendering a virtual cube by means of EMS (extracted from [Lopes et al., 2017]; © 2017 Lopes and co-authors).

connects to a handheld controller via an actuated string to simulate weight [Achberger et al., 2022]. Besides actuation-based solutions, also body-grounded systems based on breaks exist, such as *Wolverine* and *Grability* by Choi et al. [2016, 2017]. Brake-based mechanisms have also been integrated into exoskeletons, such as in the force feedback gloves *Dexmo* [Gu et al., 2016] (using mechanical locking) and *DextrES* [Hinchet et al., 2018] (using electrostatic brakes) to render the sensation of grasping objects in VR. Moreover, the VR controller *CapstanCrunch* by Sinclair et al. [2019] uses capstan-based breaking to render the sensation of grasping virtual objects of different compliance. In addition, Strasnick et al. [2018] showed with the breaking-based *Haptic Links* how the perceived realism when interacting with virtual two-handed objects in VR can be enhanced.

Apart from that, Lopes et al. [2015, 2017] proposed to provide the kinesthetic sensations of virtual impacts, virtual walls, and heavy virtual objects by activating the user's own muscles through EMS. When triggered, the EMS leads to involuntary muscle contraction, simulating force feedback.

A limitation of body-grounded solutions is that they cannot restrict the body to move relative to the real world. To account for this, Steed et al. [2020] recently proposed *Docking Haptics*, i.e., to dock worn interfaces to world-grounded devices. Within their workspace, this allows for rendering forces relative to the world, for example to simulate weight with a robotic arm docked to a glove.

Ungrounded Active Kinesthetic Feedback Ungrounded kinesthetic techniques are usually based on air jets, propellers, or inertia to convey forces.

Air jet-driven interfaces, such as the *AirGlove* [Gurocak et al., 2003], the handheld *AirWand* [Romano and Kuchenbecker, 2009], *JetController* [Wang et al., 2021], and *AirRacket* [Tsai et al., 2022a] release compressed air to produce repulsion forces. Most related to our work in Part III is the concept investigated by Suzuki et al. [2002] and Suzuki and Kobayashi [2005] shown in Figure 2.10. In their air jet-driven system users carry paddle-like objects (called air receivers) to interact with the IVE. When the paddle makes contact with a virtual object inside the IVE,



Figure 2.10: Three examples of ungrounded kinesthetic AHF. Left: A system releasing air from an air nozzle, which is felt through a handheld air receiver (extracted from [Suzuki et al., 2002]; © 2002 Suzuki and co-authors). Center: *Thor's Hammer* providing propeller-driven feedback (extracted from [Heo et al., 2018]; © 2018 Heo and co-authors). Right: *iTorqU 2.0* torque feedback device (extracted from [Winfree et al., 2009]; © 2009 IEEE).

spatially registered nozzles in the real environment release air streams against the air receiver inducing the feeling of a soft contact force.

Alternative approaches to air-based haptic feedback take advantage of propeller propulsion. Prominent examples are *Thor's Hammer* [Heo et al., 2018] depicted in Figure 2.10, *Wind-Blaster* [Je et al., 2018], *LevioPole* [Sasaki et al., 2018], *Aero-plane* [Je et al., 2019], and *PropellerHand* [Achberger et al., 2021]. While propeller-based techniques do not require compressed air, their limitations encompass latency, significant power requirements, and the produced wind and noise.

Besides air-based techniques, ungrounded kinesthetic solutions can also be based on the physics of inertia. The *iTorqU 2.0* by Winfree et al. [2009] shown in Figure 2.10 is an example of a handheld device that provides torque feedback leveraging a flywheel spinning inside an actuated gimbal. The gyroscopic effect has also been used to augment handheld mobile devices [Badshah et al., 2012] as well as HMDs [Gugenheimer et al., 2016] to provide kinesthetic feedback. Moreover, Shimizu et al. [2021] recently proposed *Unident*, a handheld VR controller that leverages one-dimensional weight shifts to provide the sensation of virtual impacts. The authors state that the design of *Unident* was inspired by the weight-shifting interface presented in Part III of this thesis. In contrast to our work, which leverages slow weight shifts inside the device, *Unident* leverages fast weight shifts to provide the impression of a virtual obstacle (e.g., a virtual tennis ball) colliding with the virtual object in the user's hand (e.g., a virtual tennis racket) by rapidly changing the controller's inertial state.

Active Tactile Feedback

While kinesthetic interfaces target the receptors in our muscles, tendons, and joints to convey forces and torques, tactile solutions stimulate the cutaneous receptors in our skin. For this, most active tactile interfaces directly contact the surface of the user's skin, but there also exist approaches that work contactless.

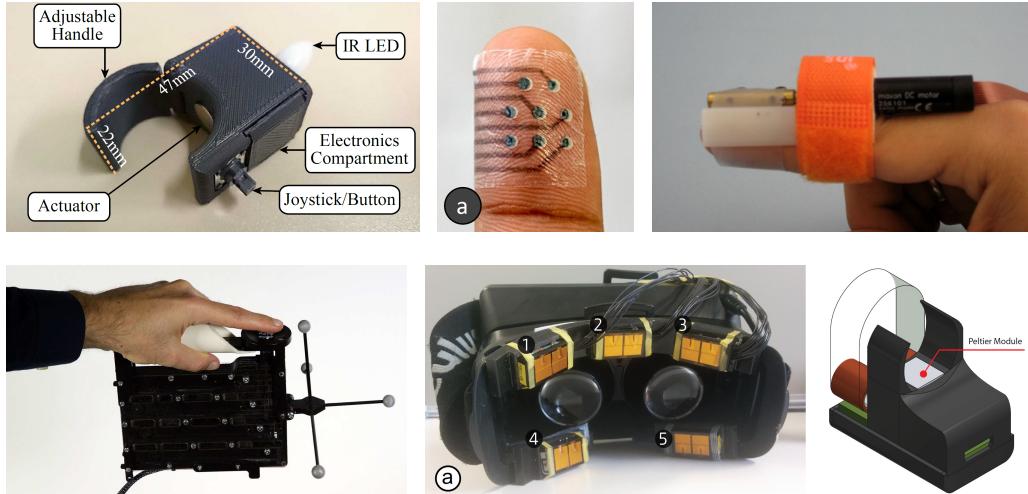


Figure 2.11: Six examples of contact-based tactile AHF. Upper left: *HapRing* (extracted from [Ariza, 2020]; © 2020 Ariza). Upper center: *Tacttoo* (extracted from [Withana et al., 2018]; © 2018 Withana and co-authors). Upper right: *GravityGrabber* simulating weight through skin-stretch (extracted from [Minamizawa et al., 2007]; © 2007 ACM). Lower left: *TextureTouch* (extracted from [Benko et al., 2016]; © 2016 Benko and co-authors). Lower center: *ThermoVR* HMD (extracted from [Peiris et al., 2017]; © 2017 ACM). Lower right: A multimodal device combining skin-stretching and thermal feedback (extracted from [Murakami et al., 2017]; © 2017 Murakami and co-authors).

Contact-Based Active Tactile Feedback Most active tactile interfaces leverage actuators that directly interact with the skin to convey vibration, pressure, or temperature. While some tactile interfaces aim to convey only tactile information such as texture details [Strohmeier and Hornbæk, 2017], others aim to convey kinesthetic impressions such as weight [Minamizawa et al., 2007; Choi et al., 2017; Lim et al., 2021; and references therein], forces [Tanaka et al., 2020], or torques [Provancher, 2014] by simulating the tactile component of such interactions.

The most common solution to actively provide tactile feedback is vibrotactile actuation. Most widely used are Eccentric Rotating Mass (ERM) actuators, Linear Resonant Actuators (LRAs), piezo-electric actuators, and voice coil actuators. Through precisely controlled asymmetric vibration patterns, the illusion of a force pulling or pushing the tactile device in a specific direction can be generated [Rekimoto, 2013; Choi et al., 2017; Tanaka et al., 2020]. Moreover, vibrotactile feedback can be used to convey different textures [Burdea, 1999a; Strohmeier and Hornbæk, 2017] or softness impressions of a physical proxy object [Choi et al., 2021]. Vibration constitutes the only active feedback channel of most commercial VR controllers and besides conventional programming and modeling of feedback patterns, vibrotactile feedback can also be designed through vocalization inside the IVE [Degraen et al., 2021a]. Moreover, vibrotactile feedback has been combined with active kinesthetic feedback, for example in the *CLAW* controller [Choi et al., 2018] and the *Canetroller* [Zhao et al., 2018]. Past research also investigated vibrotactile actuation integrated into different haptic wearables [Ariza,

2020; and references therein] such as rings [Ariza et al., 2015], thimbles [Ariza et al., 2018], gloves [Ariza et al., 2016; Hinchet et al., 2018], shoes [Strohmeier et al., 2020], and HMDs [de Jesus Oliveira et al., 2017]. Also, vibrotactile sensations can be provided through tactile vests like the *TactSuit X40* by bHaptics³³. Furthermore, Khosravi et al. [2021] recently combined vibrotactile feedback with a physically-based virtual hand to improve virtual mass perception.

Besides mechanical vibration, also electrotactile stimulation has been used to provide tactile feedback in VR. Here, electrodes placed on the skin stimulate the cutaneous receptors by means of electric current [Withana et al., 2018; and references therein]. Examples are the wearable devices proposed by Yem et al. [2016] and Hummel et al. [2016], as well as the concept of *Tacttoo* [Withana et al., 2018], which integrates electrotactile taxels into a temporary tattoo.

Actively controlled skin deformation constitutes another approach to tactile AHF. This concept has been explored by Minamizawa et al. [2007] with *Gravity Grabber*, a system leveraging an actuated belt wrapped around the finger to produce normal forces (i.e., pressure) and tangential forces (i.e., shearing) to convey different virtual weights. Recently, Williams et al. [2022] explored how normal, shear, and torsion feedback can be produced leveraging a wearable tactile device based on origami fabrication techniques. Another example for feedback based on skin-deformation is *Tasbi* [Pezent et al., 2019], which is a wrist-worn device that delivers both squeezing and vibrotactile feedback. Provancher [2014] applied the concept of skin-stretching feedback to VR controllers proposing handheld devices with sliding plates on the grip to provide force and torque cues. A related concept was also proposed by Sun et al. [2019] with *PaCaPa*. To render virtual shapes, researchers further proposed using extrudable and tiltable platforms as employed in the *NormalTouch* controller [Benko et al., 2016], or actuated pin arrays as used in the controllers *TextureTouch* [Benko et al., 2016] and *PoCoPo* [Yoshida et al., 2020]. Additionally, with *Tactile Drones*, Knierim et al. [2017] proposed to use aerial drones (e.g., caged quadcopters) to deliver tactile feedback.

Contact-based interfaces can also convey temperature with implementations commonly relying on Peltier elements to provide thermal cues on the skin. *Ambiotherm* by Ranasinghe et al. [2017], for example, utilizes thermal actuators attached to the user's neck in combination with a wind module to simulate the ambient temperature inside an IVE. Peiris et al. [2017] directly integrate thermal Peltier modules into an HMD in their *ThermoVR* system as shown in Figure 2.11 to provide directional cues and to simulate different weather conditions.

While most active tactile devices are specialized on a specific modality, also a few multimodal solutions exist, such as the *Sword of Elements* [Chen et al., 2016]. The device complements the vibrotactile feedback of commercial controllers with sensations of a rotating eccentric mass, wind, and temperature. Moreover, Murakami et al. [2017] combined temperature rendering with the belt mechanism of *Gravity Grabber* [Minamizawa et al., 2007] as shown in Figure 2.11.

³³*TactSuit X40* by bHaptics Inc. <https://bit.ly/3zFDmTf>

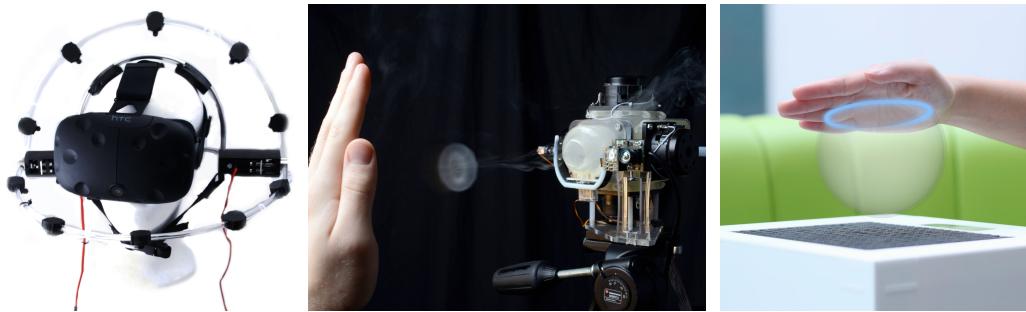


Figure 2.12: Three examples of contactless tactile AHF. Left: *VaiR* airflow simulation (extracted from [Rietzler et al., 2017]; © 2017 Rietzler and co-authors). Center: *AIREAL* shooting a vortex (extracted from [Sodhi et al., 2013]; © 2013 ACM). Right: Illustration of ultrasound haptic rendering (extracted from [Long et al., 2014]; © 2014 ACM).

Contactless Active Tactile Feedback Some active interfaces exist that can deliver tactile feedback without touching the user's skin.

Air-based tactile rendering systems leverage air streams directed at the user, generated by remote devices not in contact with the user. Moving air masses are produced either through fans or controlled exhaust of compressed air [Ranasinghe et al., 2017; Rietzler et al., 2017]. When impacting the skin, the air excites tactile receptors. The head-mounted *VaiR* system by Rietzler et al. [2017] depicted in Figure 2.12, for example, leverages compressed air to simulate air sources inside the IVE. Another approach, also shown in Figure 2.12, is to shoot rings of air (called vortices) onto the skin as proposed with *AIREAL* [Sodhi et al., 2013].

Researchers also managed to realize contactless tactile feedback through ultrasound [Rakkolainen et al., 2021; Long et al., 2014; and references therein]. Such techniques use arrays of ultrasound transducers that emit ultrasound waves, i.e., pressure waves which propagate through the air. In order to produce tactile sensations in mid-air, the activation of the transducers is algorithmically controlled. If timed correctly, the pressure waves of multiple transducers constructively interfere at desired focal points in space, which can be perceived through tactile receptors in the skin. Solutions for ultrasound haptics have already become commercially available (e.g., offered by the company Ultraleap³⁴).

Advantages and Drawbacks

We conclude our review of active haptics by summarizing the general advantages and drawbacks of the approach.

Advantages AHF for VR is a highly generalized approach [MacLean, 2000]. Relying on computer-controlled actuation, active techniques are generally flexible in terms of when and to what extent feedback is provided. Within their actuation

³⁴Ultraleap homepage. <https://bit.ly/3yHu9Z9>

range, VR system have full control over the feedback patterns, frequencies, and amplitudes delivered. Active interfaces can dynamically react to actions of the user within the IVE and convey dynamic (e.g., animated) virtual objects. Moreover, active kinesthetic and tactile feedback techniques have been extensively studied in the past, which resulted in a diverse set of approaches delivering verified effects that range from force and torque feedback to tactile sensations.

Drawbacks The advantages of active haptics, however, come at a cost. Active solutions are usually complex and costly in terms of mechanics and computation. They require continuous power supply and control electronics. Actuators, mechanical parts, electronics, and the power supply take up space, rendering many solutions bulky, immobile, and often imply severe workspace restrictions [Conti and Khatib, 2005]. Some types of actuators, like motors capable of applying significant forces and torques, further entail safety risks as failures or tracking problems might harm users [Burdea, 1999a]. Software-wise, and in line with the notion of *haptic virtuality*, AHF rendering requires exhaustive modeling of the physical properties of virtual objects, such as their shapes, textures, or materials. Moreover, AHF algorithms can necessitate complex simulations of the interaction between the user and the IVE at high update rates. Actuators need to execute commands rapidly and with low latencies so as to enable interactive experiences [Burdea, 1999a]. Finally, while solutions catering different modalities have been proposed, it remains difficult to provide multimodal feedback with active interfaces [Wang et al., 2020a] – a central limitation of AHF.

2.3.3 Passive Haptic Feedback

Passive Haptic Feedback (PHF) represents the second extreme end of the Active-Passive Haptics Continuum and corresponds to *haptic reality* [Jeon and Choi, 2009]. Here, we review the concept of proxy-based haptic feedback – the pivotal approach to VR haptics, which we advance in this thesis.

General Concept

While AHF simulates haptic stimuli and conveys them through computer-controlled actuators, Passive Haptic Feedback (PHF) completely refrains from any computational simulation and actuation. Instead, systems that employ PHF rely on *haptic proxy objects*, i.e., tangible artifacts in the real environment, also called *haptic props*, to convey the feel of virtual objects [Hinckley et al., 1994; Insko, 2001]. A haptic proxy is a physical object that provides haptic feedback for a registered virtual object simply through its inherent physical properties such as its shape, material, and weight [Lindeman, 1999]. Usually, proxies are spatially aligned with the virtual objects they represent and can be touched, picked up, or interacted with by immersed users. In contrast to AHF, proxies naturally provide multimodal haptic feedback, i.e., both kinesthetic and tactile sensations.



Figure 2.13: Two IVEs with passive haptic feedback. Left: A virtual, high-fidelity kitchen is represented by real, low-fidelity haptic proxies (extracted from [Kohli, 2013a]; originally from [Insko, 2001]). Right: A virtual cliff with the virtual ledge being augmented by a passive haptic proxy (extracted from [Insko, 2001]). All images © 2001 Insko.

The concept of PHF roots back to work by Hinckley et al. [1994], who proposed to leverage real-world objects in order to implement enhanced 3D UIs. The authors introduced a system for neurosurgical planning that utilized a tracked physical doll head and a plexiglass rectangle to take advantage of the user’s natural object manipulation skills. A little later, Lindeman et al. [1999] transferred the concept of physical proxies to fully-immersive VR, investigating how props can facilitate interaction with UI elements inside an IVE and found PHF to be of great value.

Insko [2001] finally extended the concept of PHF to large-scale IVEs as shown in Figure 2.13. In a series of experiments, Insko [2001] used low-fidelity proxies to physically resemble IVEs that were visually displayed in high-fidelity. The proxies used in this early work were rapidly constructed using readily available materials such as styrofoam, particle board, and plywood. Such proxies are considered low-fidelity as they do not perfectly match their virtual counterparts in terms of material or mass, or fine shape details. As Insko’s results show, however, due to visual dominance, even low-fidelity proxies can add to the user experience. Specifically, the addition of PHF to a stress-inducing IVE, as illustrated in the two rightmost images in Figure 2.13, was found to significantly increase presence according to subjective, behavioral, and physiological measures [Insko, 2001; Meehan et al., 2002]. Moreover, in a spatial navigation and learning task, the augmentation of a virtual maze with low-fidelity PHF walls significantly improved cognitive mapping of the IVE and training transfer to the real world [Insko, 2001]. Furthermore, research by Hoffman [1998] could show that once users experienced an object augmented with PHF, they continue to perceive other objects inside the IVE, which they have not interacted with, as more realistic. This indicates that PHF can benefit arbitrarily large IVEs provided that the objects the user is going to interact with are represented by appropriate physical proxies.

Realization

Compared to AHF, the realization of PHF is usually less complex and less costly. A VR system employing PHF requires physical proxy objects which can either be crafted for a specific IVE [Insko, 2001], repurposed everyday objects [Hoffman,

1998; Daiber et al., 2021; Makhsadov et al., 2022], or general purpose interaction objects like VR controllers. Compared to the virtual objects they represent, proxies can be of varying degree of mismatch [Simeone et al., 2015; Münder et al., 2019]. Discrepancies can be of a tactile nature, such as mismatching materials [Kitahara et al., 2010], or kinesthetic, such as mismatching weights [White et al., 2019]. Proxies can be high-fidelity physical replicas of their virtual counterparts (e.g., physical tools [Franzluebbers and Johnsen, 2018; Strandholt et al., 2020]), low-fidelity approximations (e.g., made out of building bricks [Münder et al., 2019] or materials like styrofoam [Insko, 2001]), or abstract representations not resembling the virtual counterpart at all or only very coarsely [Münder et al., 2019].

Secondly, proxies need to be spatially registered with their virtual counterparts. This can be done in a calibration step before, or at the start of the VR application if proxies are static [Insko, 2001]. If proxies are moved by the user, they need to be tracked to retain spatial registration with the IVE, for example by attaching tracking hardware or through optical tracking [Taylor and Cosker, 2020].

Moreover, VR systems that use PHF commonly track the user's hands and display virtual hand models inside the IVE to facilitate interaction. For this, motion capturing systems or specialized hand trackers (e.g., Leap Motion³⁵) can be used. Hand tracking also enables visual dominance-based manipulation techniques, which we introduce in section 2.6. To achieve visual-haptic synchronization when touching proxy objects despite tracking inaccuracies, additional capacitive sensing on the proxy's surface can be utilized [de Tinguy et al., 2021].

Kinesthetic and Tactile Proxy Design

Past research studied the importance of matching the basic haptic properties of inertia, shape, and material for proxy-based haptics in VR and MR.

Inertia Of particular importance to this thesis is research on inertial proxy properties. Since we infer many object qualities like weight, material, or shape from the inertial response of objects, our own previous research suggests that carefully designing the inertia of a proxy is of importance in order to control how users perceive virtual objects represented by proxies [Zenner, 2016]. We previously conducted two experiments studying the impact of different weight distributions inside a proxy and corresponding visual-haptic discrepancies on virtual object perception. Our results highlighted the weight distribution of proxy objects to be a highly influential proxy property that lends itself to generate perceptual illusions. Specifically, we found that by varying the weight distribution inside a proxy, other properties of the virtual counterpart can be simulated on a perceptual level, such as the virtual object's length and weight. Furthermore, mismatches in weight distribution lead to errors during interaction, and the direction of weight shift was found to be important for perceived realism [Zenner, 2016].

³⁵Leap Motion Controller by Ultraleap. <https://bit.ly/38FAXMp>

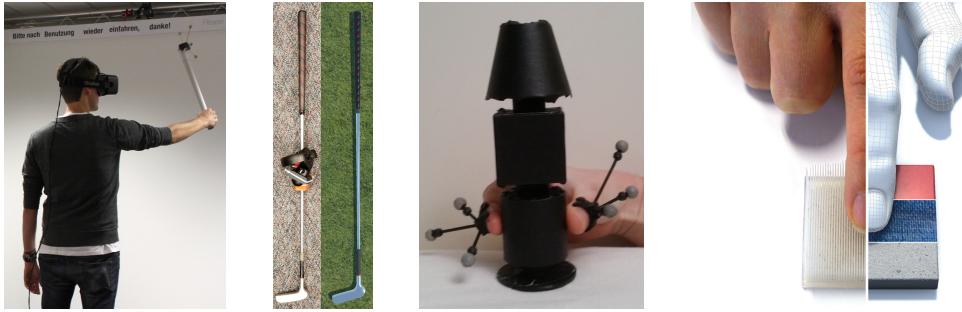


Figure 2.14: Previous works investigating proxy inertia, shape, and material. Left: Our previous experiment on proxy weight distribution (extracted from [Zenner, 2016]). Second from left: A tracked golf club proxy and the virtual club (extracted from [Franzluebbers and Johnsen, 2018]; © 2018 Franzluebbers and Johnsen). Second from right: A universal proxy with different graspable shapes (extracted from [de Tinguy et al., 2019a]; © 2019 IEEE). Right: A proxy surface featuring 3D-printed hair overlaid with different virtual materials (extracted from [Degraen et al., 2019]; © 2019 Degraen and co-authors).

Similarly, White et al. [2019] investigated the importance of absolute proxy mass. They compared a standard VR controller to three baseball bat proxies (an unweighted bat handle, a weighted bat handle, and a weighted bat handle with vibrotactile feedback) in a VR baseball simulation. Their results suggest that weighting haptic proxies correctly can significantly enhance certain aspects of the game experience and task performance. The authors conclude that for this and similar use cases “*the most important aspect is to approximate the weight of the original striking tool*” [White et al., 2019]. This conclusion is also backed by results of Franzluebbers and Johnsen [2018], who compared a standard VR controller to a tracked physical golf club with realistic weight shown in Figure 2.14 in a VR golf putting simulator. The authors found the heavier prop to be preferable in terms of both task performance and usability. In some cases, however, resembling the weight alone can be insufficient. This has been indicated, for example, by results of Martin et al. [2013], which highlight the importance of resembling both the weight and the shape of an object at the same time.

Shape Kwon et al. [2009] investigated size and shape differences when manipulating proxies in tangible AR. They found object manipulation to be more efficient when the physical and virtual objects matched in size and shape. Additionally, they found indication that maintaining matching shapes is more important than maintaining matching object sizes. Moreover, also Teng et al. [2018] and Yoshida et al. [2020] recently evaluated the extent to which the size of a virtual object can differ from the size of its proxy. Their experiments utilized the inflatable airbag proxy *PuPoP* and the shape array *PoCoPo*, respectively, as physical proxies rendering primitive shapes such as rectangles, cylinders, and spheres. Their results provide insights into acceptable visual-haptic size deviations for the tested primitives. Furthermore, de Tinguy et al. [2019a] studied how much physical and

virtual shapes pinched with thumb and index finger in VR might differ without users noticing the difference. The authors derived the Just-Noticeable Differences (JNDs) for mismatches in object width (5.75%), local surface orientation (43.8%), and local surface curvature (66.66%) relative to their reference object.

All these results indicate that perfect physical replication of virtual geometries is not always necessary and that single proxies can represent multiple different virtual objects. Concerning the choice of suitable proxies, Simeone et al. [2015] propose to pair virtual objects with those proxies that offer matching affordances and interaction possibilities at the parts users most likely interact with.

Material Several experiments investigated the perception of discrepant visual-haptic textures. The results of Iesaki et al. [2008] show that tactile impressions can be modified by manipulating the visual texture displayed during interaction. Their studies, which focused on roughness perception, also indicated the limitations of this technique as it was found that only certain degrees of mismatch yielded the intended illusions, specifically when the roughness properties of the visual and haptic textures were similar. When investigating the perception of hardness, Hirano et al. [2011] found that when users press a physical proxy surface, the perceived hardness can be manipulated by superimposing a visual animation in which the surface of the object is visually bending to different extents. Similarly, Nakahara et al. [2007] found the perceived sharpness of a physical edge on a proxy explored with the finger to be affected by the sharpness of an overlaid visual edge. By overlaying different visual materials onto various physical samples, Kitahara et al. [2010] observed that when for a visual material, different variations of that material are known to the user (e.g., through previous real-world experience), the physical properties of the haptic sample can control which variant of the visual material is perceived. For example, the authors report that *“when touching a stone plate with a rough surface while looking at a visual texture of steel, one subject had the impression of touching a steel plate that was not well polished.”* [Kitahara et al., 2010]. In our own previous work, we recently proposed using 3D-printed hair structures of different lengths (2.5mm increments) as tactile surfaces, overlaid with varying visual textures in order to evoke different material perceptions in VR [Degraen et al., 2019].

Application Areas

PHF offers many advantages over AHF, which render PHF suitable for a broad range of VR applications ranging from professional contexts to entertainment.

Jackson et al. [2013], for example, employed PHF in an interface for exploring 3D bioimaging datasets of thin fiber structures as shown in Figure 2.15. PHF has also been applied for simulation and specialist training. Martin et al. [2013], for example, explored the use of PHF for manual assembly simulation. Hasanzadeh and de la Garza [2019] investigated the risk-taking behavior of construction



Figure 2.15: Three applications that take advantage of PHF. Left: Interactive visualization of fiber structures (extracted from [Jackson et al., 2013]; © 2013 IEEE). Center: VR application for astronaut training (screenshot taken from YouTube³⁶; © 2017 NVIDIA). Right: VR system for simulated rock climbing (extracted from [Kosmalla et al., 2017]; © 2017 Kosmalla and co-authors).

workers when different levels of safety equipment are available leveraging a VR system with PHF. Moreover, Cooper et al. [2018] employed PHF in a VR system where users can change the wheel of a virtual racing car. Another example is the use of PHF for training astronauts³⁶. Besides that, PHF has also been taken advantage of for immersive sports. Examples here include the use of proxy-based haptics for VR (and MR) systems for rock climbing [Kosmalla et al., 2017; Zenner et al., 2018a; Schulz et al., 2019; Tiator et al., 2018], baseball [Gray, 2017; White et al., 2019], golf [Franzluebbers and Johnsen, 2018], and sailing³⁷.

Apart from that, PHF has been employed for immersive storytelling [Harley et al., 2017] and has seen numerous applications in entertainment (e.g., props used in VR theme parks³⁸ or escape rooms [Figueroa et al., 2019]).

Alternative Approaches to Passive Haptic Feedback

Departing from the conventional concept of proxy-based PHF as defined by Insko [2001], some works explored PHF following a second, alternative definition: *energetically passive haptic interfaces*. Such interfaces are defined by Swanson [2003] as “*interfaces [that] use energetically passive actuators which may in general only remove, store, or redirect kinetic energy within the system. They cannot add mechanical energy; all motive power must come initially from the human user, and the interface steers or dissipates that energy*”.

In the context of VR, Achibet et al. explored techniques that fall under this definition. Specifically, they studied how passive, kinesthetic force feedback can be conveyed through devices and mechanisms based on elasticity. With the

³⁶Video about NASA’s Virtual Reality Lab For Astronaut Training by NVIDIA. <https://bit.ly/2Wsfaope>

³⁷XRNAUT homepage. <https://bit.ly/3zr0Qvh>

³⁸Video *First look at THE VOID illustrating the use of PHF by The VOID*. <https://bit.ly/2WVxcjk>

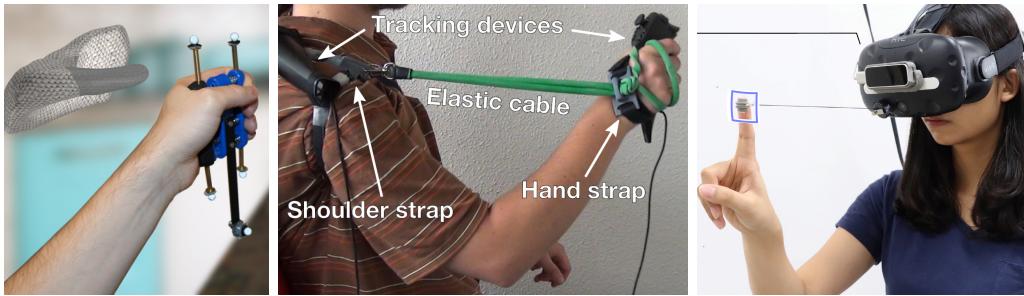


Figure 2.16: Three alternative PHF techniques. Left: *Virtual Mitten* leveraging the spring-based force feedback of a tracked hand exerciser (extracted from [Achibet et al., 2014]; © 2014 IEEE). Center: *Elastic-Arm* setup (extracted from [Achibet et al., 2015]; © 2015 IEEE). Right: *HapticSphere* setup (extracted from [Wang et al., 2019a]; © 2019 IEEE).

Virtual Mitten shown in Figure 2.16, they investigated how springs inside a proxy can provide sensations of different levels of effort [Achibet et al., 2014]. With *FlexiFingers*, a hand-worn device featuring bendable metal strips that connect the back of the hand with the fingertips, the concept of elastics-based force feedback was later generalized to individual fingers [Achibet et al., 2017].

Achibet et al. [2015] also extended the idea of elastics-based PHF to the human scale with the *Elastic-Arm* concept – a body-mounted armature that connects the user’s hand and shoulder through an elastic band, which, when stretched as the user reaches out with the hand, provides resisting forces. Following-up on this work, recently, Wang et al. [2019a] proposed *HapticSphere*, a similar interface that employs a non-elastic passive string attached to either the HMD, the user’s neck, or the shoulder to provide abrupt resistance feedback.

Advantages and Drawbacks

The general advantages and drawbacks of PHF, especially following the definitions of proxy-based feedback proposed by Lindeman [1999] and Insko [2001] applied in this thesis, can be summarized as follows:

Advantages PHF is a low-complexity approach to VR haptics that is characterized by being cheap in terms of equipment and computation. Props can be made out of inexpensive material, e.g., using 3D printing [Degraen et al., 2019, 2021b]. Since PHF props do not involve actuation, none or only minimal modeling of the physical qualities of virtual objects is required, and no computational simulation of haptic stimuli is needed. Such knowledge, however, is required to a certain extent when building props, which usually happens offline (i.e., before runtime). Proxies do not need to house actuators and, except for tracking, need no electronics or power supplies. As such, PHF, in contrast to AHF, does not entail significant safety risks and does not suffer from the same workspace limitations. Handheld proxies even support mobile and large-scale VR applications.

Previous research could show PHF to enhance presence, cognitive mapping, and training transfer [Insko, 2001]. Also, Lok et al. [2003] found the performance of tasks executed with PHF to better approximate real-world task performance. Props can enhance the “*sense of realism*” [Schulz et al., 2019] and even make objects not interacted with be perceived more realistically [Hoffman, 1998]. Finally, an essential advantage of PHF is that proxies simultaneously provide high-quality kinesthetic and tactile sensations – a level of multimodality that AHF fails to deliver. Taking advantage of users’ object manipulation skills [Hinckley et al., 1994], PHF is an intuitive haptic interface that frees users from the burden of wearing specialized haptic hardware [Wang et al., 2019b].

Drawbacks With the feedback of passive proxies being static, PHF is usually tailored to a specific IVE and the number of required props increases proportionally with the number of interactable objects in VR. As a consequence, the central shortcoming of PHF is its poor scalability and generality, which leads to PHF being used primarily for highly specialized applications. Moreover, VR systems employing PHF have only limited control over the haptic effects perceived by the user, cannot adapt the feedback to user actions, but can only provide fixed feedback patterns. To assemble suitable props, knowledge of the physical properties of virtual objects is required in advance and changes to the IVE require manual adjustments of the proxies. Additionally, PHF fails to provide appropriate feedback for animated objects that move inside the IVE (a notable exception to this is the pendulum approach by Cheng et al. [2018b]).

The following section 2.4 will identify the two central challenges that must be tackled by a VR system to fully realize the potential of proxy-based haptic feedback while overcoming the limitations of PHF.

2.3.4 Conclusion

Our review reveals AHF and PHF to be two contrasting approaches to VR haptics. As apparent from our summary in Table 2.2, the two opposite poles of the Active-Passive Haptics Continuum bear potential to complement each other. This intuition is also backed by early research of Clark and Bailey [2002] and Borst and Volz [2003], who both found the combination of passive tangible proxies and active feedback to be beneficial.

Acknowledging the great potential of proxies for rich, low-cost, realistic, immersive, and interactive tactile and kinesthetic feedback, this thesis aims to contribute to the field of proxy-based VR haptics. To advance the field, we explore techniques that help overcome the central drawback of PHF: its inherent inflexibility. By this, we aim to make proxy-based haptics less *specific* while maintaining *realistic* feedback (following the terminology of the *Haptic Fidelity Framework* by Münder et al. [2022]). Achieving this goal, however, requires solving two fundamental challenges, which we introduce in the following section.

	Active Haptics	Passive Haptics
Generality	generalized	specialized
Interaction	often through tools	direct with the hands
Feedback (Multimodality)	mostly unimodal	multimodal
Feedback (Variability)	variable	fixed
Supported Virtual Objects	static, interactable, animated	static, interactable
Costs (Equipment)	high	low
Costs (Computation)	high	low
Workspace	often limited	unlimited
Safety Risks	higher	very low
Tracking	required	partly required
Actuators	required	not required
Embedded Electronics	required	not required
Power Supply	required	not required
Physical Properties of IVE	required at runtime	required at development time
Body of Research	since 1960s	since 1990s
Classification¹	haptic virtuality	haptic reality
Haptic Fidelity²	low/medium	high
Versatility²	medium/high	low

Table 2.2: A high-level comparison of active and passive haptic feedback for VR. ¹ refers to the *Haptic Reality-Virtuality Continuum* by Jeon and Choi [2009]. ² refers to the *Haptic Fidelity Framework* by Münder et al. [2022].

2.4 Central Challenges for Proxy-Based Haptic Feedback

Proxy-based VR haptics can deliver high-quality tactile and kinesthetic feedback if sufficiently many and appropriate proxies are used. But how many proxies are sufficient and what proxies are appropriate?

For use cases in which users only interact with a very limited set of virtual elements, such as in VR simulators of aircraft or ship cockpits, conventional PHF, i.e., the naïve approach of employing (nearly) perfect physical replications is often a good option. Complete replication of the IVE, however, is only feasible and economically reasonable in such specific use cases [Insko, 2001]. In more general cases where users, for example, might want to freely explore large-scale IVEs within the confined tracking spaces at home, physically re-building the IVE and all contained objects is not an option. Lohse et al. [2019] summarizes that “*the utility of physical props decreases in proportion to the number of virtual objects the user can interact with*”. Moreover, as Insko [2001] states, “*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*”.

Hence, research has focused on improving the generality and scalability of proxy-based VR haptics with the goal to enable interactive experiences even with a limited set of proxies while sustaining a high feedback quality. Techniques that aim to provide successful proxy-based haptic feedback for VR thereby need to solve two high-level challenges, which we introduce in the following. These two

criteria have first been identified and mentioned by Lohse et al. [2019] and later by Strandholt et al. [2020]. Since then, we reviewed and classified these criteria and related research in two of our own previous works [Nilsson et al., 2021a,b].

This and the next three sections are based on the the following two publications of ours. The content provided here extends our previously published reviews and presents a substantially more comprehensive overview of the related work:

Nilsson, N. C., Zenner, A., and Simeone, A. L. (2021a). Popping Up Virtual Reality With Haptic Proxies. *IEEE Computer Graphics and Applications*, 41(05):104–112. © 2021 IEEE. Final published version available in the IEEE Xplore® Digital Library.

DOI: [10.1109/MCG.2021.3097671](https://doi.org/10.1109/MCG.2021.3097671)

Nilsson, N. C., Zenner, A., Simeone, A. L., Degraen, D., and Daiber, F. (2021b). Haptic Proxies for Virtual Reality: Success Criteria and Taxonomy. In *Proceedings of the Workshop on Everyday Proxy Objects for Virtual Reality at the ACM Conference on Human Factors in Computing Systems*, EPO4VR'21, pages 1–5. © 2021 Niels Nilsson and co-authors. Final published version available on the EPO4VR'21 workshop website.

Link: <https://bit.ly/3HK274S>

2.4.1 The Challenge of *Similarity*

The first challenge to be solved to provide successful proxy-based haptic feedback relates to the haptic qualities of the proxies being used. To deliver appropriate haptic sensations, VR systems need to solve the *Challenge of Similarity*:

“All haptic proxies touched by the user should feel sufficiently similar to their virtual counterparts with respect both to material properties (e.g., texture, hardness, and temperature) and geometric properties (e.g., shape, size, and weight)”

– [Nilsson et al., 2021a]

This criterion is only insufficiently fulfilled, for example, when a physical proxy is too abstract to provide convincing feedback for a virtual object, breaking the illusion of interacting with the virtual representation. The criterion of *Similarity* is perfectly fulfilled, however, when the physical object delivers haptic sensations during the interaction as convincing as an exact replica. As reviewed in the preceding sections, the degree to which proxy and virtual object can differ depends on the haptic property considered and the type of interaction.

2.4.2 The Challenge of *Colocation*

The second challenge to be solved by any VR system employing proxy-based haptics is concerned with the spatio-temporal matching of the physical and virtual environments. It can be summarized as the *Challenge of Colocation*:

“When the user touches a virtual object, it should be colocated with a haptic proxy in a way that allows for seamless interaction (e.g., the transformation of the virtual object should correspond to the position and orientation of the haptic proxy)”

– [Nilsson et al., 2021a]

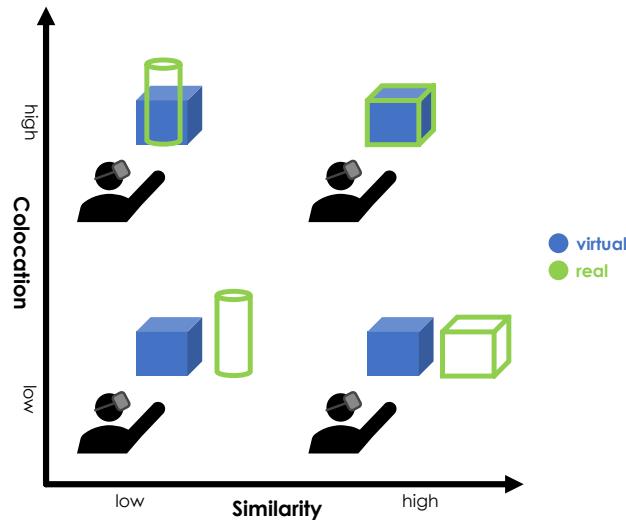


Figure 2.17: *Similarity* and *Colocation* Orthogonality (adapted from [Nilsson et al., 2021a]).

If a user, for example, reaches out to touch a virtual object but in the real environment only grasps for thin air, the VR system failed to establish sufficient colocation. The challenge of *Colocation* is perfectly solved, however, if a user can seamlessly interact with objects in the IVE, i.e., when the user perceives physical contact in synchronization with the virtual hands touching virtual objects. Some previous works investigated the impact of dislocation between physical and virtual objects. Gall and Latoschik [2018], for example, found offsets between a real and a virtual table of 20cm to significantly decrease self-reported presence compared to a condition with perfect alignment. Moreover, results of Fremerey et al. [2020] suggest that spatial offsets exceeding 1cm may increase user confusion and impair task performance. However, additional research on the tolerance of spatial misalignment in different scenarios is still required [Nilsson et al., 2021a].

2.4.3 Orthogonality & Classification of Solutions

Strandholt et al. [2020] identified the two challenges of *Similarity* and *Colocation* to be conceptually orthogonal as illustrated in Figure 2.17. One of the two criteria can be sufficiently fulfilled, while the other is not. If, for example, a VR system features property-wise perfect, but spatially offset, physical replicas of interactable objects in the IVE, the challenge of *Similarity* is solved but *Colocation* is only insufficiently established. On the other hand, if every interactable virtual object is represented by a perfectly colocated proxy, which greatly differs in terms of its haptic properties like shape, weight, or material, the system achieves *Colocation* but fails to achieve sufficient *Similarity*.

Furthermore, when both the challenges of *Similarity* and *Colocation* are fulfilled for specific objects in the IVE, the VR system can further achieve what is referred

to by Strandholt et al. [2020] as “*compelling contact forces*”. In this case and as a result of the physical interaction between the proxies, convincing contact and collision forces can arise, adding to the believability of the experience. It is also worth mentioning that the degree to which the criteria of *Similarity* and *Colocation* need to be fulfilled likely depends on the application. VR simulations in training contexts, for example, might require higher fidelity and more complete *Similarity* and *Colocation* than applications in, for example, the entertainment context, which might have relaxed requirements [Nilsson et al., 2021b].

As this thesis contributes by proposing and studying approaches tackling the challenges of *Similarity* and *Colocation*, in the following sections, we provide a review of related techniques introduced by previous research. While most of the reviewed techniques only propose a solution for either the *Similarity* or the *Colocation* challenge, some tackle both criteria at the same time [Nilsson et al., 2021a]. Following the taxonomy we proposed previously [Nilsson et al., 2021a,b], we distinguish techniques based on two dimensions:

1. according to *which reality* they operate in – are they manipulating the *physical* or the *virtual* environment?
2. according to *when* they operate – are they employed *offline* during development, or in *real-time* during application runtime?

The four resulting categories are distinct but not mutually exclusive. Some techniques are complementary and can be combined [Nilsson et al., 2021a].

2.5 Physical Approaches to Similarity and Colocation

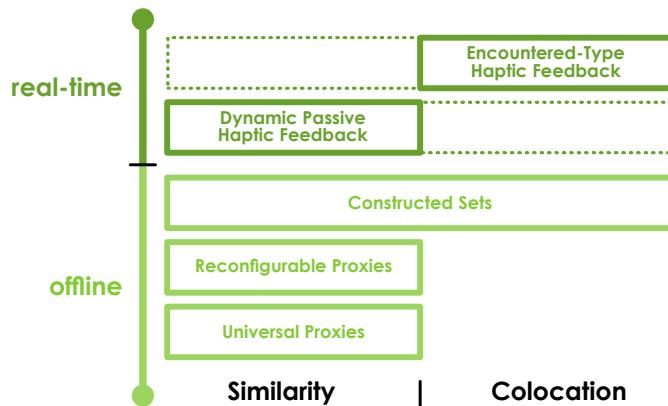


Figure 2.18: Offline and real-time approaches operating in the *physical* environment to achieve *Similarity* and *Colocation* (adapted from [Nilsson et al., 2021a]).

We start by reviewing *offline* and *real-time* techniques operating in the *physical* environment. Figure 2.18 provides an overview.

2.5.1 Offline Physical Strategies

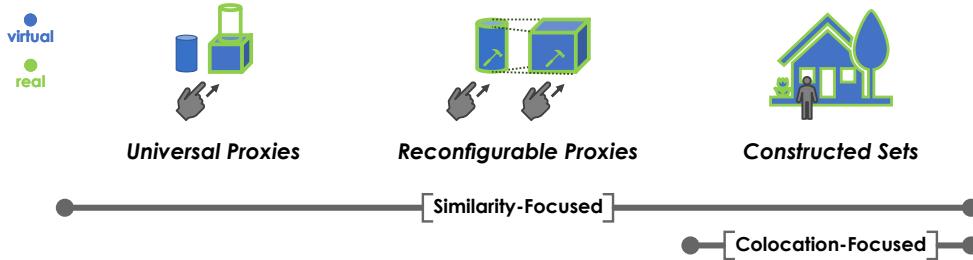


Figure 2.19: Overview of *offline physical* strategies.

Offline physical strategies to establish *Similarity* and *Colocation* make use of multi-purpose proxies or carefully prepared physical environments.

Universal Proxies

One strategy to achieve *Similarity* is to employ *universal proxies*. Such proxies offer various haptic features, such as different surface geometries, which make them suitable counterparts for multiple virtual objects in the IVE.

Cheng et al. [2017b], for example, proposed the use of sparse haptic proxies that consist of different surface primitives at different orientations and act as general-purpose touch surfaces in the real environment. The authors propose to combine sparse haptic proxies with haptic retargeting (introduced in paragraph 2.6.2) to steer the user's hand towards suitable primitives when reaching out to touch a virtual object. The concept of universal proxies has also been explored by de Tinguy et al. [2019a,b]. Figure 2.14 shows an example of a universal proxy by de Tinguy et al. [2019a]. The authors propose an algorithm that analyzes the geometries of virtual objects and universal props to find pinching locations on the real and virtual objects that haptically match best [de Tinguy et al., 2019b]. Moreover, they propose to automatically generate universal proxy designs (e.g., suitable for 3D-printing) based on analyses of the virtual scene and knowledge about tolerable discrepancies [de Tinguy et al., 2019a,b].

Reconfigurable Proxies

Besides universal proxies, which represent multiple virtual objects through a single physical state, an alternative approach to *Similarity* is to use *reconfigurable proxies*. Such proxies can be modified into different physical configurations, each representing different virtual objects. An example is *Haptobend* shown in Figure 2.20, a foldable proxy consisting of four rigid segments connected through hinges, which can be bent to represent different shapes [McClelland et al., 2017].



Figure 2.20: Three examples of reconfigurable proxies. Left: The foldable *Haptobend* (extracted from [McClelland et al., 2017]; © 2017 McClelland and co-authors). Center: A bending manipulator of the *TanGi* toolkit (extracted from [Feick et al., 2020a]; © 2020 IEEE). Right: The foldable board in the *iTurk* VR experience (extracted from [Cheng et al., 2018b]; © 2018 Cheng and co-authors).

To construct and reconfigure proxies offline, several toolkits have been proposed. The *VirtualBricks* toolkit by Arora et al. [2019], for example, is based on snap-together building bricks in combination with custom feature bricks that enable interactive functionality like rotation or translation of proxy parts. Zhu et al. [2019] proposed the *HapTwist* toolkit to construct proxies based on Rubik’s Twist³⁹. Moreover, with *TanGi*, Feick et al. [2020a] recently proposed a toolkit for rapid construction of haptic proxies leveraging 3D-printed composable shape primitives and manipulators for rotation, translation, bending (shown in Figure 2.20), and stretching that can be reconfigured with Velcro tape.

Reconfigurable proxies have also been explored in real-time contexts by Cheng et al. [2018b] with *iTurk*. In their story-driven VR experience, users travel through several virtual rooms. A single reconfigurable proxy (the foldable board shown in Figure 2.20) represents different virtual objects in each room (either a suitcase, a fuse cabinet, a railing, or a seat). For reconfiguration, the story leads users to interact with the proxy in a way that reconfigures it for the next room. The authors abstract this pattern in “*reconfigure-use-remap*” cycles [Cheng et al., 2018b].

Constructed Sets

A physical approach that can be applied offline and ensures both *Similarity* and *Colocation* is constructing a physical set that resembles the IVE. Such *constructed sets* consist of proxies representing walls, columns, furniture, and interactable objects. Due to the effort involved, this approach is primarily encountered in research or VR theme parks. An example is the work by Insko [2001] who physically recreated virtual scenes with low-fidelity props. Figure 2.13 illustrates two examples: a virtual kitchen and the upper level of a virtual pit room including a physical ledge. Pair et al. [2003] explored constructed sets for spatial AR.

³⁹*Rubik’s Snake* on Wikipedia. <https://w.wiki/4uCT>

To achieve *Similarity* between props and their virtual counterparts, set designers have two options: (1) creating physical replicas that perfectly match their virtual counterparts with respect to *all* properties, or (2) creating physical proxies that match *only perceptually relevant* properties of the virtual object, while accepting mismatches with respect to other properties. An example for the latter approach is the work by Fujinawa et al. [2017]. Following the theories of dynamic touch [Turvey, 1996; Kingma et al., 2004], Fujinawa et al. [2017] proposed a system for the construction of handheld proxies that intentionally differ from their virtual counterpart in shape and size. Designed based on a data-driven shape perception model, these proxies, however, carry weights so as to yield haptic perceptions during wielding that match those of the virtual objects. Thus, for example, even though the user might physically carry a short, weighted sword proxy, the illusion of wielding a longer sword can be achieved.

2.5.2 Real-Time Physical Strategies

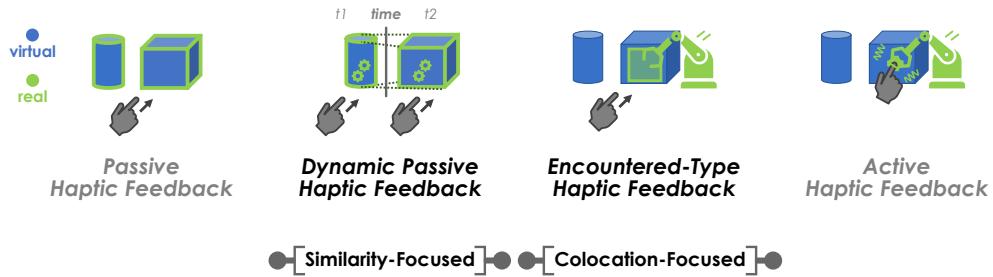


Figure 2.21: Overview of *real-time physical* strategies.

Real-time physical strategies are advanced versions of offline physical strategies [Nilsson et al., 2021a]. To adapt the physical environment while the VR application is running, they incorporate actuation. As such, they are classified as *Mixed Haptics* or *haptic mixed reality* [Jeon and Choi, 2009].

Two main real-time physical strategies tackling the challenges of *Similarity* and *Colocation* exist. Both strategies combine the advantages of AHF (i.e., flexibility) and PHF (i.e., feedback quality) to compensate for their respective drawbacks. They differ, however, in the degree to which they take advantage of actuation.

Dynamic Passive Haptic Feedback

Dynamic Passive Haptic Feedback (DPHF) is proposed and defined in Part III [Zenner and Krüger, 2017], and constitutes one of the main contributions of this thesis. While we introduce our own work on DPHF in later parts, we here review both earlier work classifiable as DPHF but conducted before the concept was formulated, as well as follow-up work on DPHF by other researchers.

DPHF enables a single proxy to represent different virtual objects with sufficient *Similarity*. To achieve this, DPHF employs *dynamic* proxy objects. These can, in contrast to *passive* proxies, dynamically adapt their passive haptic properties. Adaptation is accomplished by integrating actuators into the props that change the proxy's physical configuration. In contrast to AHF, DPHF does not use the actuators to actively exert forces on the user. Instead, actuation is used only to change the proxy itself and with this its passive haptic feedback.

From a conceptual point of view, DPHF is located close to the PHF-end of the Active-Passive Haptics continuum (as sketched in Figure 2.7). The feedback perceived by the user when DPHF is employed is still dominated by the real, physical qualities of the proxy. Moreover, DPHF rendering only requires basic knowledge about the haptic properties of virtual objects in the scene. Hence, DPHF can be regarded a form of *haptic Augmented Reality (AR)* [Jeon and Choi, 2009]. The actuation augments the PHF by matching it to the properties of different virtual objects – increasing the generality of the props.

Mass- and Mass Distribution-Changing Proxies Research outside the domain of VR has proposed devices that change their mass or mass distribution to provide haptic feedback. An early example is the *TorqueBAR* by Swindells et al. [2003], a 1050g construction with a motor that can move along a rail. The movement shifts the Center of Mass (Com) of the device to provide ungrounded kinesthetic feedback. The concept of shifting weights for haptic feedback has also been explored by Hemmert et al. [2010] in the context of mobile devices.

Schneider et al. [2005] first mentioned weight shifting in the context of VR haptics. The authors propose the handheld input device shown in Figure 2.22 to control a fork lift truck in a VR simulation. When picking up or releasing a virtual object with the truck, a heavy weight inside the controller moves to shift the balance of the device. Already in this early work, Schneider et al. mention that weight shifts can yield the impression of increased or decreased weights of the input device itself. They recognize the potential of haptic feedback based on weight shifting and propose as an avenue for future research an integration of this approach “*into a pointing device like it is often used in virtual reality applications*” [Schneider et al., 2005] to simulate the weight of virtual objects. In Part III, we take exactly this step and fill the gap in research by proposing two weight-shifting VR controllers.

In follow-up research, alternative implementations of 1D weight-shifting VR controllers have been investigated. Both Krekhov et al. [2017] as well as Kim and Baek [2021] integrated weight-shifting DPHF into gun-shaped VR controllers. Also, robotic arms have been mounted to handheld VR controllers in order to convey different virtual objects and interactions through weight shifting⁴⁰. Furthermore, Park et al. [2021] recently derived a length perception model that predicts the perceived length of a handheld object based on the VR controller's

⁴⁰UploadVR article about the weight-shifting *Nyoibo* prototype by Manuel Rosado. <https://bit.ly/3zXFTIK>



Figure 2.22: Three examples of weight-shifting DPHF. Left: 1D mass-shifting input device for virtual fork lifts (extracted from [Schneider et al., 2005]; © 2005 Schneider and co-authors). Upper right: *Transcalibur* shifting weights in 2D (extracted from [Shigeyama et al., 2019]; © 2019 Shigeyama and co-authors). Lower right: *SWISH* simulating 3D mass shifts of fluids (extracted from [Sagheb et al., 2019]; © 2019 ACM).

moment of inertia and diameter. Shigeyama et al. [2019] extended weight shifting DPHF to 2D with *Transcalibur*, shown in Figure 2.22. The authors use a data-driven approach to derive a “computational perception model” [Shigeyama et al., 2019], mapping device configurations to perceived 2D shapes. For this, they recorded the virtual shapes perceived by participants when interacting with different states of *Transcalibur* and fit a linear regression model to the obtained data pairs – effectively modeling how users *perceive* different device states. This model is then used for rendering target shapes. Extending the concept further, Sagheb et al. [2019] proposed *SWISH*, shown in Figure 2.22, a proxy leveraging weight shifts in 3D to simulate objects that contain fluids. Dynamic movements of the CoM can also be conveyed with the DPHF proxy *ElastOscillation* [Tsai et al., 2020], which modifies the oscillation properties of a metal weight to render the sensation of shaking elastic objects or handling virtual objects containing liquid.

Besides changing the mass distribution, researchers also explored devices that change their absolute mass. Niyama et al. [2014] proposed a weight-changing device that changes its mass by pumping liquid metal into or out of the object. The authors combine the weight change with projections of different visual textures in a spatial AR scenario to simulate various materials. Cheng et al. [2018a] transferred the concept of weight change to VR, proposing *GravityCup*. *GravityCup*’s weight is dynamically changed by transferring water between the proxy and a water tank carried by the user. Wang et al. [2022] combined the concepts of changing mass and mass distribution by proposing a VR controller add-on that can fill and empty two separate fluid chambers.

Stiffness-Changing Proxies To achieve *Similarity* when touching or grasping virtual objects, some DPHF proxies change their stiffness. Murray et al. [2018] proposed the VR controller handle shown in Figure 2.23, which can transform between two different states, each offering a different shape and level of compliance. Here, DPHF is realized through pneumatic actuation and an elastomer

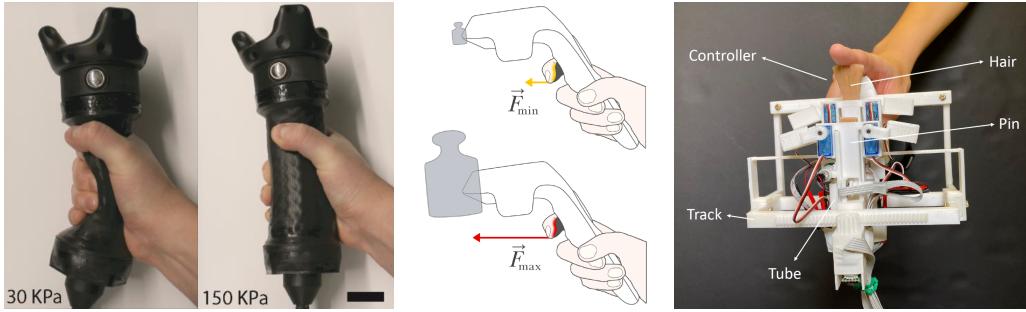


Figure 2.23: Three examples of stiffness-changing DPHF. Left: Pneumatic VR controller handle (extracted from [Murray et al., 2018]; © 2018 IEEE). Center: Conveying virtual weight through adaptive trigger resistance (extracted from [Stellmacher, 2021]; © 2021 IEEE). Right: *HairTouch* controller (extracted from [Lee et al., 2021]; © 2021 ACM).

lattice inside the handle, which determines the handle’s deformation properties. Alternatively, Stellmacher proposed *Triggermuscle*, a VR controller that can dynamically change the resistance of its trigger button to convey different weights as illustrated in Figure 2.23 [Stellmacher, 2020, 2021; Stellmacher et al., 2022]. Further examples for stiffness-changing DPHF are *Mouillé* by Han et al. [2020] and the *HairTouch* controller by Lee et al. [2021] shown in Figure 2.23, which renders roughness, stiffness, and surface elevation by changing the length and orientation of integrated hairs.

Shape-Changing Proxies Shape-changing interfaces have been proposed in diverse HCI disciplines (e.g., for mobiles [Roudaut et al., 2013]) and lend themselves also to providing haptic feedback for virtual objects [Alexander et al., 2018]. Zhao et al. [2017], for example, proposed a multi-robot system shown in Figure 2.24 for robotic assembly of low-resolution proxy shapes based on Zooids [Le Goc et al., 2016]. Another way of shape-changing DPHF was proposed with *PuPoP* [Teng et al., 2018], a palm-worn proxy that consists of inflatable airbags. The airbags are designed so as to yield the shape of the virtual object interacted with when pneumatically actuated, as illustrated in Figure 2.24. A solution to rapidly change the shape of dynamic proxies has been proposed by Gonzalez et al. [2021a] with *X-Rings*. Here, a VR controller with a built-in shape display serves as a shape-adapting proxy capable of rendering the surface of virtual objects through four expandable rings. Also *Adaptic* is an example for shape-changing DPHF [Gonzalez et al., 2021b] as the device can be considered a dynamic version of the reconfigurable *Haptobend* [McClelland et al., 2017] shown in Figure 2.20. Yet another concept for DPHF-based shape rendering was introduced by Kataoka et al. [2019], who proposed moving a length-changing proxy against a table surface to convey virtual shapes – a concept recently transferred to a finger-worn solution by Tsai et al. [2022b].

Besides that, proxy shapes can also be changed to modify other perceivable passive haptic properties, such as weight distribution or air resistance. We



Figure 2.24: Three examples of shape-changing DPHF. Left: Assembly of proxy shapes through small robots (extracted from [Zhao et al., 2017]; © 2017 Zhao and co-authors). Center: Inflatable *PuPop* (extracted from [Teng et al., 2018]; © 2018 ACM). Right: *ShapeSense* VR controller (extracted from [Liu et al., 2019]; © 2019 Liu and co-authors).

propose such an approach in Part III. Apart from our own implementation of this concept, a second realization was proposed by Liu et al. [2019] with *ShapeSense*. In contrast to our implementation and as shown in Figure 2.24, *ShapeSense* uses rigid sliding plates to shift weight and to adjust the proxy's air resistance.

Temperature-Changing Proxies In terms of tactile modalities, DPHF can also change the temperature properties of proxies. Ziat et al. [2014], utilized an electric heating pad and a Thermoelectric Cooler (TEC) unit to change the temperature of stainless steel cups, which served as tangible proxies in their VR demonstration scenario at UIST 2014. Moreover, *Mouillé* is a temperature-changing DPHF proxy, which, in addition to varying its stiffness, can also vary its surface temperature. This is realized by built-in TEC plates and was shown to successfully induce the illusion of interacting with wet virtual objects [Han et al., 2020].

Texture-Changing Proxies Besides temperature, also the surface texture of a proxy can be dynamically changed. *HairTouch* [Lee et al., 2021], for example, provides different impressions of surface roughness by adjusting the angle of inte-



Figure 2.25: Two examples of texture-changing DPHF. Left: *Haptic Palette* (extracted from [Degraen et al., 2020]; © 2020 Degraen and co-authors). Right: *Haptic Revolver* (extracted from [Whitmire et al., 2018]; © 2018 Whitmire and co-authors).

grated hairs. In our own previous work, we proposed the *Haptic Palette* [Degraen et al., 2020] shown in Figure 2.25, a handheld proxy that carries different material samples on a servo-actuated rotatable disk. The rotation mechanism can change the material presented to the user when reaching out to touch a virtual object. The design of the *Haptic Palette* was inspired by the *Haptic Revolver* proposed by Whitmire et al. [2018]. The *Haptic Revolver* is a handheld VR controller capable of rendering different textures to the index finger leveraging an actuated wheel as shown in Figure 2.25. While providing texture-changing DPHF, the *Haptic Revolver* additionally uses actuation for active rendering of contact and shear forces. As such, the device is conceptually located at the intersection of DPHF and Encountered-Type Haptic Feedback (ETHF) introduced next.

Encountered-Type Haptic Feedback

The second real-time physical strategy is *Encountered-Type Haptic Feedback* (ETHF) [Yokokohji et al., 1996]. ETHF has been proposed by McNeely [1993], Hirota and Hirose [1993], and Tachi et al. [1994], and is also referred to as *robotic graphics*. Mercado et al. [2021c] provides a recent and comprehensive review of ETHF research. ETHF primarily aims to solve the *Colocation* challenge by allowing users to establish physical contact with virtual objects in the IVE when and where they want. To achieve this, ETHF makes heavy use of actuators, such as robots, to move physical proxies or so-called *surface displays* in the real environment to locations where they are physically *encountered* by the user when touching a virtual object inside the IVE [Mercado et al., 2021a]. These relocated proxies can be the robots themselves or props attached to the robotic actuators [Mercado et al., 2021a,b]. Systems proposed in the past can be classified as either world-grounded, body-grounded, ungrounded, or based on human actuation.

Similar to DPHF, also ETHF combines aspects from PHF (i.e., proxy elements encountered by the user) and AHF (i.e., robotic systems). However, in contrast to DPHF, ETHF employs actuation also to actively render properties of encountered objects like rigidity. Hence, ETHF is located close to the AHF-end of the Active-Passive Haptics Continuum as illustrated in Figure 2.7. ETHF is highly dynamic and requires accurate tracking, as well as rapid and precise actuation [Burdea, 1999a]. As such, ETHF is, compared to DPHF, computationally and mechanically more complex and necessitates more knowledge about the virtual objects for active rendering. One can see ETHF as an evolution of AHF, augmenting active interfaces with the qualities of haptic proxies. Hence, ETHF can be considered *haptic Augmented Virtuality* (AV) [Jeon and Choi, 2009; Lee et al., 2020].

World-Grounded Systems Most ETHF systems make use of world-grounded actuators to present proxy elements. Tachi et al. [1994], for example, proposed shape rendering with a robot carrying a “*Shape Approximation Device*” featuring convex and concave edges, planes, and curves to simulate virtual shapes. Other aspects like inertia, viscosity, or stiffness are actively rendered by the

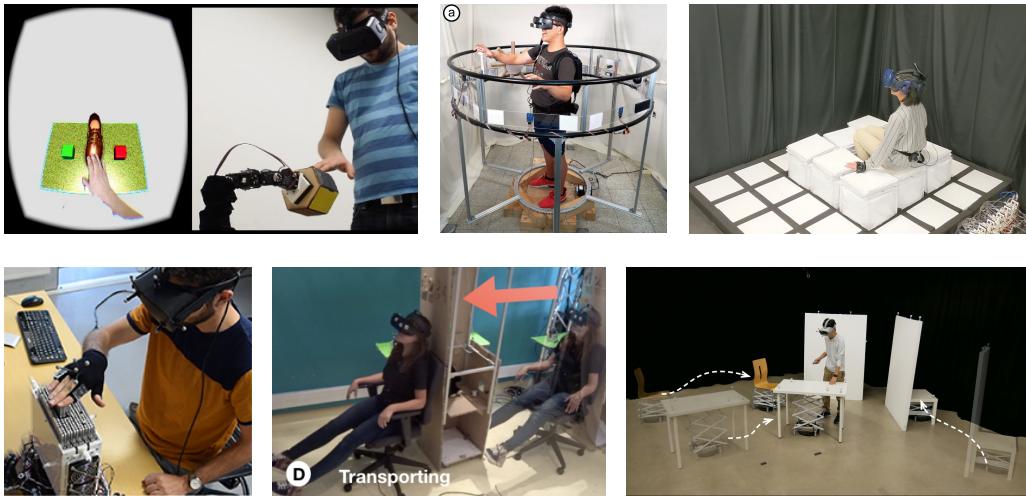


Figure 2.26: Six examples of world-grounded ETHF. Upper left: *Snake Charmer* (extracted from [Araujo et al., 2016]; © 2016 Araujo and co-authors). Upper center: *Haptic-Go-Round* (extracted from [Huang et al., 2020]; © 2020 ACM). Upper right: *TilePoP* (extracted from [Teng et al., 2019]; © 2019 ACM). Lower left: *shapeShift* (extracted from [Siu et al., 2018]; © 2018 Siu and co-authors). Lower center: The ceiling-mounted CoVR interface (extracted from [Bouzbib et al., 2020]; © 2020 ACM). Lower right: *RoomShift*'s robot swarm (extracted from [Suzuki et al., 2020]; © 2020 Suzuki and co-authors).

robot. Another example is the *Snake Charmer* by Araujo et al. [2016] shown in Figure 2.26. Their ETHF system features a commodity robotic arm and various proxy endpoints that the robot can autonomously exchange to render the position, orientation, texture, weight, or shape of virtual objects. Some endpoints even simulate physical buttons or haptically render airflow or temperature. Furthermore, the *Cobot*⁴¹ can render different virtual materials in the automotive context. Besides robotic arms, researchers also considered less complex actuation approaches. An example is the *Haptic-Go-Round* system by Huang et al. [2020] featuring a 1-DoF rotating platform that surrounds the user as shown in Figure 2.26.

Limitations of ETHF like constrained contact areas were approached by Mercado et al. [2020, 2021b] through interaction techniques and prop rotation. Moreover, Vonach et al. [2017] accounts for limited workspaces by combining ETHF with a locomotion platform. Another approach is to distribute actuators across the tracking space, as done in the *TilePoP* system by Teng et al. [2019]. Here, a floor of pneumatically actuated tiles is used, each holding stacked, cube-shaped airbags that can be inflated to form low-resolution proxy objects as shown in Figure 2.26.

Furthermore, some ETHF systems use mobile actuators to maximize their workspace. Examples for systems that move on tabletops are the mobile shape display *shapeShift* [Siu et al., 2018] depicted in Figure 2.26, and the *PhyShare* system [He et al., 2017] for collaborative VR experiences. Examples for room-scale ETHF systems are the *CoVR* [Bouzbib et al., 2020] interface shown in Figure 2.26, which

⁴¹Video about the *Cobot* by CLARTE-LAB. <https://bit.ly/3k3JHmb>

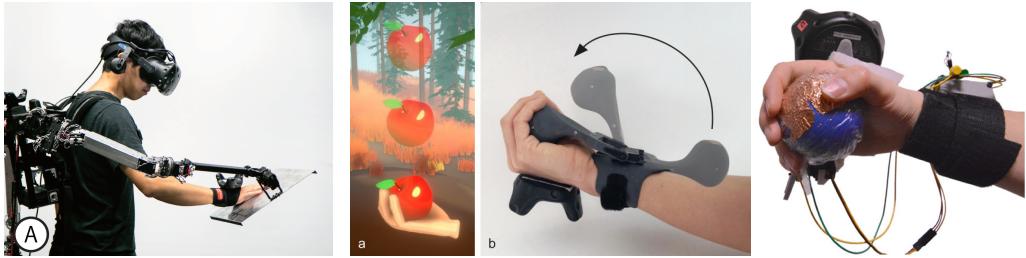


Figure 2.27: Three examples of body-grounded ETHF. Left: The *EncounteredLimbs* system (extracted from [Horie et al., 2021]; © 2021 IEEE). Center: A virtual apple caught by the user. *PIVOT* provides a proxy encountered for haptic feedback (extracted from [Kovacs et al., 2020]; © 2020 Kovacs and co-authors). Right: *WeATAViX* (extracted from [de Tinguy et al., 2020]; © 2020 de Tinguy and co-authors; Cropped; License: [CC BY 4.0](#)).

is based on a physical column moved by a ceiling-mounted robot, or *RoomShift* [Suzuki et al., 2020]. Unlike *CoVR*, *RoomShift* employs a swarm of mobile robots augmented with scissor lifts that can relocate furniture to adapt the physical environment to the IVE as shown in Figure 2.26. Moreover, Wang et al. [2020b] explored how cleaning robots carrying proxies can provide ETHF in VR.

Body-Grounded Systems ETHF can also be body-grounded. As shown in Figure 2.27, *EncounteredLimbs* [Horie et al., 2021] delivers ETHF in room-scale VR setups through a robotic arm attached to the user’s back that actuates a proxy surface. For kinesthetic ETHF, however, a drawback of body-worn systems is that users cannot lean against encountered objects such as virtual walls.

For tactile feedback, in contrast, body-groundedness is usually less of a limitation. The *Haptic Revolver* shown in Figure 2.25, for instance, can simulate contact sensations by lifting its wheel up against the finger, and conveys the sensation of sliding the finger over a virtual surface by rotating the wheel during contact. Moreover, both de Tinguy et al. [2020] (with *WeATAViX*) and Kovacs et al. [2020] (with *PIVOT*) developed wearable devices that move a proxy in or out of contact with the palm. This technique conveys the sensation of grasping or releasing a virtual object when closing or opening the hand as illustrated in Figure 2.27.

Ungrounded Systems With aerial drones, ungrounded ETHF systems can be realized that offer theoretically unlimited workspaces and forces relative to the physical world. Leveraging flying drones for ETHF was first proposed by Yamaguchi et al. [2016]. The authors attached a sheet of paper to a quadcopter as shown in Figure 2.28, which served as an encountered haptic surface. Later, Abdullah et al. [2017] explored using drones for rendering stiffness and weight. They propose leveraging safe-to-touch drones as shown in Figure 2.28 for ETHF, which actively throttle up or down to render forces. The results of Hoppe et al. [2018] indicate that when drones are used as ETHF displays, they are most suitable for rendering lightweight or compliant virtual objects.

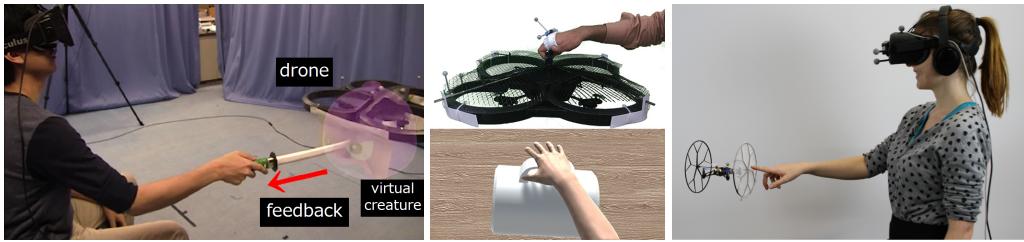


Figure 2.28: Three examples of ungrounded ETHF. Left: A proxy drone (extracted from [Yamaguchi et al., 2016]; © 2016 ACM). Center: *HapticDrone* (extracted from [Abdullah et al., 2017]; © 2017 Abdullah and co-authors). Right: Another proxy drone (extracted from [Hoppe et al., 2018]; © 2018 Hoppe and co-authors).

Human Actuation Systems Finally, Cheng et al. [2014, 2015, 2017a, 2018b] proposed substituting computer-controlled mechanical actuation with computer-instructed human actuation to implement ETHF for VR.

The *TurkDeck* system [Cheng et al., 2015] shown in Figure 2.29 employs multiple *human actuators*, i.e., additional people that manually relocate proxies in a just-in-time manner, instructed through projections on the floor and auditory feedback. In follow-up work, Cheng et al. [2017a] proposed to also make the human actuators themselves immersed VR users. They demonstrate this concept of *Mutual Turk* with a VR experience that lets one immersed user interact with a prop in such a way that the very same prop can be encountered in a second VR application simultaneously experienced by another user, as shown in Figure 2.29. Moreover, with the use of reconfigurable props in their *iTurk* system, Cheng et al. [2018b] demonstrated how immersed users can themselves serve as actuators that provide ETHF for later segments of their own VR experience.

ETHF based on human actuation profits from the flexibility, strength, and workspace of human operators, while facing challenges due to the limited speed and accuracy of humans, as well as the complexity of instructing and timing them.

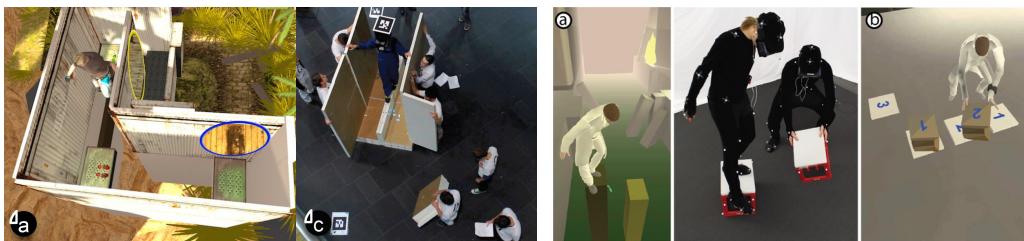


Figure 2.29: Two examples of ETHF based on human actuation. Left: A user of the *TurkDeck* system experiencing ETHF provided by a team of human actuators (extracted from [Cheng et al., 2015]; © 2015 ACM). Right: Two users, each immersed in their own VR application, share the same physical space and props. The right user places a prop to solve a puzzle. This prop is stepped on by the left user once the pillar in the left IVE is raised (extracted from [Cheng et al., 2017a]; © 2017 ACM).

2.5.3 Conclusion

When employing proxy-based haptic feedback for VR, *Similarity* for multiple scenes can be ensured with offline physical strategies, for example by utilizing universal or reconfigurable proxies. Moreover, *Colocation* can be achieved by optimizing the physical location of proxies in constructed sets. These offline techniques require careful planning of the IVEs and preparations of the physical counterparts. When offline strategies are applied, complexity is moved to the preparation phase and as a result, the flexibility to handle changes to the IVE at runtime remains limited.

More flexibility is usually provided by real-time physical techniques. Both Dynamic Passive Haptic Feedback (DPHF) and Encountered-Type Haptic Feedback (ETHF) generalize proxy-based haptics and can adjust the provided feedback to various IVEs at runtime and without additional preparation. This flexibility comes at the cost of increased hardware and software complexity. Real-time techniques can be regarded as automated variants of their respective offline counterparts. For example, DPHF proxies can be seen as actuated reconfigurable proxies, and ETHF can be interpreted as a dynamic and self-transforming version of a constructed set.

2.6 Virtual Approaches to Similarity and Colocation

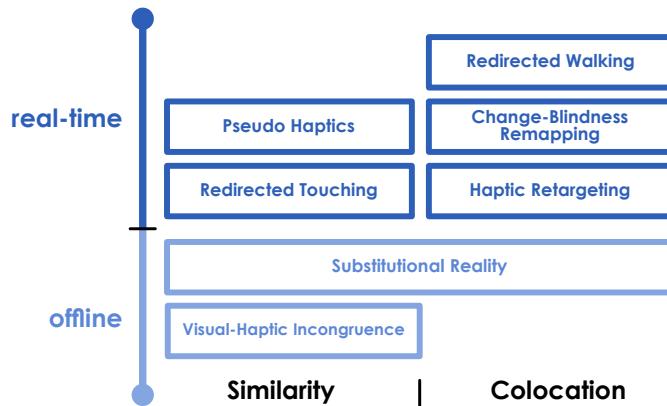


Figure 2.30: *Offline* and *real-time* approaches operating in the *virtual* environment to achieve *Similarity* and *Colocation* (adapted from [Nilsson et al., 2021a]).

Following our summary of physical strategies, we outline techniques operating in the *virtual* environment to achieve *Similarity* and *Colocation* for proxy-based haptics. Based on our taxonomy in Nilsson et al. [2021a], we again distinguish between *offline* and *real-time* strategies. Figure 2.30 provides an overview.

2.6.1 Offline Virtual Strategies

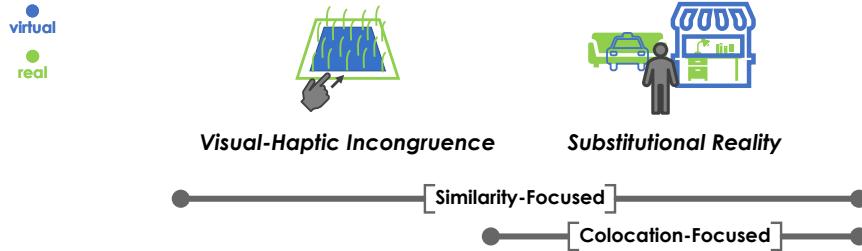


Figure 2.31: Overview of *offline virtual* strategies.

Virtual strategies towards *Similarity* and *Colocation* that are applied *offline* carefully prepare the visual appearance of the IVE and virtual objects touched by the user.

Visual-Haptic Incongruence

When touching a virtual object represented by a physical proxy, both the visual appearance and the proxy's haptic qualities influence how the object is perceived due to sensory integration. Techniques based on *visual-haptic incongruence* change the user's overall perception of an object by modifying the way it looks to achieve *Similarity*. A common application for visual-haptic incongruence is to modify an object's tactile impression. This can be achieved either by overlaying the proxy surface with a different visual texture, or by visually animating how the surface responds when being touched. Such visual-haptic mappings can be defined offline. For details on how to alter tactile perceptions in MR regarding material impressions, roughness, hardness, and edge sharpness, the results of Kitahara et al. [2010], Iesaki et al. [2008], Hirano et al. [2011], and Nakahara et al. [2007] can be consulted. Moreover, we demonstrated the potential of such techniques for VR recently by investigating how 3D-printed hair structures can convey diverse materials when being overlaid with different visual textures [Degraen et al., 2019].

Substitutional Reality

Another offline technique to achieve both *Similarity* and *Colocation* has been proposed by Simeone et al. [2015]. *Substitutional Reality* (SR) is the virtual counterpart to the physical strategy of constructed sets. Yet, instead of physically replicating the IVE, SR systems adapt the IVE to the existing physical environment. Each physical item in the tracking space, such as furniture, architectural features, or objects, is substituted by a virtual counterpart. Figure 2.32 shows an example.

Various levels of mismatch between physical objects and virtual substitutions are accepted, each representing a different trade-off between *Similarity* and design



Figure 2.32: An example of SR. Physical objects in the real room (center) are substituted by virtual counterparts, forming either a medieval courtyard (left) or a spaceship (right) in VR (extracted from [Simeone et al., 2015]; © 2015 Simeone and co-authors).

freedom. Simeone et al. [2015] distinguish between replication (max. *Similarity*), aesthetic mismatch, structural mismatch (addition/subtraction), functional mismatch (differences in affordances), and categorical mismatch (no *Similarity* between real and virtual objects). Their findings reveal those objects to be most suitable for substitution that have similar affordances [Simeone et al., 2015].

By substituting objects with virtual counterparts, SR also achieves *Colocation*. One way to construct SRs is by manually authoring them with a desired level of mismatch, such as in our own previous work on SR rock climbing [Kosmalla et al., 2017, 2020; Zenner et al., 2018a]. Specific 3D UIs can support the substitution process [Garcia et al., 2018]. Another way is to dynamically generate SRs based on 3D scans of the physical environment [Shapira and Freedman, 2016; Sra et al., 2016; Hettiarachchi and Wigdor, 2016; Hsu and Lin, 2021]. Moreover, the high degree of *Colocation* achieved in SR systems allows projecting the substitutions onto the physical objects using visible light projectors – effectively creating spatial AR experiences [Jones et al., 2013, 2014; Wiehr et al., 2016]. We explored this approach in our own previous work to let bystanders participate in an immersed user’s SR climbing experience [Zenner et al., 2018a, 2019a].

2.6.2 Real-Time Virtual Strategies

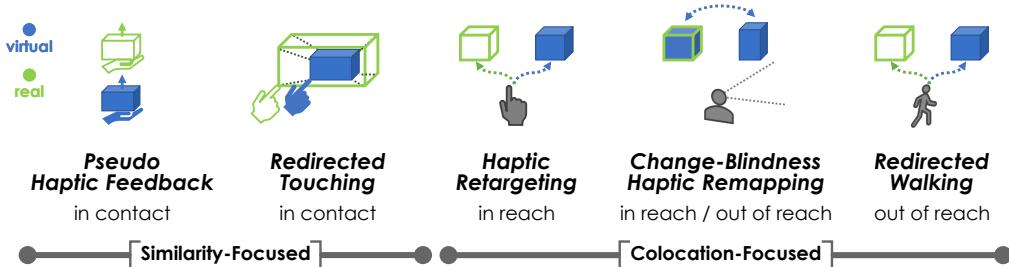


Figure 2.33: Overview of *real-time virtual* strategies.

Real-time virtual strategies achieve *Similarity* and *Colocation* by adapting the visual rendering at runtime, taking advantage of various perceptual phenomena.

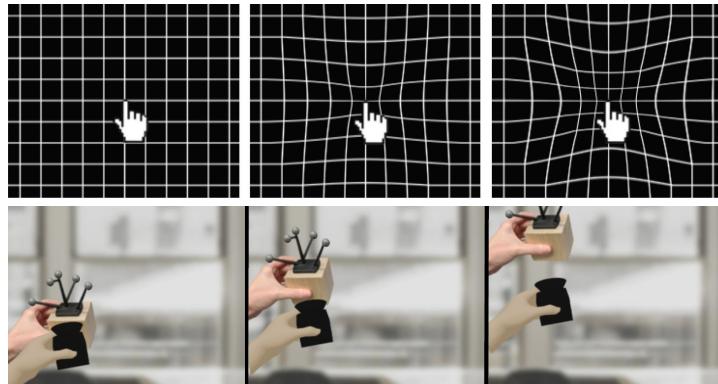


Figure 2.34: Two examples of pseudo-haptic feedback. Top: Conveying the elasticity of 2D images through visual deformation (extracted from [Argelaguet et al., 2013]; © 2013 ACM). Bottom: Conveying weight in proxy-based VR interaction through C/D ratio manipulation (extracted from [Samad et al., 2019]; © 2019 Samad and co-authors).

Pseudo-Haptics

Similarity can be achieved with *pseudo-haptic* feedback, a concept introduced by Lécuyer et al. [2000]. Pseudo-haptic effects are rather perceptual illusions than haptic feedback in the classical sense. Lécuyer [2017] captures this succinctly, writing that pseudo-haptic feedback creates “*a haptic illusion – that is, the perception of a haptic property that differs from the one present in the real environment*”.

To achieve such illusions, pseudo-haptic techniques are based on multisensory integration and take advantage of the fact that in VR, the haptic and visual senses can be independently stimulated. Specifically, pseudo-haptic feedback requires visual-haptic discrepancies [Lécuyer, 2009, 2017]. The concept grounds itself on the dominance of vision when we perceive certain spatial aspects of an interaction like distances, positions, sizes, and alike. Sensory integration resolves sensory conflicts introduced by pseudo-haptic techniques and arrives at a percept that differs from the haptic signal originally sensed according to the well-controlled, discrepant visual feedback. Many techniques utilize non-1-to-1 mappings from tracked to displayed user movement, i.e., Control/Display (C/D) ratio manipulation, to elicit pseudo-haptic effects. Such C/D manipulations accelerate or decelerate the virtual movements of the tracked cursor, hand, or object controlled by the user.

In their pioneering work, Lécuyer et al. [2000] demonstrate the pseudo-haptic rendering of friction and stiffness. Since then, pseudo-haptics has been used to simulate a variety of further haptic properties like mass, mass distribution, and textures. Lécuyer [2009] and Lim et al. [2021] provide reviews of related research.

In 2D, pseudo-haptic methods have been used to simulate virtual texture details like bumps and holes [Lécuyer et al., 2004], or elasticity [Argelaguet et al., 2013] as shown in Figure 2.34. In 3D, pseudo-haptics was employed to add feedback to purely virtual 3D UI elements [Gaucher et al., 2013; Speicher et al., 2019], or to

simulate virtual object weight [Hummel et al., 2013]. Furthermore, pseudo-haptic feedback was used in conjunction with the elastics-based PHF devices *Virtual Mitten* and *FlexiFingers* to convey different levels of grasping effort and stiffness, respectively [Achibet et al., 2014, 2017].

Importantly, pseudo-haptics can also be used in combination with haptic proxies to ensure *Similarity* [Lécuyer, 2009] (usually at the expense of *Colocation*). For example, a classic approach to change the perceived mass properties of a proxy is to modify the C/D ratio of the proxy's translations [Dominjon et al., 2005] or rotations [Yu and Bowman, 2020], with accelerated (decelerated) movements conveying lighter (heavier) objects [Samad et al., 2019]. Alternative approaches use similar visual cues such as virtual rubber bands connecting the tracked proxy and the virtual object [Palmerius et al., 2014].

Samad et al. [2019] investigated the perceptual aspects of pseudo-haptic weight simulation in detail. Based on the theories of multisensory integration, the authors derived a predictive model of perceived mass as a function of applied C/D ratio for interactions with a cubic proxy as shown in Figure 2.34. Rietzler et al. [2018] showed a similar pseudo-haptic approach to weight simulation based on overt tracking offsets to enhance presence even when lifting a purely virtual object with the VR controller instead of using a matching prop. Furthermore, Kim et al. [2022] investigated the benefits of combining pseudo-haptics and EMS-based feedback when rendering virtual weight. In addition, Maehigashi et al. [2021] recently found indication of the size-weight illusion to also occur in VR for a C/D ratio of 1 as well as for decelerated virtual movements. Moreover, Jauregui et al. [2014] investigated how modifications of the user's full-body avatar movements can convey weights when lifting virtual dumbbells with a proxy stick. Besides simulating absolute mass, Yu and Bowman [2020] contributed two pseudo-haptic techniques to convey the mass distribution inside a virtual object. Apart from VR, pseudo-haptic methods have also been employed to enhance the perception of mass in MR [Issartel et al., 2015].

It is noteworthy that pseudo-haptic perceptions depend on the interaction context [Lécuyer, 2017]. For example, slowing down a virtual movement by means of C/D manipulation can be interpreted as a change in friction [Lécuyer et al., 2000] or in mass [Samad et al., 2019], depending on the context – rendering pseudo-haptics a versatile instrument to tackle the *Similarity* challenge. A central drawback of pseudo-haptics, however, is that the magnitude of achievable effects is often limited, as indicated by the results of Samad et al. [2019].

Hand Redirection-Based Approaches

With *haptic retargeting* and *redirected touching* two further strategies exists to establish *Colocation* and *Similarity*. Both are based on the same real-time virtual technique known as *hand redirection*. Because of its importance for haptic retargeting and redirected touching, we dedicate Part IV to its most common variant.

Hand Redirection Hand redirection is a technique that grants the VR system control over the user's real hand movement. Specifically, when the user intends to move the virtual hand inside the IVE towards a virtual target, hand redirection can be applied to lead the user's real hand along a trajectory that is different from the naïve trajectory (i.e., towards the virtual target) the user's hand would follow if no hand redirection was applied. To achieve this, three different types of hand redirection techniques have been proposed [Kohli, 2013a; Azmandian et al., 2016a], each relying on the effect of visual dominance: *body warping*, *world warping*, and *hybrid warping*. Each of these three techniques introduces subtle visual manipulations in real-time during interaction, which force the user to adapt their real hand movement in order to reach their target. Yet, the different types of hand redirection differ in *what* is manipulated [Azmandian et al., 2016b]:

- *body warping* manipulates only *virtual body parts* but not the IVE, e.g., by offsetting the position of the virtual hand from the position of the real hand in order to provoke compensatory movements of the user's real hand.
- *world warping* manipulates only the *IVE* but not the user's virtual body, e.g., by translating the virtual scene, including the virtual target, in order to provoke an adjustment of the real hand trajectory.
- *hybrid warping* combines both concepts and applies both *body warping* and *world warping* at the same time. Thus, it manipulates *both virtual body parts and the IVE*, e.g., by introducing hand offsets and translations of the virtual scene simultaneously.

All three implementations of hand redirection lead to the user adapting their real hand trajectory during goal-directed hand movements inside the IVE. Most important for the work in this thesis is the common approach of body warping-based hand redirection, which we will contribute to in Part IV and thus review in detail in the following. Research results on the approaches of world and hybrid warping are discussed below in the context of haptic retargeting. Moreover, world warping is also used in the context of Redirected Walking (RDW) – a related real-time virtual technique introduced further below in this chapter.

Body Warping-Based Hand Redirection In the context of proxy-based haptics, a popular type of hand redirection is *Body Warping-based Hand Redirection* (HR⁴²). This type of hand redirection is implemented by introducing mismatches between the position and/or orientation of a tracked body part and its virtual representation in VR. Most commonly, the rendering of the user's virtual hand is spatially offset from the tracked position of the physical hand [Kohli, 2013a; Azmandian et al., 2016b] as illustrated in Figure 2.35.

⁴²As our work in this thesis will focus primarily on Body Warping-based Hand Redirection, for the sake of brevity, we will in the remainder of this thesis refer to this particular type of hand redirection only as HR, and denote references to other types of hand redirection explicitly by writing world warping-based HR or hybrid warping-based HR.

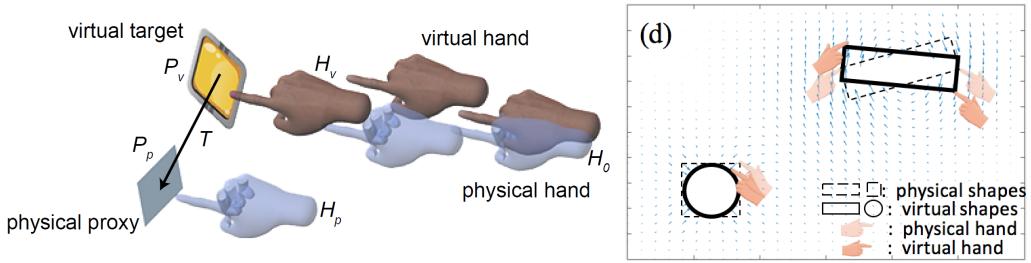


Figure 2.35: *Body warping-based HR* offsets the virtual from the physical hand. Left: Body warping-based HR as commonly implemented for haptic retargeting to establish *Colocation*: the redirection maps the physical point P_p to the virtual point P_v (extracted from [Cheng et al., 2017b]; © 2017 Cheng and co-authors). Right: Body warping-based HR as commonly implemented for redirected touching to establish *Similarity*: the redirection maps the physical shapes to their distorted virtual counterparts (extracted from [Zhao and Follmer, 2018]; © 2017 Zhao and Follmer).

HR techniques can go unnoticed by the user if applied within perceptual limits [Burns et al., 2006] and are restricted to HMD-based VR systems that deprive users of the view of their real hands [Han et al., 2018]. By refraining from a 1-to-1 Real-to-Virtual (R/V) mapping, body warping introduces visual-proprioceptive conflicts. When these are resolved under the influence of visual dominance, users tend to perceive their hand to be in the pose it is *shown in* rather than the pose it *physically is in*. As a consequence, users start issuing motor commands to the real hand that correspond to the virtual hand's spatio-temporal state – effectively compensating for any R/V hand offsets. By controlling these offsets, e.g., during goal-directed reaching, HR algorithms can *redirect* the user's real hand trajectory. Redirected touching and haptic retargeting techniques take advantage of this to make proxy-based interactions seem coherent, despite mismatches in terms of *Similarity* and/or *Colocation*.

Different implementations of HR have been proposed in the past. Algorithms for computing hand offsets range from simple approaches like constant offset vectors [Suhail et al., 2017; Han et al., 2018; Benda et al., 2020] or constant offset change rates [Burns et al., 2006; Feuchtner and Müller, 2018; Strandholt et al., 2020], over step-functions [Strandholt et al., 2020], geometric calculations [Ban et al., 2014], functional optimization [Zhao and Follmer, 2018], model predictive control [Gonzalez et al., 2022], space partitioning approaches [Montano Murillo et al., 2017], spatial interpolations [Kohli, 2013a; Spillmann et al., 2013], to combinations thereof [Strandholt et al., 2020], with linear interpolation [Azmandian et al., 2016b; Cheng et al., 2017b; Han et al., 2018; Matthews et al., 2019] being the most commonly-used approach. Most HR algorithms gradually increase or decrease hand offsets during interaction, while some implementations also involve instantaneous offset changes [Strandholt et al., 2020].

In most HR techniques, offsets are defined as a mapping from real to virtual space and applied to the hand rendering at runtime. Hence, depending on

their application context, HR algorithms were designed to either map a single physical point to a corresponding spatially offset virtual point as shown on the left in Figure 2.35 [Azmandian et al., 2016b; Cheng et al., 2017b; Han et al., 2018; Matthews et al., 2019], or map a set of physical coordinates (i.e., a shape or space) to a corresponding set of virtual coordinates (i.e., the distorted shape or space) as illustrated on the right in Figure 2.35 [Kohli, 2013a; Spillmann et al., 2013; Zhao and Follmer, 2018]. When interacting with grabbing tools or two fingers, some specialized techniques map physical grasp sizes to virtual grasp sizes [Yang et al., 2018; Bergström et al., 2019].

To compute the correct offsets, many techniques require knowledge about which object the user is going to interact with next. This information can be provided through scripted interactions, user selections, or predictions [Azmandian et al., 2016b; Cheng et al., 2017b; Matthews and Smith, 2019]. Other techniques, such as those that map complete spaces (also known as space warping approaches), are more flexible and do not require such predictions [Spillmann et al., 2013]. Moreover, when moving the hand from one redirected object to the next, it is desirable to prevent sudden offset changes that could result in visible jumps of the hand rendering. To this end, some techniques apply warps only within predefined spaces [Debarba et al., 2018b] or depend on resets to take place before the user can reach for another redirected target [Azmandian et al., 2016b; Matthews et al., 2019, 2021]. Such resets usually require the hand to return to a location where no offset is applied, commonly called the HR *origin*, and Matthews et al. [2021] recently proposed various adaptive reset strategies to optimize redirection sequences. Other HR techniques are more flexible and do not require resets as they smoothly blend offsets of subsequent redirections [Spillmann et al., 2013; Cheng et al., 2017b; Yang et al., 2018]. Still others directly apply the complete offset when the hand enters the interaction zone and only then display the virtual hand [Suhail et al., 2017].

Users can accustom to HR over time [Han et al., 2018] and to achieve increased presence and ownership when using a virtual hand representation, considering the user's gender and rendering a realistic avatar hand is preferable over using an abstract hand representation [Argelaguet et al., 2016; Schwind et al., 2017]. Moreover, while Matthews et al. [2022] consider the minimum distance between hand and target to determine offsets, most HR techniques compute hand offsets according to only a single, fixed reference point on the hand (e.g., the location of the physical fingertip) and then apply this offset rigidly to the complete hand model [Kohli, 2013a; Ban et al., 2014]. Some works applied HR to handheld tools in the same way [Spillmann et al., 2013], and only few approaches support multi-finger interaction [Zhao and Follmer, 2018], bimanual redirection [Gonzalez and Follmer, 2019], or are designed to work specifically with two reference points [Yang et al., 2018; Bergström et al., 2019]. Also, only few works introduced hand orientation mismatches [Kohli, 2013a; Azmandian et al., 2016b].

Table 2.3 summarizes research on different HR techniques. It is, to the best of our knowledge, the first systematic overview over this topic area.

		Hand Redirection Technique			Requirements			was applied ...			Support	
Category	Reference	R/V-Mapping	Offset Behavior	Offset Computation	Prediction	Reset	... to	... with	... for	Context	HaRT	
Redirected Touching	Kohli [2013a]	shapes/space	grad. changing	spatial interpolation	no	no	hand	PHF	Similarity	Cockpit Procedures	variant	
	Spillmann et al. [2013]	shapes/space	grad. changing	spatial interpolation	no	no	tool	PHF	Both	Surgical Sim.		
	Ban et al. [2014]	shapes/space	grad. changing	geometric calculation	no	no	hand	DPHF	Similarity	Basic Research		
	Abtahi and Follmer [2018]	point	grad. changing	geometric calc. & gain	yes	yes	hand	PHF & AHF	Both	Shape Displays		
	Yang et al. [2018]	grasp size	grad. changing	linear interpolation	yes	no	tool (2p.)	PHF	Similarity	Basic Research		
	Zhao and Follmer [2018]	shapes/space	grad. changing	functional optimization	no	no	hand	PHF	Both	Basic Research		
	Bergström et al. [2019]	grasp size	grad. changing	linear interpolation	yes	no	hand (2f.)	PHF	Similarity	Basic Research		
	Strandholt et al. [2020]	point	inst. & grad. chg.	linear interp. & step func.	no	yes	tool	PHF	Both	Carpentry Sim.		
Haptic Retargeting	Strandholt et al. [2020]	point	grad. changing	linear interpolation	no	yes	tool	PHF	Both	Carpentry Sim.	Carpentry Sim.	
	Strandholt et al. [2020]	point	grad. changing	constant rate	no	n.a.	tool	PHF	Both	Carpentry Sim.		
	Azmandian et al. [2016b]	point	grad. changing	linear interpolation	yes	yes	hand	PHF	Colocation	Basic Research		
	Carvalheiro et al. [2016]	hand position	grad. changing	gain	no	no	hand	PHF	Colocation	Basic Research		
	Carvalheiro et al. [2016]	shapes/space	grad. changing	spatial interpolation	no	no	hand	PHF	Colocation	Basic Research		
	Cheng et al. [2017b]	point	grad. changing	linear interpolation	yes	no	hand	PHF	Both	Basic Research		
	Suhail et al. [2017]	constant	constant	always on	yes	no	hand	PHF	Colocation	Basic Research		
	Han et al. [2018]	point	constant	always on	yes	no	hand	PHF	Colocation	Basic Research		
Ergonomic Interfaces	Han et al. [2018]	point	grad. changing	linear interpolation	yes	yes	hand	PHF	Colocation	Basic Research	yes	
	Abtahi et al. [2019] ²	point	grad. changing	linear interpolation	yes	no	hand	ETHF	Colocation	Drone Hovering		
	Gonzalez and Follmer [2019] ^{1,2}	point	grad. changing	linear interpolation	yes	yes	2 hands	PHF	Colocation	Basic Research		
	Matthews et al. [2019] ^{1,2}	point	grad. changing	linear interpolation	yes	yes	hand	PHF	Both	User Interfaces		
	Gonzalez et al. [2020] ¹	point	grad. changing	linear interpolation	yes	yes	hand	ETHF	Colocation	Basic Research		
	Fang and Harrison [2021] ¹	hand position	grad. changing	linear interpolation	no	no	hand	body	Colocation	Self-Haptics		
	Geslain et al. [2021] ³	point	grad. changing	linear interpolation	yes	yes	hand & tool	PHF	Colocation	Training		
	Gonzalez et al. [2021b] ¹	point	grad. changing	linear interpolation	yes	yes	hand	DPHF	Colocation	Basic Research		
Threshold Investigation	Matthews et al. [2021] ²	point	grad. changing	linear interpolation	yes	yes	hand	PHF	Colocation	Adaptive Resets	yes	
	Gonzalez et al. [2022]	point/path	grad. changing	model predictive control	yes	no	hand	VTs	Colocation	Basic Research		
	Matthews et al. [2022] ²	shapes	grad. changing	linear interpolation	yes	no	hand	(PHF)	Colocation	Shape Aware		
	Montano Murillo et al. [2017]	shapes/space	grad. changing	space partitioning & bijective mapping	no	no	hand	VTs	Colocation	Ergonomics		
	Feuchtmayer and Müller [2018]	hand position	grad. changing	constant rate	no	yes	hand	VTs	Colocation	Ergonomics		
	Wentzel et al. [2020]	hand position	grad. changing	non-linear interp.	no	no	tool	VTs	Colocation	Ergonomics		
	Matsuoka et al. [2002]	finger flexion to progress bar	grad. changing	linear interpolation	n.a.	n.a.	hand	AHF	n.a.	Rehabilitation		
	Burns and Brooks [2006]	hand speed	grad. changing	gain	n.a.	n.a.	hand	VTs	n.a.	Basic Research		
Threshold Investigation	Burns et al. [2006]	hand position	grad. changing	constant rate	n.a.	n.a.	hand	VTs	n.a.	Basic Research	yes	
	Lee et al. [2015]	hand position	constant	always on	n.a.	n.a.	hand	AHF	n.a.	Tracking Errors		
	Debara et al. [2018b]	shapes/space	grad. changing	geometric calculation	no	no	hand	PHF	Colocation	Basic Research		
	Debara et al. [2018a]	point	grad. changing	linear interpolation	no	no	hand	VTs	Colocation	Basic Research		
	Gonzalez et al. [2019]	point	grad. changing	linear interpolation	n.a.	n.a.	hand	PHF	Colocation	Basic Research		
	Zenner and Krüger [2019b]	point	grad. changing	geometric calculation	yes	yes	hand	VTs	Colocation	Basic Research		
	Benda et al. [2020]	hand position	constant	always on	n.a.	n.a.	tool	VTs	n.a.	Basic Research		
	Esmaeili et al. [2020]	hand speed	grad. changing	gain	n.a.	n.a.	hand	VTs	n.a.	Basic Research		
Threshold Investigation	Feick et al. [2021, 2022]	hand speed/path	grad. changing	gain	n.a.	n.a.	hand	PHF	Similarity	Basic Research	yes	
	Lebrun et al. [2021]	point	grad. changing	linear interpolation	yes	yes	hand	PHF	Colocation	Basic Research		
	Ogawa et al. [2021]	hand position	grad. changing	geometric calculation	n.a.	yes	hand	VTs	n.a.	Basic Research		
	Zenner et al. [2021b]	point	inst. changing	toggle	yes	yes	hand	VTs	Colocation	Basic Research		
	Zenner et al. [2021b] ²	point	inst. & grad. chg.	linear interp. & toggle	yes	yes	hand	VTs	Colocation	Basic Research		
	Hartfill et al. [2021] ^{4,5}	hand speed	grad. changing	gain	n.a.	yes	hand	VTs	n.a.	Rehabilitation		
	Patras et al. [2022] ¹	hand position	grad. changing	linear interpolation	yes	yes	hand	PHF	Colocation	vs. C. B. Remapping		

Table 2.3: Overview of HR research (VTs = Virtual Targets; R/V = Real-to-Virtual; 2p. = 2 points; 2f. = 2 fingers; chg. = changing; n.a. = not available/applicable). ¹ = based on Azmandian et al. [2016b]. ² = based on Cheng et al. [2017b]. ³ = based on Han et al. [2018]. ⁴ = based on Zenner and Krüger [2019b]. ⁵ = based on Esmaeili et al. [2020]. Blue entries represent research conducted in Part IV. The last column indicates techniques implemented in the *Hand Redirection Toolkit* (HaRT).

Detectability of Body Warping-Based Hand Redirection To successfully employ HR, it is important to know about the maximum degree to which the technique can be applied without negatively influencing the perception of the interaction. Offsetting the hand too much can cause semantic violations that break the illusion [Gonzalez-Franco and Lanier, 2017]. Hence, VR developers integrating HR can consider three key parameters: tolerance thresholds, self-attribution thresholds, and Detection Thresholds (DTs).

Tolerance and self-attribution thresholds describe how much an interaction can be manipulated while still being at least *tolerated* or even *self-attributed* by users. Tolerable redirection has been explored, for example, by Cheng et al. [2017b] who found redirections of up to 40° to be acceptable. The work of Debarba et al. [2018a,b] informs about the limits of self-attribution. Their results indicate that users tend to self-attribute redirected motions if they make a reaching task easier. Tolerating a certain mismatch, however, does not rule out that one is aware of its presence – highlighting the need for research on the noticeability of HR.

To this end, DTs measure the degree to which HR can be employed without users even *noticing* it. For estimating DTs, research usually resorts to experimental methods of psychophysics. Such experiments can be based on classic psychophysical methods such as the method of limits, the method of constant stimuli, or adaptive methods like staircase procedures⁴³. During a threshold experiment, participants typically experience different levels of manipulation (i.e., stimuli) and report whether they detect a stimulus or not according to some sort of Alternative Forced-Choice (AFC) task. In this thesis, we classify such tasks following the definitions of Kingdom and Prins [2016b; pp. 24 ff.]. In the literature about DTs related to HR, common tasks are 1AFC tasks in which only a single level of manipulation is presented per trial and users either answer whether any manipulation was perceived or not (known as “yes/no” 1AFC; used for example by Yang et al. [2018]) or chose the direction of the manipulation (e.g., redirection towards the right or left; known as “symmetric” 1AFC; used for example by Bergström et al. [2019]).

In 2AFC tasks, participants experience two stimuli per trial (as used for example by Lee et al. [2015]): a baseline (i.e., no redirection) and a manipulated stimulus (i.e., redirection with a certain magnitude). Participants then need to identify the manipulated stimulus among the two. A special case are tasks that encompass the assessment of stimuli pairs as either being equal or different (used for example by Matsuoka et al. [2002]), which are classified as “same/different” 1AFC tasks [Kingdom and Prins, 2016c; p. 41].

Tables 2.4 and 2.5 summarize the research on HR thresholds. As apparent from the tables, DTs can vary between techniques, with the apparatus employed, the application scenario, and even the estimation methodology [Grechkin et al., 2016].

⁴³Psychophysics on Wikipedia. <https://w.wiki/453S>

Hand Redirection Thresholds (1/2)		Experiment			Haptic Stimuli ...			Results	
Type of Threshold	Reference	Method	AFC Task	Distraction Involved	... during HR	... on arrival	Thresholds	Findings	
Detection	Matsuoka et al. [2002]	Constant Stimuli	1AFC (s./d.)	no	yes linear spring (AHF)	yes (hard stop)	distortion of visual finger position: up to 36%	Unnoticeable distortion of R/V finger position mapping possible.	
	Burns and Brooks [2006]	Adaptive Staircase	1AFC (sym.)	no	no	no	hand gain [speed up; down]: left: [0.43; 0.09], right: [0.40; 0.05] down: [0.38; 0.27], up: [0.51; 0.16] away: [0.68; -0.10], toward: [0.63; 0.45]	Large unnoticeable gains found when hand rendered as sphere.	
	Burns et al. [2006]	Limits (asc.)	n.a.	yes	no	no	hand offset: user primed: 19.1° or 19cm user not primed: 45.4° or 42cm	Indication that distraction allows for large distortions.	
	Lee et al. [2015]	Constant Stimuli	2AFC	no	yes tactile interactions (AHF)	no	fingertip offset: no haptics: up to 5.24cm with haptics: up to 6.21cm	Indication that haptics increase thresholds.	
	Abtahi and Follmer [2018]	Constant Stimuli	1AFC (yes/no)	no	yes contour & pin (PHF & AHF)	yes (constant contact)	HR while contour following: vir. contour angle: up to 49.5° phy. contour length: up to 1.9x vir. vert. pin speed: up to 3.3x	Thresholds for short HR movements with haptics. Further evidence that haptics can yield large thresholds.	
	Yang et al. [2018]	Constant Stimuli	1AFC (yes/no)	no	yes resistance (PHF)	yes (Grabbers closing)	VR Grabbers C/D ratio and mov. distortion: [0.71, 1.77] and [-1.5cm, 2.0cm]	Unnoticeable manipulation of 2-stick grasping.	
	Bergström et al. [2019]	Constant Stimuli	1AFC (sym.)	no	no	yes (grasping proxy)	vir. scaling of phy. cuboid: 3cm to [2.7cm, 4.4cm] 6cm to [5.4cm, 7.3cm] 9cm to [7.0cm, 9.2cm]	Unnoticeable manipulation of 2-finger grasping.	
	Gonzalez and Follmer [2019]	Constant Stimuli	1AFC (yes/no)	no	no	yes (touching proxy)	vert. HR (right & left hand): [-16.4°, 17.1°] & [-16.2°, 18.5°] 2 hands vert. HR (same & opp. dir.): [-19.5°, 21.4°] & [-12.3°, 14.3°]	Bimanual HR more noticeable when hands are redirected in different directions.	
	Gonzalez et al. [2019]	Constant Stimuli	1AFC (yes/no)	no	no	yes (touching proxy)	HR with PHF on arrival: left: 9.5cm +/- 1.4cm	Thresholds for desktop-scale HR towards the left with PHF.	
	Zenner and Krüger [2019b]	Constant Stimuli	1AFC (sym.)	yes	no	no	HR in mid-air: hor./vert.: [-4.5°, +4.5°] gain: [0.88, 1.07]	Conservative thresholds for desktop-scale HR.	
	Benda et al. [2020]	Constant Stimuli	1AFC (yes/no)	no	no	no	constant hand offsets: left: 10.27cm, right: 9.40cm up: 12.83cm, down: 13.37cm far: 13.25cm, close: 7.83cm	Unnoticeable constant hand offsets.	

Table 2.4: Thresholds research on HR (1/2) (s./d. = same/different; sym. = symmetric; asc. = ascending; vir. = virtual; phy. = physical; hor. = horizontal; vert. = vertical; dir. = direction; opp. = opposite). Blue entries represent contributions of Part IV.

Hand Redirection Thresholds (2/2)			Experiment			Haptic Stimuli ...		Results	
Type of Threshold	Reference	Method	AFC Task	Distraction Involved	... during HR	... on arrival	Thresholds	Findings	
Detection	Esmaeili et al. [2020]	Constant Stimuli	1AFC (yes/no)	yes	no	no	hand gain (no distraction): hor.: [0.81, 1.31]; vert.: [0.87, 1.52] depth: [0.78, 1.38] gain warping (low & high distraction): [0.80, 1.39] & [0.76, 1.43]	Unnoticeable scaled hand movements.	
	Feick et al. [2021]	Adaptive Staircase	1AFC (yes/no)	no	yes proxy & resistance (PHF)	yes (constant contact)	C/D manipulation of slider lin. translation (vir. 7cm): [0.70; 1.62] lin. translation (vir. 14cm): [0.76; 1.50] lin. stretching (vir. 7cm): [0.75; 1.54] lin. stretching (vir. 14cm): [0.80; 1.42]	Thresholds for short HR movements with haptics. Indication that increased distance & resistance decreases range of unnoticeable C/D ratio manipulation.	
	Lebrun et al. [2021]	Constant Stimuli	1AFC (sym.)	no	no	yes (touching proxy)	HR in mid-air: hor.: [-4.52°, +2.27°]	Confirmed thresholds for desktop-scale HR.	
	Ogawa et al. [2021]	Adaptive Staircase	1AFC (yes/no)	no	no	no	hand offset [left, right]: realistic hand: [-4.5cm, +4.9cm] abstract hand: [-3.5cm, +5.0cm]	Indication that with realistic virtual hands, more HR can go unnoticed compared to abstract virtual hands.	
	Zenner et al. [2021b]	Adaptive Staircase	1AFC (yes/no)	no	no	no	hand offset (BSHR+0%): right: 2.6cm, down: 3.8cm towards user: 3.3cm	Unnoticeable HR possible leveraging only blink-induced change blindness.	
	Zenner et al. [2021b]	Adaptive Staircase	1AFC (yes/no)	no	no	no	hand offset (BSHR+100%): right/towards: 4.3cm down: 5.4cm hand offset (Cheng et al. [2017b]): right: 5.8cm, down: 5.6cm towards user: 4.6cm	Adding gradual warping to instantaneous warping during blinks increases range of unnoticeable HR. Indication of unnoticeable HR achievable with Cheng et al. [2017b].	
	Hartfill et al. [2021]	Adaptive Staircase	1AFC (yes/no)	no	no	no	decelerating gain thresholds in: [0.552, 0.727]	Indication that thresholds differ with hand motion direction.	
	Feick et al. [2022]	Adaptive Staircase	1AFC (yes/no)	no	yes handheld proxy (PHF)	yes (constant contact)	accelerating gain thresholds for different grasp types and linear/circular trajectories: [1.38, 1.46] different proxy weights: [1.31, 1.36]	Indication that limited DoFs during interaction allow for more discrepancy. Indication that users' sensitivity to detect HR correlates with VR experience.	
Self-Attribution	Patras et al. [2022]	Constant Stimuli	1AFC (yes/no)	no	no	yes (touching proxy)	haptic retargeting: hor. HR: 7.9cm change blindness remapping: hor. remapping: 9.7cm	Indication that change blindness remapping is less noticeable than HR.	
	Debarba et al. [2018b]	Constant Stimuli	1AFC (yes/no)	no	no	yes (touching proxy)	ID change ratio (tapping task) helping threshold: -0.39 hindering threshold: +0.28	Users tend to self-attribute manipulated movements that make a pointing task easier.	
	Debarba et al. [2018a]	Adaptive Staircase	1AFC (yes/no)	no	yes handheld proxy (PHF)	yes (constant contact)	hand gain [speed up; down]: left: [0.86; 0.13], right: [0.84; 0.21] down: [0.90; 0.27], up: [0.65; 0.18]	Users tend to self-attribute manipulated movements that make a reaching task easier.	
Tolerance	Cheng et al. [2017b]	Rating Scale	n.a.	no	no	yes (touching proxy)	HR with PHF on arrival: ~40°	Users might tolerate HR even above detection thresholds.	

Table 2.5: Thresholds research on HR (2/2) (s./d. = same/different; sym. = symmetric; asc. = ascending; vir. = virtual; phy. = physical; hor. = horizontal; vert. = vertical; dir. = direction; opp. = opposite). Blue entries represent contributions of Part IV.

Notable examples are works investigating the detectability of hand position and speed discrepancies [Burns et al., 2006; Burns and Brooks, 2006]. Moreover, Lee et al. [2015] found haptic feedback at the fingertip to enlarge the range of unnoticeable discrepancy between real and virtual fingertip positions. Similarly, Abtahi and Follmer [2018] studied the detectability of HR-based illusions when haptic signals are present at the hand and Feick et al. [2022] explored the impact of grasping type, movement trajectory, and object mass on the detectability of HR when handling proxies. The extent to which HR goes unnoticed when interacting with a manipulatable part of a prop (e.g., a slider) has also been subject of our own recent work: For gain-based HR during interactions involving linear translation, Feick et al. [2021] found movement distance and the presence of force feedback to impact DTs, with shorter movements allowing for greater manipulations and resisting forces increasing users' sensitivity to detect manipulations.

Finally, several investigations followed up on our work presented in Part IV, studying DTs for HR in mid-air [Zenner and Krüger, 2019b]. For example, Esmaeili et al. [2020] and Hartfill et al. [2021] investigated the detectability of scaled hand movements, while Benda et al. [2020] studied the detection of constant hand offsets. In particular, Esmaeili et al. [2020] and Benda et al. [2020] considered detectability in less conservative scenarios. Their results indicate thresholds to differ with offset direction. Gonzalez and Follmer [2019] found that when redirecting two hands at the same time and with the same magnitude (bimanual retargeting), redirecting them in opposite directions leads to the redirection being more noticeable compared to redirecting them in the same direction. Moreover, Lebrun et al. [2021] recently confirmed our results presented in chapter 8 concerning DTs for horizontal HR and elaborated a hand trajectory model for desktop-scale HR based on a Bézier curve with four control points. Furthermore, Ogawa et al. [2021] showed that for hand offsets towards the left, rendering a realistic human avatar hand can lower users' sensitivity to detecting HR by 31.3% compared to an abstract hand representation like a sphere. Their results suggest this likely being an effect of a greater weighting of visual information during multisensory integration when realistic avatars are used.

Haptic Retargeting A first way to utilize hand redirection for proxy-based haptics is to solve the challenge of *Colocation* when reaching for a prop. Specifically, hand redirection can be employed to redirect the user's real hand towards a haptic proxy when it is not perfectly colocated with its virtual counterpart. One commonly refers to this strategy as *redirected reaching* or haptic retargeting.

The concept of haptic retargeting was introduced by Azmandian et al. [2016b] and is based on the fundamental ideas underlying redirected touching introduced in the next section [Kohli, 2013a]. Azmandian et al. [2016b] showed that the technique enables a single physical proxy to represent multiple, dislocated virtual objects, effectively tackling the *Colocation* challenge for objects within the user's reach. Moreover, the results of Azmandian et al. [2016b] show that any type of hand redirection is suitable for realizing haptic retargeting.



Figure 2.36: Two examples of haptic retargeting. Left: HR enabling a single proxy to provide haptic feedback for three displaced virtual cubes (extracted from [Azmandian et al., 2016b]; © 2016 ACM). Right: The real hand and proxy (blue) are offset from their virtual counterparts (colored) (extracted from [Han et al., 2018]; © 2018 IEEE).

To establish *Colocation* through world warping-based HR, for example, Azmandian et al. [2016b] proposed to translate or rotate the IVE until the virtual object's location coincides with that of the proxy. Inspired by techniques employed for RDW, such IVE manipulations can be performed in an unnoticeable way, e.g., during head movements. Besides relocating the complete IVE, world warping can also be applied locally. An example for this is the interface warp technique by Matthews et al. [2019] which maps multiple virtual buttons onto a single physical button by shifting only the virtual UI.

Most common and most relevant for this thesis is the approach of body warping-based HR to establish *Colocation*. Here, the user's hand is redirected when reaching for a virtual object to ensure that the real hand ends up touching the corresponding but spatially offset proxy instead of thin air as illustrated in Figure 2.36. Besides a physical object, such a proxy can also be another part of the user's own body [Fang and Harrison, 2021]. A variety of algorithms have been proposed for HR. Most gradually apply the offset from proxy to virtual counterpart (\vec{T} in Figure 2.35) to the virtual hand rendering while the hand approaches the object. As the user intends to reach the virtual object with their virtual hand, this offset leads to a *Colocation* of the real hand with the physical proxy. Table 2.3 provides a comprehensive overview of different techniques.

A common implementation for haptic retargeting based on HR is the algorithm by Azmandian et al. [2016b], which is effective and low in complexity. Referring to the example in Figure 2.35, the algorithm interpolates the offset vector ($\vec{T} = P_V - P_P$) between proxy ($P_P \in \mathbb{R}^3$) and virtual counterpart ($P_V \in \mathbb{R}^3$) to displace the virtual hand model. This interpolation is subject to an interpolant given by the hand's progress from the origin of the redirection (where no offset is applied) towards the physical target P_P . Cheng et al. [2017b] later extended this algorithm in a way that allows for on-the-fly switching of the redirection target. Also Cheng et al.'s algorithm trades off simplicity and effectiveness very well and is one of the most frequently used HR techniques today. We will investigate and extend Cheng et al.'s HR technique in Part IV of this thesis.

To ensure that the hand is redirected towards the correct location, HR techniques that only map a single physical point to an offset virtual point require information about which virtual object the user is going to reach for next. This can be solved with scripted interactions following a fixed sequence [Matthews and Smith, 2019]. However, since scripted interactions are often too restrictive, researchers introduced various strategies for user selection and automatic prediction of reaching targets. Matthews and Smith [2019], for example, proposed the selection of desired virtual targets through head gaze. Azmandian et al. [2016b] automatically predicts the next targets based on hand velocity and the direction of potential targets relative to the hand. Similarly, Cheng et al. [2017b] employed an analysis of hand motion and eye gaze to infer reaching targets. Another approach is to use neural networks trained to predict reach targets based on the beginning of the user’s reach trajectory [Clarence et al., 2021]. Moreover, apart from HR algorithms that map individual points, also space warping approaches have been used for haptic retargeting [Carvalheiro et al., 2016].

Furthermore, in unwarped settings, hand speed profiles and hand arrival times of an ongoing reach can be predicted with mathematical models like the Minimum-Jerk model. Recently, Gonzalez et al. [2019, 2020] investigated how well this model holds when HR is applied. Their findings indicate that for redirected reaches, the Minimum-Jerk model’s prediction quality and fit is decreased when HR above DTs is applied; emphasizing the importance of identifying DTs for HR.

Finally, Azmandian et al. [2016b] also explored hybrid warping-based HR, i.e., the combination of world and body warping, which can reduce the required amount of body warping-based HR by shifting the virtual target towards the proxy. Specifically, they proposed performing as much world warping as users’ head rotations allow for, while bridging remaining offsets with body warping. Similarly, Matthews et al. [2019] combined body warping and their proposed interface warp according to a fixed warp proportion parameter. Azmandian et al. [2016b] found hybrid warping to yield the highest reported presence and realism scores when compared to body warping and world warping. Moreover, Matthews et al. [2019] found hybrid warping to yield improved task times and reduced error rates when remapping interfaces compared to body warping.

HR can be used, for example, to reduce the number of haptic proxies required for successful VR training [Geslain et al., 2021]. Yet, besides haptics, HR can also be employed when interacting with purely virtual objects and 3D UIs, for example, by redirecting the user’s reach to improve ergonomics. In this context, Montano Murillo et al. [2017] proposed the *Erg-O* technique that employs HR for mapping interactive virtual elements in the user’s surroundings to physical locations that are ergonomically more comfortable to interact at. Similarly, Feuchtner and Müller [2018] proposed to redirect the real hands during prolonged overhead interactions in VR towards lower and more comfortable locations with their *Ownershift* technique. Besides that, Wentzel et al. [2020] recently explored how non-linear HR bound to the user’s natural reach distance can improve ergonomics in VR while maintaining body ownership.

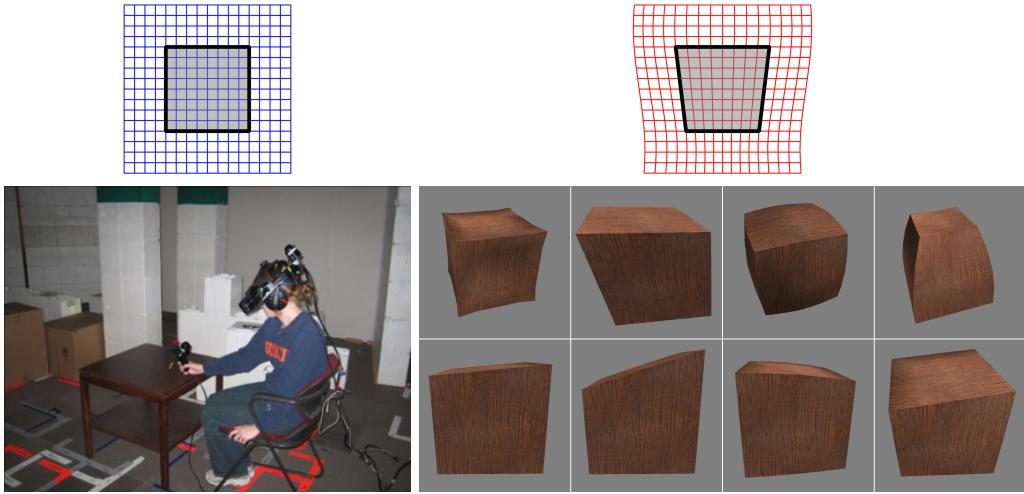


Figure 2.37: An example of redirected touching. Left: The proxy table touched by the user. Right: The warped virtual space in which both the table’s geometry and the user’s virtual hand is displaced (extracted from [Kohli, 2013a]; © 2013 Kohli).

Redirected Touching The second relevant real-time virtual strategy based on hand redirection is redirected touching [Kohli, 2013a]. Here, hand redirection is employed not to reach a physical proxy, but while touching and exploring a prop that the user is already in contact with. Similar to pseudo-haptics, redirected touching also focuses on solving the *Similarity* challenge. However, instead of simulating properties like stiffness, friction, or mass, redirected touching is specialized on *contour following* [Lederman and Klatzky, 2009] and conveys multiple different shapes with only a single proxy fixed in shape. Specifically, the technique is based on visual dominance and redirects the user’s hand so that it physically follows the shape of the haptic prop while virtually following the discrepant shape of a virtual object [Ban et al., 2012].

Conceptually, one can think of this approach as a distortion of the space around the haptic prop as illustrated in Figure 2.37 [Kohli, 2013a]. This distortion defines how much the virtual hand is displaced from its real-world position while exploring the surface of the prop and the space around it. The resulting HR ensures the temporal synchronization of hand-object collisions in the real and virtual world, and that the user’s virtual finger follows the contours of the virtual object while the real finger traces the proxy’s edges. Technically such space warping is realized by defining offsets at specific points in the IVE and interpolating them to determine the offsets at arbitrary hand locations (e.g., using thin plate splines⁴⁴ [Kohli, 2010; Kohli et al., 2012] or other interpolation methods [Spillmann et al., 2013]). Zhao and Follmer [2018] recently extended the approach to support arbitrary geometries leveraging functional optimization to determine offsets.

Notable work on redirected touching was conducted by Kohli, who originally

⁴⁴Thin plate spline on Wikipedia. <https://w.wiki/4542>

introduced the technique. To demonstrate the concept for the first time, Kohli [2009] mapped the flat surface of a physical table to various warped virtual surfaces as shown in Figure 2.37. A little later, Kohli [2010] presented a first calibration method to set up a warped space and extended the approach to work with Leap Motion³⁵ hand tracking [Kohli, 2013b]. Moreover, Kohli et al. [2012] formally investigated the impact of warped interactions on task performance. In a Fitts's law task⁴⁵, the authors showed that interacting with warped virtual objects yields comparable performance to interacting with non-warped virtual objects. As Kohli [2013a] concludes, "*evidence suggests that after adaptation, users can perform tasks in a discrepant VE generally no worse than in a one-to-one VE*".

The great potential of redirected touching for establishing *Similarity* was also demonstrated by Ban et al. [2012], who found that users can be tricked into perceiving the edges on a warped object to be oriented differently than the edges on its proxy. A little later, Ban et al. [2014] proposed a system that combines redirected touching with a DPHF proxy, which can dynamically add and remove a haptic edge on its surface to convey many differently-shaped virtual objects.

Furthermore, Yang et al. [2018] and Bergström et al. [2019] explored *resized grasping*, a related concept concerned with the redirection of two reference points to simulate grasping differently-sized virtual objects. Yang et al. [2018] presented an approach compatible with passive haptic grabbing tools like chopsticks, while Bergström et al. [2019] focused on grasping with thumb and index finger.

In terms of application areas, redirected touching has been applied in a knee arthroscopy simulator to map different bone shapes and knee joint behaviors onto a single physical proxy [Spillmann et al., 2013]. Moreover, Strandholt et al. [2020] proposed leveraging redirected touching in an application related to carpentry for implementing redirected hammering, screwing, and sawing.

Additional Remarks Redirected touching and haptic retargeting are closely related concepts. Both have in common that HR plays a starring role in their most common implementations. Yet, the two terms, as well as the terms *hand redirection* and *body warping*, are frequently used in inconsistent ways in the related literature. In fact, up to now, the conceptual relationship between these terms and concepts still remains to be laid out in a concise and coherent way.

It is with Figure 2.38 that we attempt to provide an overview of how these terms and concepts relate to each other. With the illustration in Figure 2.38 and our understanding of these terms introduced in this thesis, we aim to facilitate a common language for the research community.

In summary, we define hand redirection as a technique that allows the VR system to control the user's real hand movement. This technique of hand redirection can be implemented in three different ways, either through body warping (referred to in this thesis as HR), world warping, or hybrid warping. Independent of how

⁴⁵*Fitts's law* on Wikipedia. <https://w.wiki/4uCa>

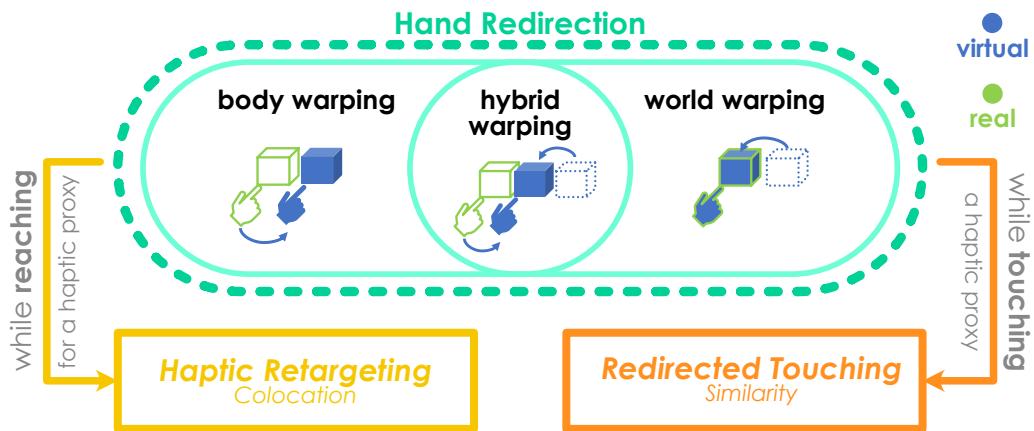


Figure 2.38: Relation between *hand redirection*, *haptic retargeting*, and *redirected touching*.

it is implemented, hand redirection can be employed for redirecting the real hand either during interaction with purely virtual elements (e.g., to improve the ergonomics of virtual interfaces as in [Montano Murillo et al., 2017; Feuchtnner and Müller, 2018]), or during proxy-based haptic interaction. In the latter case, one can distinguish two different use cases for hand redirection. When employed while the user is reaching for a haptic proxy to establish *Colocation*, we refer to it as haptic retargeting (e.g., as in [Azmandian et al., 2016b]). Alternatively, when employed while haptically exploring the shape of an object the user is already in contact with to convey *Similarity*, we refer to it as redirected touching.

In addition, we like to mention that solving the *Similarity* challenge in respect to shape through redirected touching can, on a conceptual level, also be regarded as solving infinitely many *Colocation* challenges, yielding a view of shape-*Similarity* as a *sum of Colocation problems*.

Finally, as in common interaction scenarios users first reach for an object (*Colocation* phase) before exploring it (*Similarity* phase), applications will likely first employ haptic retargeting followed by redirected touching as outlined in Figure 2.39. To realize this, body warping-based HR appears especially promising as it only affects the virtual hand (and not the entire IVE), is realizable with simple algorithms, and thus is easy to integrate in VR applications. For this reason, our research on hand redirection focuses particularly on body warping-based HR.



Figure 2.39: Order in which haptic retargeting and redirected touching are applied.



Figure 2.40: An example of change blindness haptic remapping. Left: Two virtual objects represented by a single proxy. Center: The proxy. Right: Once the target has been identified and the user looks elsewhere, the *Mise-Unseen* system remaps the target onto the proxy. All images extracted from [Marwecki et al., 2019]; ©2019 ACM.

Change Blindness Haptic Remapping

Besides pseudo-haptics and hand redirection, a third real-time virtual strategy is *change blindness haptic remapping* [Lohse et al., 2019], which establishes *Colocation* for objects in- and out-of-reach. As shown in Figure 2.40, change blindness haptic remapping translates, rotates, or blends virtual objects to the location of their proxies. Relocation takes place when users are likely to miss it due to change blindness, e.g., when their view of the object is occluded, or when it is outside the FOV [Lohse et al., 2019]. Marwecki et al. [2019] also proposed unnoticeable remapping inside the FOV by performing changes outside foveal vision, utilizing visual masking, and taking into account spatial memory. Change blindness haptic remapping can thus be considered a special form of world warping, timed specifically to take advantage of change blindness [Lohse et al., 2019].

Similar to HR, change blindness haptic remapping can provide haptic feedback for multiple objects with only a single proxy but requires knowledge about which object the user is going to interact with next [Lohse et al., 2019]. Once the next target is known, opportunities for subtle remapping must be provided either through interactions in the IVE or specific distractors. With the *Mise-Unseen* system, Marwecki et al. [2019] proposed an approach to prevent anticipation, observation, and recall of IVE-changes inside the user’s FOV. Based on eye tracking and attention models, their system monitors the user’s present and past visual attention, as well as the probability that users will notice a specific change in order to identify opportunities for change blindness haptic remapping.

The studies by Lohse et al. [2019] and Patras et al. [2022] demonstrate the potential of change blindness haptic remapping for establishing *Colocation*. Patras et al. [2022], for example, found indication that compared to HR-based haptic retargeting, change blindness remapping can be preferable in terms of virtual embodiment and can unnoticeably establish *Colocation* for greater offsets (i.e., be less noticeable). However, Lohse et al. [2019] also highlight that the technique’s utility is limited in scenarios that allow for free exploration. A reason for this is that there might not always be a suitable opportunity for remapping when the user engages a new object. With the techniques introduced by Marwecki et al. [2019], this limitation can be alleviated to a certain extent.

Redirected Walking

Concluding our review of real-time virtual strategies, *Redirected Walking* (RDW) represents a solution for establishing *Colocation* of objects outside the user's reach. The goal of RDW is to enable users to explore IVEs larger than the physical tracking space by means of physical walking. We briefly summarize the most prominent approaches to RDW and then review how RDW can be applied to establish proxy *Colocation*. Comprehensive reviews of RDW research can be found in Langbehn and Steinicke [2018] and Nilsson et al. [2018].

Redirection Techniques Two types of RDW techniques exist today [Nilsson et al., 2018]: (1) RDW techniques based on redirection gains, and (2) RDW techniques based on manipulations of the IVE layout. Some techniques are overt by design while others are subtle and can go unnoticed.

In their pioneering work on RDW, Razzaque et al. [2001] proposed to rotate the IVE about the user's head so as to lead the user along a physical arc while walking straight inside the IVE (similar to the illustration on the left in Figure 2.41). This concept was generalized to *gain-based RDW* in later research. In parallel to body warping for haptic retargeting, gain-based RDW manipulates how physical rotations and translations of the user's head are mapped to the IVE. Rotation gains that scale the user's head rotations, translation gains, curvature gains, and bending gains have been proposed to make users follow a physical trajectory that differs from the virtual path traversed in the IVE [Nilsson et al., 2018; and references therein]. To apply manipulations without users noticing it, perceptual thresholds for RDW have been established through psychophysical studies similar to research conducted on the detectability of HR. Steinicke et al. [2010b], for example, found that rotations can be scaled within [0.67, 1.24] to make users physically turn more or less than in the IVE, respectively. Similarly, scaling factors within [0.86, 1.26] for translation gains were found to go unnoticed and a circle with a radius of at least 22m was found to be required for infinite straight virtual walking with curvature gains. Moreover, some RDW algorithms take advantage of human eye blinks [Langbehn et al., 2018b] – an approach we extend to HR in Part IV. Still others leverage saccades to hide changes [Sun et al., 2018]. Also for these cases, perceptual studies inform about the magnitude of view rotations and translations that can go unnoticed. During blinks, translations of up to 4cm – 9cm as well as rotations of up to 2° – 5° when standing still [Langbehn et al., 2018b] and 9.1° when walking (compared to 2.4° when not blinking) [Nguyen and Kunz, 2018] were found to go unnoticed. Bolte and Lappe [2015] found translations of up to 50cm and rotations of up to 5° to remain undetectable during saccades. Both Langbehn and Steinicke [2018] and Nilsson et al. [2018] provide an overview of further DTs identified by previous research.

Besides RDW algorithms, additional *resetting controllers* are required when the tracking space cannot fit a full physical circle with unnoticeable RDW curvature [Nilsson et al., 2018]. Such controllers lead users to reorient towards a direction



Figure 2.41: Examples of Redirected Walking (RDW). Left: Illustration of gain-based RDW (extracted from [Steinicke et al., 2008]; © 2008 IEEE). Center & Right: Relocation of a door and corridor during change blindness redirection. The center image shows the IVE when entering the room, the right image shows the IVE after it was changed behind the user’s back (extracted from [Suma et al., 2011a]; © 2011 IEEE).

that is free to continue walking when they reach the tracking boundary. Razzaque et al. [2001], for example, used verbal instructions inside the IVE to have users look around the scene while scaling their head rotations. Other controllers have been proposed by Williams et al. [2007] with the *Freeze-Backup*, *Freeze-Turn*, and the *2:1-Turn* techniques, and by Nguyen and Kunz [2018] with the *to-corner* reset. The *Redirected Walking Toolkit* [Azmandian et al., 2016a] offers implementations of common gain-based RDW techniques (also considering DTs) and resetting controllers, as well as analysis features to support RDW research.

Related to resetting controllers are *relocation techniques*, which also enable users to explore IVEs by means of real walking [Nilsson et al., 2018]. In contrast to conventional RDW, however, relocation techniques transport users to different locations in the IVE without physical walking, for example through virtual teleportation [Bowman et al., 1997]. Once at the virtual destination, users then explore their immediate surroundings by means of real walking. Relocation techniques are widely used in commercial VR applications designed for domestic VR setups with limited tracking space [Frommel et al., 2017; Simeone et al., 2020].

The second class of RDW techniques modifies the IVE instead of the R/V mapping of the user’s view. Hence, it is the equivalent to world warping for haptic retargeting and does not introduce visual-proprioceptive conflicts. Approaches like *change blindness redirection* [Suma et al., 2011a] and *impossible spaces* [Suma et al., 2012] compress large IVEs into smaller physical spaces by modifying the virtual architecture as the user walks through it, or by implementing self-overlapping spaces, respectively. Change blindness redirection, for example, changes the location of virtual doors and corridors as shown in Figure 2.41 to ensure that users remain inside the tracking space when walking from one virtual room to the next [Suma et al., 2010]. Even significant scene changes can go unnoticed when timed appropriately (e.g., happening behind the user’s back) due to change blindness [Suma et al., 2011a]. For impossible spaces, i.e., virtual rooms that are “*bigger on the inside*” [Suma et al., 2012], layouts may overlap by up to 56% without users noticing. Other techniques overtly modify the IVE. In our own previous work, for example, we introduced the *Space Bender* technique

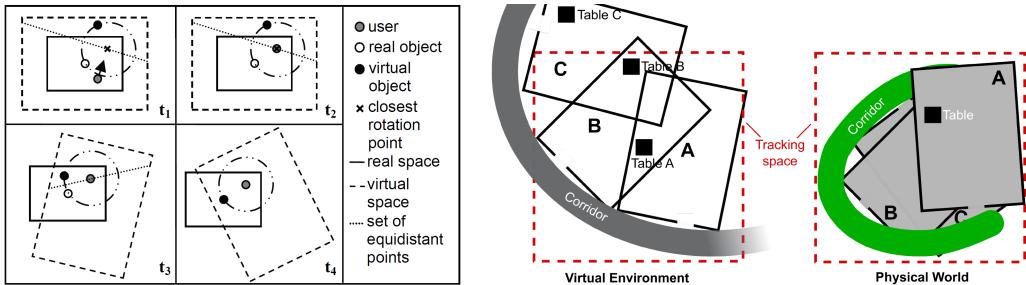


Figure 2.42: Examples of RDW for proxy *Colocation*. Left: Rotating the IVE to align virtual object and proxy (extracted from [Kohli et al., 2005]; © 2005 HIT Lab NZ, University of Canterbury). Right: Combining bending gains and impossible spaces to map three virtual tables to a single proxy (extracted from [Langbehn et al., 2018a]; © 2018 IEEE).

that keeps the user walking inside the tracking space by visually bending virtual paths when the user approaches the tracking boundary [Simeone et al., 2020].

Finally, the two approaches to RDW can also be combined, similar to hybrid warping for haptic retargeting. Suma et al. [2012] presented a system where users can walk through a virtual village of $22m \times 44m$ leveraging a physical space of only $9m \times 9m$. This compression is achieved by combining curvature gains, translation gains, and impossible spaces. Similarly, Langbehn et al. [2018a] took advantage of bending gains and self-overlapping spaces to compress three virtual rooms into a physical space of $4m \times 4m$.

Redirected Walking for Enhanced Colocation With RDW, users follow physical paths that differ from the paths traversed in the IVE, just like hand redirection leads the hands along different trajectories. Hence, RDW *reuses* physical space for different parts of a scene and enables the reuse of proxies to establish *Colocation*.

Kohli et al. [2005] first demonstrated such a solution in a system where users could touch multiple virtual pedestals distributed across the IVE. The scene could be explored by real walking, with only a single symmetric prop providing PHF for all virtual pedestals. To repeatedly redirect users to this prop, a story was guiding users to locations that were equidistant to both the virtual and the physical pedestals. At these locations, head rotations were incentivized during which the IVE rotated about the user's location to align the virtual pedestal with its proxy as illustrated on the left in Figure 2.42. Following-up on this work, Steinicke et al. [2008] presented a general approach for enabling proxy *Colocation* through RDW. The authors consider target prediction, virtual and real object locations, approach angles, and obstacle avoidance to derive suitable paths for RDW that lead the user to the virtual object while appropriately approaching the proxy in the real environment. Langbehn et al. [2018a] further demonstrated how a single proxy table can provide PHF for tables in three different virtual rooms as sketched on the right in Figure 2.42 by combining bending gains and impossible spaces. Additionally, Sait et al. [2018] introduced two methods specialized for

domestic VR setups, which align a proxy with different virtual objects when the user remains in a fixed physical position (e.g., seated at a table on which the proxy resides). Furthermore, RDW can also ensure *Colocation* of PHF walking surfaces. Suma et al. [2011b], for example, leveraged change blindness redirection in combination with a physical gravel pathway to simulate walking from one virtual building to the next via an arbitrarily long virtual gravel corridor.

Finally, while RDW can enhance proxy-based haptics, proxies can also benefit RDW, as demonstrated by Matsumoto et al. [2019]. In their *Unlimited Corridor* installation, users perceive PHF from a physical handrail during RDW. The authors could show that the feedback provided by the handrails leads to users perceiving their walking path as more straight compared to RDW without PHF.

2.6.3 Conclusion

Virtual techniques are valuable as they are realized solely through software. Yet, their main drawbacks are perceptual limits beyond which manipulations become noticeable and potentially detrimental to the VR experience.

When certain proxies (e.g., a chair) or proxy properties (e.g., a surface with a specific material) are available, offline virtual techniques such as Substitutional Reality (SR) or the use of visual-haptic incongruence can yield compelling haptic experiences leveraging visual dominance. Most relevant for this thesis, however, is how *Similarity* and *Colocation* can be achieved through real-time virtual techniques. While pseudo-haptics and Redirected Walking (RDW) are suitable for scenarios in which the user already is in contact with an object or the object is out-of-reach, respectively, this thesis is concerned with the case of establishing contact with an object in reach distance; specifically, the strategy of Body Warping-based Hand Redirection (HR). HR is a versatile approach that can establish both *Similarity* (when used for redirected touching) and *Colocation* (when used for haptic retargeting). However, it seems crucial to apply HR in an unnoticeable way so as to create the best user experience possible, especially in light of research that indicates visual-haptic illusions to impact task performance and accuracy of an interaction [Feick et al., 2021].

2.7 Combined Approaches to Similarity and Colocation

The orthogonal nature of physical and virtual approaches allows for *combined strategies*, which promise to further enhance *Similarity* and *Colocation*. Offline techniques are frequently interwoven, such as when constructed sets rely on visual-haptic incongruence to convey an IVE. Yet, only few works studied techniques that mix real-time physical and virtual approaches.

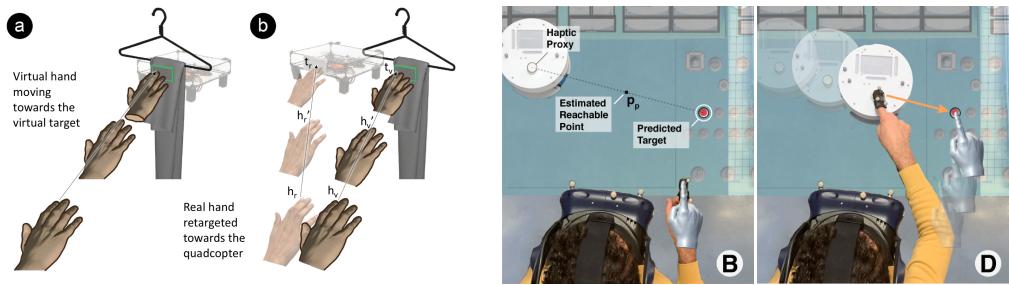


Figure 2.43: Two combinations of ETHF and haptic retargeting to ensure *Colocation*. Left: HR being used to compensate for limited drone accuracy (extracted from [Abtahi et al., 2019]; © 2019 Abtahi and co-authors; Unchanged; License: [CC BY-NC-ND 4.0](https://creativecommons.org/licenses/by-nd/4.0/)). Right: Redirecting the user’s real hand to a location reachable by the robotic proxy to ensure a proper encounter (extracted from [Gonzalez et al., 2020]; © 2020 ACM).

2.7.1 Encountered-Type Haptics & Haptic Retargeting

Recent research revealed considerable potential for the combination of physical and virtual real-time techniques. Haptic retargeting, for example, promises to relax workspace constraints, i.e., *Colocation* problems, inherent to ETHF.

Abtahi et al. [2019] achieved improved *Colocation* in their quadcopter-based ETHF system. Here, HR compensated for the limited accuracy of the drone encountered by users as shown in Figure 2.43. Moreover, Lee et al. [2020] proposed a world warping algorithm to solve *Colocation* issues in an ETHF system based on a grounded robotic arm. Furthermore, Gonzalez et al. [2020] combined a roving proxy with HR as shown in Figure 2.43, and proposed the *REACH+* framework to account for workspace limitations and for limited robot speeds. Their empirical results indicate this method to enhance the performance of ETHF involving low-speed robots by boosting on-time arrival rates by up to 25%.

2.7.2 Dynamic Passive Haptics & Haptic Retargeting

The combination of DPHF and HR has been left largely understudied. Yet, two works started to investigate the combination of shape-changing props and HR. Ban et al. [2014] demonstrated the potential of combining a cylinder with a dynamically appearing haptic edge on its surface for solving the *Similarity* challenge when rendering shapes. Leveraging a combined technique, here, the authors could convey various virtual shapes consisting of either one or multiple surfaces using only a single dynamic prop and body warping. Moreover, Gonzalez et al. [2021b] recently investigated how the shape-changing DPHF prop *Adaptic* can provide haptic feedback for several dislocated virtual objects of various shapes, using HR to ensure *Colocation*.

2.7.3 Conclusion

Studying combinations of real-time physical and virtual techniques for enhanced proxy-based haptics is a novel field of research. Up to now, only very little research has been conducted to systematically investigate the capacity of combining DPHF with HR and no research yet studied the combination of HR with the specific DPHF concepts proposed in this thesis. Such a conceptual partnership, however, appears promising in light of the benefits achievable when combining ETHF with HR – hence, motivating our work in Part V.

2.8 Research Methodology

To conclude this chapter, we outline the methods from the fields of HCI, VR, and psychophysics used to evaluate the systems and techniques of this thesis.

Apart from two *thought experiments* in chapter 12, the majority of our evaluations was realized as *user studies*, generally collecting both *quantitative* and *qualitative* data. Most evaluations were *controlled laboratory experiments*, which are suitable for early-stage research and controlled comparisons (e.g., when assessing the haptic feedback of our hardware prototypes in Part II and Part III). Additionally, lab studies facilitate the collection of quantitative and objective data (e.g., when assessing perceptual DTs in Part IV and Part V) [Dix et al., 2003; pp. 358 f.]. An exception is the evaluation of the Hand Redirection Toolkit (HaRT) in chapter 10, which was conducted as a *remote user study*. Here, expert users solved prepared tasks in their own working environment while the experimenter carried out a remote *observation*. Table 2.6 summarizes the assessment methods used throughout our work – ranging from task-related measurements to *query techniques* like questionnaires and interviews [Dix et al., 2003; pp. 348 ff.]. Each evaluation assessed the aspects that were most important to answer the respective research questions and the corresponding chapters outline the employed methods in detail.

For the implementation of some experiments, we utilized the *Unity Experiment Framework (UXF)*⁴⁶ by Brookes et al. [2020]. Moreover, some experiments acquired questionnaire data with the *VRQuestionnaireToolkit*⁴⁷ by Feick et al. [2020b].

2.8.1 Theoretical Aspects

The approach we propose in Part V was first evaluated on a purely conceptual level, before it was implemented. For this, we conducted two thought experiments⁴⁸. Thought experiments are conducted only theoretically by systematically *thinking through* a scenario based on established causalities, such as physical laws.

⁴⁶Unity Experiment Framework (UXF) on GitHub. <https://git.io/JzlqR>

⁴⁷VRQuestionnaireToolkit on GitHub. <https://git.io/Jzlq1>

⁴⁸Thought experiment on Wikipedia. <https://w.wiki/46cG>

Category	Aspect	Assessment Method
Theoretical	Concept	Thought Experiments
Technical	Prototype Specification	Physical Measurements
VR	Presence	Slater-Usoh-Steed Presence Questionnaire (SUS)
	User Experience	User Experience Questionnaire (UEQ) (UEQ-S)
	Sickness	Simulator Sickness Questionnaire (SSQ)
Task	Task Performance	Interaction Timing, Task-Specific Metrics
	Task Load	NASA Task Load Index (NASA TLX)
Perception	Haptic Realism	Custom Questionnaires
	Detection Thresholds	Psychophysical Experiments
Other	Demographic	Demographic Questionnaire
	Experiment-Specific	Custom Questionnaires, Observations, Interviews
	Unanticipated	

Table 2.6: Overview of evaluated aspects and assessment methods.

The advantage of thought experiments is that ideas can be tested for validity and their potential without the efforts and risks involved in actually implementing them. Often, thought experiments are precursors of actual experiments. We applied the same strategy in Part V: encouraged by the outcomes of our thought experiments, we subsequently conducted a user experiment to validate that the effects predicted in the thought experiments also occur in actual implementations.

2.8.2 Technical Aspects

To provide details about the devices built for this thesis, we documented their most important technical specifications. For this, characteristics such as the size, weight, or actuation speed of the prototypes were assessed by means of *physical measurements*. This approach is common in the field of VR haptics and HCI [Benko et al., 2016; Choi et al., 2018], and allows for comparisons with related devices proposed in the literature.

2.8.3 Virtual Reality-Related Aspects

Following established research practices, we assessed VR-related aspects like presence, User Experience (UX), and sickness through standard questionnaires.

Presence Several questionnaires exist to measure presence [Van Baren and IJsselsteijn, 2004]. For this thesis, we chose to measure presence with the *Slater-Usoh-Steed Presence Questionnaire (SUS)* [Slater et al., 1994, 1995; Usoh et al., 1999], as it is comparably short and one of the most widely used presence questionnaires. The SUS consists of six questions, each answered on a scale from 1 to 7. Two ways of computing a participant's presence score are common in the VR research

community: the *SUS count* (used, for example, in [Lok et al., 2003]) and the *SUS mean* (used, for example, in [Suma et al., 2011a]). While the SUS count is computed by counting the number of questions to which the participant responded with a 6 or a 7, the SUS mean is computed by averaging the responses to all six questions. Consequently, the SUS count results in a value in [0, 6] and the SUS mean in a value in [1, 7] with higher values indicating higher presence.

User Experience As for presence, also several questionnaires exist to measure User Experience (UX). For this thesis, we administered the *User Experience Questionnaire (UEQ)*⁴⁹ [Laugwitz et al., 2008] and its short version (*UEQ-S*) [Schrepp et al., 2017] as it is a well-established tool to measure both usability and UX. The full-length version of the UEQ consists of 26 questionnaire items [Schrepp, 2019]. Each item presents a pair of contrasting attributes the agreement to which is to be rated on a 7-point scale (e.g., 3 = attractive to -3 = unattractive). From this data, six subscales are computed, quantifying a system's *attractiveness*, *perspicuity*, *efficiency*, *dependability*, *stimulation*, and *novelty*. In addition, a score for *pragmatic*/“*goal-directed*” *quality* can be computed from the perspicuity, efficiency, and dependability results. A score for a system's *hedonic*/“*not goal-directed*” *quality* can likewise be derived from the stimulation and novelty ratings. All scales are in the range [-3, +3] with higher values indicating more positive assessments. We used the UEQ to evaluate the system presented in chapter 3.

The short version of the questionnaire (UEQ-S) allows for an even more efficient assessment and consists of only eight items. The UEQ-S, however, does not evaluate the six detailed subscales of the full-length version. Instead, it is used to derive only an overall scale and values for pragmatic and hedonic quality [Schrepp et al., 2017]. We employed the UEQ-S to evaluate the HaRT in chapter 10.

Sickness To assess potential negative effects of our VR systems, we administered the *Simulator Sickness Questionnaire (SSQ)* [Kennedy et al., 1993]. The SSQ is one of the most widely used questionnaires in research for measuring cybersickness issues. Sickness scores are collected post-exposure by having participants rate the severity of 16 symptoms on a 4-point scale (0 = none, 1 = slight, 2 = moderate, 3 = severe). Using the evaluation scheme outlined by Kennedy et al. [1993], four indicators are computed from the participant's ratings: a *nausea* subscore (range [0, 200.34]), an *oculomotor* subscore (range [0, 159.18]), a *disorientation* subscore (range [0, 292.32]), and a *total score* (range [0, 235.62]). The total score serves as an overall indicator for sickness problems with higher values indicating greater severity. The subscores can help identify the technical aspects of the simulator responsible for inducing sickness [Kennedy et al., 1993; Hösch, 2018].

⁴⁹User Experience Questionnaire. <https://bit.ly/3nPNEwT>

2.8.4 Task-Related Aspects

The experiments in chapter 3 and chapter 10 involved predefined tasks to be performed by the user. We evaluated them by measuring task-related aspects.

Task Performance

Where applicable, the performance of users in solving the experimental tasks was measured. *Efficiency* was assessed by taking the time that it took participants to complete a task. *Effectiveness* was measured by experiment-specific metrics, such as percentage-correct or rating scales. The details of the metrics and evaluations employed are described for each case separately in the respective chapter.

Task Load

In chapter 3, we also assessed the workload experienced when solving the experimental task with the *NASA Task Load Index (NASA TLX)*⁵⁰ [Hart and Staveland, 1988]. The NASA TLX requires participants to first rate six different aspects of workload (*Mental Demand*, *Physical Demand*, *Temporal Demand*, *Performance*, *Effort*, and *Frustration*) each on a scale from 0 to 100 in 5-point increments. In a second step, users are presented all 15 pairs of workload aspects and asked for each pair to select the aspect that contributes more to their perceived workload of the task. From this data and following a fixed evaluation scheme, an *overall task load index* in the range [0, 100] is computed. This overall index takes into account the user's ratings of each of the six aspects, weighted according to the aspect's importance concerning workload. Higher values represent greater workload.

2.8.5 Perception-Related Aspects

Evaluating the visual-haptic perception of users is crucial for the work in this thesis. In this context, we assessed the high-level aspect of perceived haptic realism, as well as low-level aspects like perceptual thresholds.

Haptic Realism

Perceived haptic realism is commonly assessed with *custom questions* rated on a 1-to-5 [Whitmire et al., 2018] or a 1-to-7 [Sinclair et al., 2019] scale. These questions are tailored to the experiment and ask users how well the perceived haptic stimuli match their expectation based on the visual feedback provided in VR. For example, Whitmire et al. [2018] asked users "*How closely did the haptic rendering match your visual impression of the scene?*" to assess haptic realism on a 1-

⁵⁰We used Keith Vertanen's online version of the NASA TLX. <https://bit.ly/39n7sZc>

to-5 scale. We follow this approach of custom questions and asked participants to report haptic realism on a rating scale when assessing our prototypes in Part III.

Perceptual Detection Thresholds

Several of our experiments derive perceptual Detection Thresholds (DTs). DTs help us to uncover how much HR can go unnoticed in Part IV, and quantify how well our technique proposed in Part V solves the *Colocation* challenge.

To derive DTs, we adopted established methods from the field of psychophysics and conducted psychophysical experiments. Different variants of threshold experiments exist and have been used by related work as outlined in paragraph 2.6.2. In this thesis, we employed the *method of constant stimuli* and the *staircase method* (also known as *adaptive up/down method*). With both approaches, we collect empirical perception data to derive estimates for absolute DTs, i.e., “*the magnitude of a stimulus that can be just discriminated from its null*” [Kingdom and Prins, 2016b; p. 31]. To derive such estimates, the performance of humans in detecting the presence of a stimulus is commonly modeled by means of a psychometric function [Leek, 2001]. This function maps stimulus intensities to the probability of the stimulus being detected by a human observer. The DT is commonly defined as the stimulus magnitude that results in equal probabilities for detecting and for failing to detect the presence of a stimulus [Leek, 2001; Steinicke et al., 2010b]. The following sections outline the two methods we applied to estimate this magnitude.

Method of Constant Stimuli In chapter 8, we explore the detectability of HR using the method of constant stimuli with a *symmetric 1-Alternative Forced-Choice (AFC) task*. In the experiment, we first collect data with which the psychometric function (modeling the users’ stimulus detection performance) can be approximated. Then, we use this psychometric function to compute the DTs [Kingdom and Prins, 2016c; p. 52].

To approximate the psychometric function, experiments applying the method of constant stimuli sample the performance of participants in detecting a stimulus within a certain range of stimulus intensities. This range spans from undetectable to well-detectable stimuli [Leek, 2001]. Participants are presented multiple times with intensities drawn randomly from a set of predefined values within this range. For each presented intensity, participants indicate their detection of the stimulus. From this data, the percentage of correct detection per intensity (approximating the probability of detection) is computed and a psychometric function is fit to the data. The intersection of the psychometric function with the probability that corresponds to the threshold yields the DT estimate. Since the psychometric function is reconstructed, the method of constant stimuli also informs about other intensities, such as the Point of Subjective Equality (PSE) [Steinicke et al., 2010b].

Utilizing a 1AFC task [Kingdom and Prins, 2016b; p. 26], participants in our

experiment experienced a single stimulus in each trial, corresponding to a specific magnitude of HR applied along a certain axis (e.g., horizontal). Since we utilized a symmetric design, participants were to indicate the direction of the experienced HR (e.g., responding either with “left” or “right”). If a stimulus in such a symmetric design is not detected, participants need to guess (forced choice). Symmetric designs can be considered bias-free [Kingdom and Prins, 2016c; p. 42].

Adaptive Staircase Method In chapter 9 and Part V, we use the staircase method (also known as up/down method) with a *yes/no-based 1AFC task* to derive DTs [Kingdom and Prins, 2016a; pp. 120 ff.]. In contrast to the method of constant stimuli, the adaptive staircase method does not reconstruct the psychometric function but directly and efficiently approximates the threshold itself [Leek, 2001].

Our staircase experiments present a single stimulus intensity per trial and participants indicate whether they notice the stimulus by responding “yes” or “no”. The intensities presented in each trial are based on the participant’s response in the previous trial [Kingdom and Prins, 2016a; p. 120]. Specifically, if a participant notices (fails to notice) a stimulus, the intensity of the stimulus in the following trial is reduced (increased) by a certain step size. The resulting stimulus sequence converges towards a specific percentage-correct value on the underlying psychometric function (i.e., the threshold) and starts to oscillate around it – sampling the participant’s perception most where it matters most for the threshold calculation [Leek, 2001]. Once a termination criterion is met, the sequence stops and the last reversal points, i.e., the intensities at which the sequence changed direction from upwards to downwards or vice versa, are averaged to approximate the threshold. The threshold targeted by a staircase is defined by the up- and downward step sizes, and by the rules governing intensity changes [Kingdom and Prins, 2016a; pp. 122 ff.]. We outline details like step sizes and termination criteria for each of our experiments separately in the respective chapters. For our staircase experiments, we developed the open-source *Unity Staircase Procedure Toolkit*.

2.8.6 Other Aspects

Other aspects considered in our evaluations include demographic information and experiment-specific aspects. Moreover, we intended to also capture aspects of the user’s experience that were unanticipated.

Following common HCI practices, data about the participants was collected through *demographic questionnaires* [Sharp et al., 2019; pp. 278 ff.]. These questionnaires assessed aspects such as the participants’ age, gender, and handedness, as well as previous experience with VR or other relevant domains such as 3D gaming, and the absence of experiment-critical health issues.

To capture experiment-specific and unanticipated aspects, we administered addi-

tional custom questionnaires and observed participants during the experiments. Furthermore, we conducted post-experiment interviews with the participants. The experiment-specific questions we used varied from study to study and are outlined in each study's corresponding chapter.

Observations [Sharp et al., 2019; pp. 287 ff.] of participants' interactions with the tested systems were carried out by the experimenter. For this, the experimenter took notes during the studies if any interesting observations were made so as to follow-up on those in later conversations with the participants.

Also, comments from participants allowed for further insights into their subjective impressions. Comments were provided either in written form through free-text fields in our custom questionnaires, or communicated verbally after the experiment. To this end, we conducted *semi-structured interviews*, which we also refer to as *(verbal) debriefings* in the remainder of this thesis. Such interviews provide a general structure to the dialogue with the participants while still providing the flexibility to follow up on interesting (and sometimes unanticipated) comments with further questions [Sharp et al., 2019; pp. 269 ff.]. To evaluate the semi-structured interview conducted in the context of chapter 10, we performed a basic *thematic analysis* [Braun and Clarke, 2012]. Interesting insights gained from observations and interviews were incorporated into the discussions of the respective studies.

Part II

Proxy-Based Haptics for VR

A Novel Application



... getting in *touch*

Research Question 1	<i>Applicability</i>
<p>This part addresses RQ 1: How can proxy-based haptic feedback support the domain of process model exploration?</p>	

Chapter 3

Immersive Process Model Exploration

In this chapter, we start advancing the field of proxy-based haptic feedback by exploring a new application area. In the context of **RQ 1**, we showcase the **Applicability** of haptic proxies in the domain of business process model exploration, which focuses on learning in professional contexts and offers great freedom in terms of how haptic interactions and involved proxies can be designed. We first provide a brief introduction into this domain and then present the concept of *Immersive Process Model Exploration* alongside the implementation of a VR system realizing it. At the end of the chapter, we report on an evaluation of our system and the impact of proxy-based haptics on the exploration of processes.

A video⁵¹ about the work presented in this chapter is available online. This chapter is based on the following two publications. Images and parts of the text in this chapter, as well as the presented concepts, implementations, and results have been published previously therein:

Zenner, A., Klingner, S., Liebemann, D., Makhsadov, A., and Krüger, A. (2019b). Immersive Process Models. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'19, pages 1–6. ACM. © 2019 André Zenner and co-authors. Final published version available in the ACM Digital Library.

DOI: [10.1145/3290607.3312866](https://doi.org/10.1145/3290607.3312866)

Zenner, A., Makhsadov, A., Klingner, S., Liebemann, D., and Krüger, A. (2020c). Immersive Process Model Exploration in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(5):2104–2114. © 2020 IEEE. Final published version available in the IEEE Xplore® Digital Library. DOI: [10.1109/TVCG.2020.2973476](https://doi.org/10.1109/TVCG.2020.2973476)



Video Link⁵¹

⁵¹Immersive Process Model Exploration Video. <https://bit.ly/3oMFkNw>

3.1 Introduction

In conventional desktop applications, abstract data is commonly conveyed to the user by means of text, pictures, videos, charts, graphs, or one of the many other representation formats that we have become used to interpreting. While many of these formats are successful in communicating abstract data to us, the rise of VR motivates research exploring how abstract data can be communicated in novel, alternative ways that have not been possible in the past. In this chapter, we investigate the communication of graph-based data through VR. Specifically, we focus on business process models that are used in a broad range of professional contexts and investigate how proxy-based haptic feedback can add to this domain. By proposing a multisensory VR system designed to support learning, we aim to make understanding processes easier, more enjoyable, more interactive, and more memorable compared to traditional interfaces.

3.1.1 Immersive Data Analysis

VR technology enables users to immerse themselves in IVEs that simulate realistic environments and in IVEs that visualize abstract data. The latter has spawned the field of immersive data analysis, where researchers have used large-scale projection systems (e.g., CAVEs) [Kuhlen and Hentschel, 2014; Laha et al., 2012] and HMDs [Zielasko et al., 2016, 2018; Sousa et al., 2017] to explore data in novel ways. In this context, previous work highlighted the importance of immersion for VR applications that focus on data investigation [Kuhlen and Hentschel, 2014; Laha et al., 2012], which motivates our research on multisensory feedback in this chapter. Related to the approach we will propose, first steps towards including proxy-based haptic feedback for immersive data analysis have been taken, for example, by Sousa et al. [2017]. In their immersive display of tomographic data for radiodiagnosis, the authors integrated haptic feedback during data exploration as users touched a physical desk to perform gestural input. Furthermore, Zielasko et al. [2017] proposed to represent physical objects, like desks, in the IVE during data exploration to provide tangibility.

While previous research started to investigate how VR can support the communication of data, the potential of multisensory VR as an immersive interface to communicate information about business processes has not yet been sufficiently studied. To fill this gap, we will present the concept and implementation of a novel VR system that employs proxy-based haptic feedback and enables the immersive exploration of business process data.

3.1.2 Process Models

In the domains of information professionals and business process management, real-world processes are often described as *process models*. Such process models formalize real-world procedures in a concise, abstracted format [Recker and Dreiling, 2011] and hold information about involved operational steps, stakeholders, decisions and dependencies. In professional domains, a process model might, for example, be used to formally describe the steps involved in the inspection of ordered goods and their delivery in a store, or how a customer complaint is handled in a company's support center. For such formal descriptions of business processes, several representation formats exist. In this chapter, we will consider a widely used standard format which depicts processes as *Event-Driven Process Chains* (EPCs) [Keller et al., 1992]. EPC is a graphical representation format similar to alternative formats such as, for example, *Business Process Model and Notation* (BPMN) [Object Management Group, 2013]. EPCs lay out the flow of a process, involved steps, and stakeholders in a graph structure, as can be seen in Figure 3.1. Such graph structures are human- and machine-readable, which makes them particularly suitable for a range of application areas. Hence, EPCs are widely used for education, documentation, evaluation, simulation, optimization, and worker guidance [Knoch et al., 2019].

Event-Driven Process Chains (EPCs)

Being a structured representation of a process, an EPC formally consists of *function*, *event*, and *logical connector* (*and*, *or*, *xor*) nodes. EPC nodes are connected by arrows that indicate the process flow. Further process details are added utilizing special node types like *organization unit*, *input*, *output*, or references to other process models. Different node types represent different process elements:

1. *Events* (red in Figure 3.1) have no duration and represent process states which can trigger functions.
2. *Functions* (green), in contrast, can take time as they are active elements representing process activities.
3. Through *logical connectors* (grey circles), the process flow might be split into multiple branches, or branches can be merged.
4. Additional information is attached to functions as connected information nodes like *organizational unit nodes* (yellow), *input nodes* and *output nodes* (grey rectangles), or linked processes (blue).

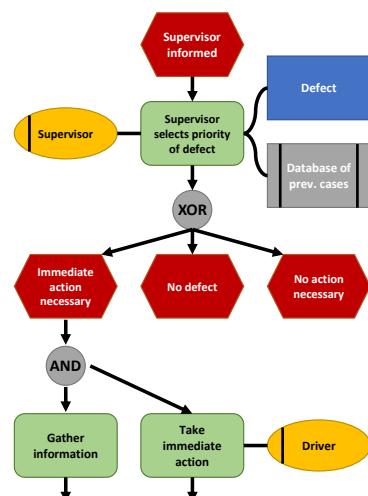


Figure 3.1: An EPC in 2D.

Each node type is visualized with a specific 2D shape and color. Complete process models consist of an alternating sequence of *events* and *functions*.

EPCs constitute a novel application domain for VR and proxy-based haptics. They are well suited for the investigation of **RQ 1**, since they represent data that is concise, consists of different node types with different meanings, as well as established shapes, colors, and logical operators. Moreover, EPCs come with an appropriate number of nodes for VR visualization. Besides that, the exploration of graph-based EPCs might further support a generalization of some of our findings to other domains that also feature graph-based data visualizations.

Process Model Understandability

In order to familiarize with novel processes, users traditionally explore the corresponding 2D models (e.g., EPCs) of the processes by looking at their graph representation on paper, or using a desktop or mobile UI. Hence, such traditional interfaces only target the user's visual channel to convey the logical connections and dependencies within an organizational process – leaving out other senses that might contribute to a better UX. As a result, understanding a process becomes more complicated and 2D charts become harder to read as the processes portrayed get more sophisticated and increase in size.

To investigate this aspect, several research projects have studied *process model understandability* in the past. In this context, Recker and Dreiling [2011] found that process model understandability is influenced by several factors including, for example, previous process modeling experience and the use of English as a second language. Moreover, Houy et al. [2012] conducted a review of 42 experiments on the understandability of conceptual models and distilled the various concepts of model understandability found in the related research in a central *conceptual model understandability reference framework* – a well established framework that we will base the evaluation at the end of this chapter on.

3.1.3 Motivation

In order to approach **RQ 1**, we question the suitability of traditional 2D UIs for scenarios in which laypersons (with regard to the domain of process modeling) need to internalize unfamiliar processes. Our motivation is sparked by representative use cases, such as when new employees in a company are acquainted with important processes as part of their onboarding (like the accounting of a business trip, or the re-ordering of new goods), or in educational scenarios where students learn about process modeling. In these scenarios, understanding a process represented as a 2D EPC graph can become difficult for inexperienced users and cause frustration as the process complexity increases. Yet, especially in situations like employee training, customer presentations, or education, ensuring a motivating and interesting UX can be of key importance.

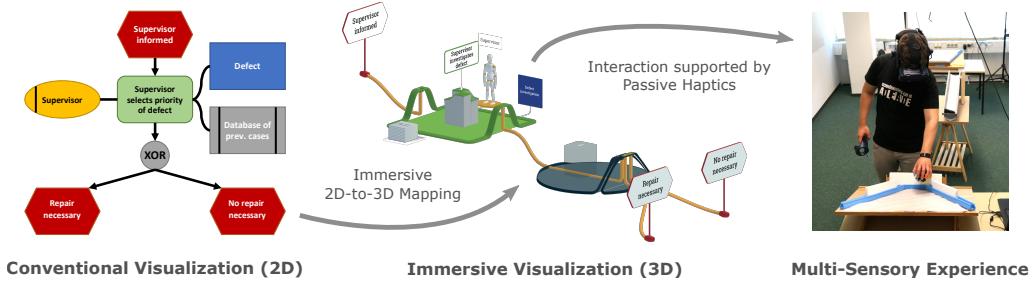


Figure 3.2: Illustration of the *Immersive Process Model Exploration* system. Left: Process models are traditionally communicated through 2D graph visualizations. Center: Our system, in contrast, automatically generates a 3D IVE from any given process model that can be interactively explored by users. Right: Through basic gamification and auditory, vibrotactile, and proxy-based haptic feedback, our system turns learning a process into a multisensory VR experience. © 2020 IEEE.

With this in mind, in this chapter, we investigate how multisensory VR can change the way users experience the exploration of process models. Based on the motivational examples above, our system targets the user group of non-specialists as they can draw less on previous EPC experience and consequently might especially benefit from a less formal way of exploring processes. Specifically, the system we will present is designed to be used in the context of education (e.g., for students learning about the concept of process models), internal company training (e.g., to teach employees about new processes important for their work), or customer communication (e.g., to present process optimization results to clients, or to explain company-internal processes to a customer).

3.2 The Immersive Process Model Exploration System

To turn the exploration of EPCs into an interactive and memorable VR experience, we developed the *Immersive Process Model Exploration* system. The leading goal of the system is to enable users to associate a personal experience with the explored process. To achieve this, the *Immersive Process Model Exploration* system lets users become an active part of the process.

We introduce a system that given an EPC as input automatically generates an IVE that is greater than room scale and represents the process in 3D as illustrated in Figure 3.2. This IVE allows users to dive into the graph, experiencing an immersive first-person walkthrough through the process chart, as shown in Figure 3.3. Leveraging a combination of locomotion techniques and gamification, our system allows users to explore the process on foot and to interact with it. Users can transport information bits through the process following its operational flow and interactively experience decisions involved in the process. To boost immersion during these interactions, our system further supports two types of haptic feedback, spanning from active to proxy-based haptics.

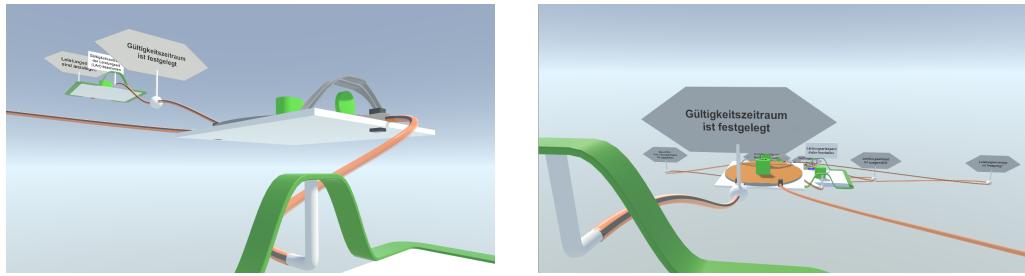


Figure 3.3: First person VR view inside an immersive process model. Left: The user looks back to the root of the process. Right: The user looks down at the remaining parts of the process. All screenshots show an early version of our proposed system and depict EPCs labeled in German.

To realize this basic concept, the *Immersive Process Model Exploration* system consists of three central components:

1. *2D to 3D Mapping* – A component responsible for generating an immersive visual 3D representation of the explored process in VR.
2. *Logical Walkthrough* – A component to motivate users to explore a process model with the primary objective to provide guidance while highlighting logical dependencies within the process flow.
3. *Haptic Interactions* – A component that supports immersion by transforming the experience into an interactive journey, allowing for haptic interaction with information bits and the process flow throughout the graph.

In the following, we describe these three main components and their implementation in more detail.

3.2.1 2D to 3D Mapping

The first component of our system is responsible for an immersive visualization of EPC models. To allow users to leverage their natural skills of spatial orientation while exploring process models, the first central component of our system spatializes the EPC to be explored. Given an EPC in a standard file format as

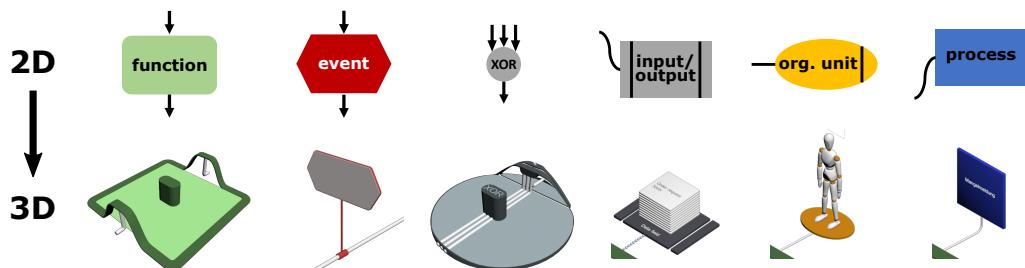


Figure 3.4: Traditional 2D EPC elements and their 3D representations. © 2020 IEEE.

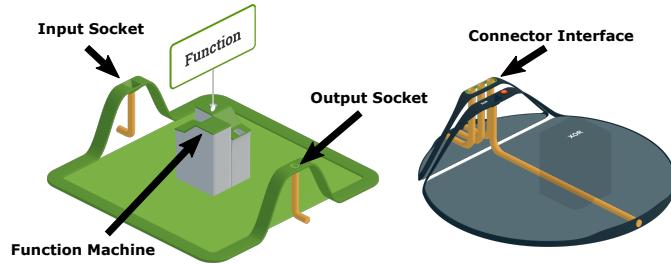


Figure 3.5: Interactive elements on function and connector platforms. ©2020 IEEE.

input (our implementation supports EPCs in .aml, .epml, and a specific .xmi data format), a parser loads the process model and a *2D to 3D Mapping* algorithm generates an immersive virtual 3D representation of the process as an output.

The mapping algorithm generates a virtual world in which the nodes of the EPC are represented by floating platforms. Functions and connectors are represented by room-sized, walkable platforms and events in the graph are displayed as virtual signs. Further node types like organization units, inputs or outputs are likewise represented by corresponding 3D objects. The visual design of the 3D elements is based on the original design of the 2D EPC elements taking shapes and colors into account in order to facilitate recognition and knowledge transfer. The individual elements of the 3D process model are connected by a virtual tube system – the 3D representation of the edges in the EPC graph. This tube system is used to transport information bits from the beginning of the process to the end of the process from element to element. Figure 3.4 depicts how different 2D EPC structures are translated into corresponding 3D objects. Subsequent process elements are placed spatially below preceding elements to visualize the flow direction of the process through a descending platform layout. Figure 3.2 (center) depicts the descending 3D environment generated by our *2D to 3D Mapping* that corresponds to the 2D EPC shown in Figure 3.2 (left).

3.2.2 Logical Walkthrough

The second central component of the introduced system handles the user's travel through the virtual process. For long-distance travel from one walkable platform to another walkable platform, our system implements the established overt relocation method of *teleportation*⁵². Figure 3.6 (a) shows a screenshot. Once the user is transported to the node of interest, we decided to let users *physically walk* within the boundaries of the corresponding virtual platform since previous research found natural walking to be superior in terms of presence compared to more stationary techniques [Usoh et al., 1999]. To avoid collisions with the physical surroundings, the size of the virtual platform corresponds to that of the physical tracking area.

⁵²SteamVR Unity Plugin. <https://bit.ly/3DXjwoh>

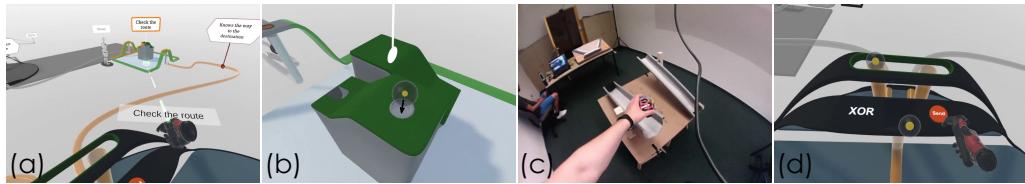


Figure 3.6: Immersed view: (a) Teleporting to the next node. (b) Putting the information packet into a function machine. (c) Real-world perspective of the interaction in (b). (d) Sending the information only to the left child of the xor node. All images © 2020 IEEE.

While freely exploring the nodes of a process model is one way to use our system, we additionally implemented a basic guidance system that enforces a logical exploration path through the graph called *Logical Walkthrough*. In the *Logical Walkthrough* mode, users need to carry information packages from the beginning of the process to the end, which are represented by a virtual sphere shown in Figure 3.6 (b). Users start at the process root, the only *unlocked* node at the beginning, and can only visit already unlocked platforms in the process (see Figure 3.7). Further process nodes can be unlocked node by node through correct interaction with function and connector platforms.

To unlock a node in the process, the information package must be transported to the respective node. Each 3D function platform contains an abstracted virtual machine (shown in Figure 3.5) that has to be operated interactively by the user. To proceed with an information package at a function platform, the user has to pick up the incoming information at the platform's input socket and drop it into the function machine on the platform as shown in Figure 3.6 (b). The machine processes the information and ejects a new information package. This processed information package is then to be picked up again by the user and sent through the virtual tube system to the next platform at the output socket to unlock the following node.

Similar to the interaction with functions, users also interact with the process on connector platforms. The system supports all logical operators: `or`, `xor`, and `and`. In contrast to function platforms, connectors can have several incoming or outgoing tubes. At connectors, the task of the user is to provide the necessary input for the connector according to its logical type. For an `and` connector with two incoming tubes, for example, the user has to send an information packet

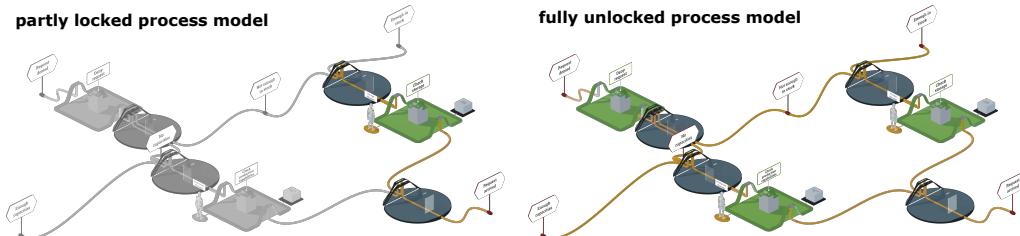


Figure 3.7: Partly locked (left) and fully unlocked (right) process. Images © 2020 IEEE.

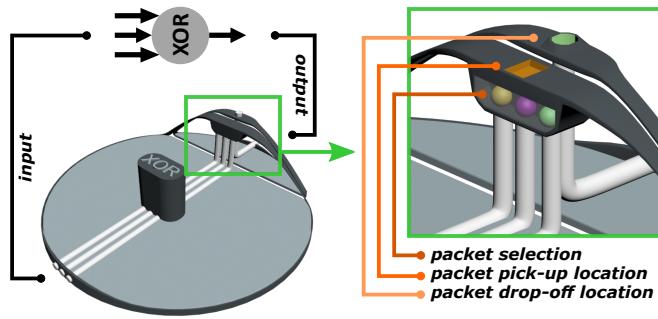


Figure 3.8: Concept of a 3D *xor* connector platform. The user can select and pick up exactly one incoming information packet. To send it to the next node, the user drops the packet again at the green drop-off location.

through each of the two incoming tubes, which leads the user to go through the corresponding previous process steps. When the input requirements of a connector are fulfilled, the user can decide how the information at the connector will flow further through the process. For this, users can interact with a connector interface on the platform. Figure 3.8 depicts an early concept of this interface. At an *xor* connector with multiple outgoing tubes, such as depicted in Figure 3.6 (d), for example, the user can select to which of the following platforms an information package is moved.

While the difference between passive events (i.e., process states that take no time) and active functions (i.e., activities that take time to execute) is only weakly communicated with traditional 2D EPC representations, the interaction in our system facilitates the perception of functions as *active* components of the process to raise awareness for the relevant process steps. Furthermore, the interactions with connector platforms that control the process flow are designed to strengthen the understanding of logical decisions and dependencies occurring in the process. In sum, all these aspects of the *Logical Walkthrough* mode guide the user through the process in a logically meaningful order. The developed system transforms the exploration of a process from a passive observation of the 2D graph to an interactive experience in a 3D world. By this, our system aims to let users associate a personal and spatial experience with the explored process.

3.2.3 Haptic Interactions

The third component builds on the visualizations generated by the *2D to 3D Mapping* and the interactions with the platforms enforced by the *Logical Walkthrough*. Large-scale setups have been used in the past for immersive data analysis where space was required to immerse users with projection systems like CAVEs [Kuhlen and Hentschel, 2014]. We propose to utilize the visual and auditory quality of modern HMDs and leverage the physical space for multisensory experiences by integrating haptic feedback.

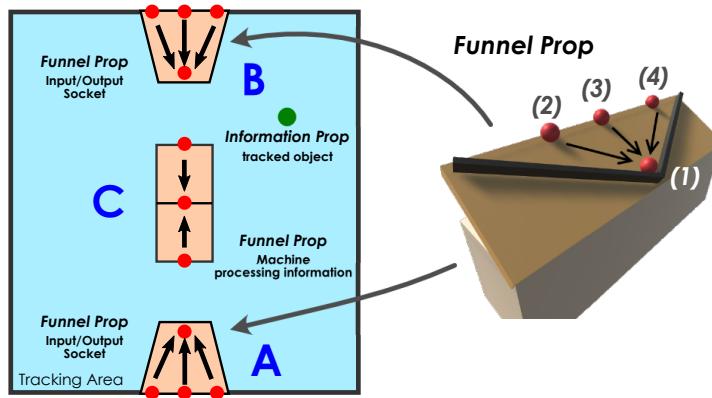


Figure 3.9: Left: Symmetric tracking area and involved physical props. The blue area maps to virtual platforms. Right: Rendering of a physical funnel prop. (1) marks the funnel and pick-up location; (2), (3), and (4) mark locations to drop an information prop.

While classical interfaces for process model exploration (e.g., 2D representations on paper or displays) only allow for visual inspection of the process model, our *Haptic Interactions* component additionally introduces the auditory and haptic dimensions. Specifically, wherever information flows into or out of a walkable node, at connectors, and at the function machine, interactive sockets can be found in the IVE as illustrated in Figure 3.5. By letting users manually transport information packets from and to these sockets, we aim to raise the awareness that function nodes represent actions that take time, and that operators impact the process flow. To this end, we implemented two levels of haptic feedback: active vibrotactile feedback and passive proxy-based haptic feedback.

Active Vibrotactile Feedback

In a first mode, users can explore the 3D process model while carrying an HTC Vive Pro Controller⁵³ (see *Controller* condition in Figure 3.12). The controller triggers two different vibration patterns during interaction to signal either a positive or a negative outcome (successful or unsuccessful interaction). For a successful interaction, a continuous vibration of 0.75s was triggered, while in the case of an unsuccessful interaction, four vibrations of 0.25s each were triggered with pauses of 0.25s in between. Similarly, basic sound effects were played to support the positive or negative feedback. This feedback mode serves as a basic “notification” in response to virtual events and interactions.

Passive Proxy-Based Haptic Feedback

In a second mode, the system leverages haptic props to increase the fidelity of the interactions with the virtual process, making them more physical and

⁵³HTC Vive Controller (2018). <https://bit.ly/3aU1CpT>

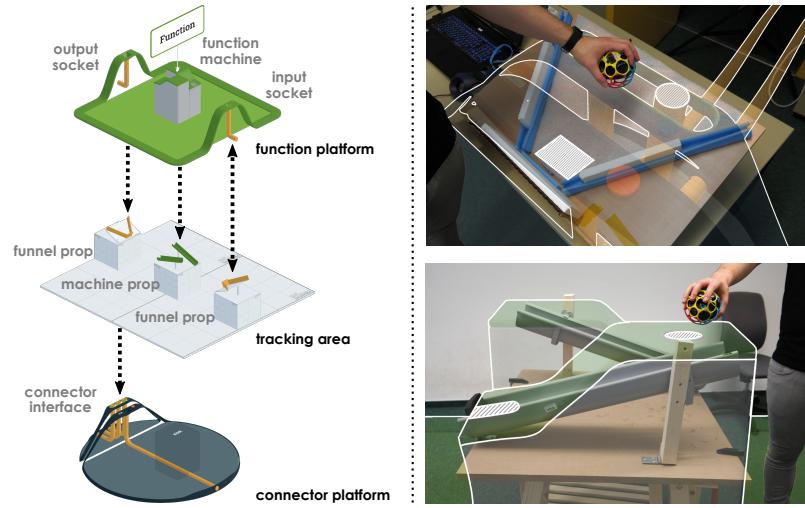


Figure 3.10: Left: Registration of the physical setup with the platforms. Top Right: Funnel prop overlaid with connector interface. A user drops information into an outgoing tube (hatched area on the right). When released, the ball rolls down the funnel where it can be picked up again later (hatched area at the bottom). Bottom Right: Machine prop overlaid with function machine. A user drops information into the machine (hatched area on the right). The ball rolls down the prop as on a marble run and stops at the machine's output (hatched area on the left). Here, it can be picked up again, now representing the processed information. All images © 2020 IEEE.

engaging. Here, the user explores the IVE with an HTC Vive Pro Controller in the non-dominant hand, leaving the dominant hand free to interact with physical props located within the physical tracking space. Conceptually, our approach is related to *iTurk* [Cheng et al., 2018b], as props are manipulated by the immersed user and reused throughout the experience. The conceptual sketch in Figure 3.9 and the image on the right of Figure 3.12 show the layout of the symmetrical real environment in this haptic feedback mode. Following the classical approach of passive haptics [Insko, 2001], virtual objects in our application are not represented by 1-to-1 replications, but by three different types of low-fidelity props – i.e., physical proxies that allow for realistic interactions while being simplified and not representing the virtual counterparts in full detail:

1. **Mesh-Ball Prop:** The information bit that is to be transported by the user through the process is represented by a physical mesh ball, shown on the right in Figure 3.10. It is made out of a toy ball containing an HTC Vive Tracker⁵⁴ and allows for robust tracking even when carried with one hand.
2. **Funnel Prop (2x):** Haptic interactions at the input and output sockets of function nodes, and at the connector interface of connector platforms, take place at two funnel props placed at opposite ends of the tracking space. Figure 3.9 shows a 3D concept of such a prop. Each funnel prop has a

⁵⁴HTC Vive Tracker (2018). <https://w.wiki/3xG2>

tilted surface registered with the outgoing tubes in VR, and two funneling wooden slats. When dropping the information bit into any outgoing tube, the physical mesh-ball prop will drop onto the surface of the funnel prop at the location of the tube. Pulled down by gravity, the ball prop will roll down to the bottom of the funnel where it will be picked up again and reused later in the experience (e.g., at the input socket of the next platform).

3. **Machine Prop:** The machine in the center of function platforms is represented by the symmetrical prop shown in the bottom right image in Figure 3.10, holding two tilted gutters. The upper ends of the gutters are registered with the input slots of the virtual function machine. When dropping the information packet into these slots, as shown in Figure 3.6 (b), the physical mesh ball will roll down the gutter and end up at the lower end on the opposite side of the machine. The lower end is registered with the machine’s output where the information prop can be picked up again later. This means that upon termination of the virtual machine’s processing, it is the same mesh ball that physically represents the new information packet ejected by the machine.

Enabling Proxy Reuse Figure 3.10 illustrates the real-to-virtual registration and the physical room setup. To enable exploration of arbitrarily large and complex EPCs, we take advantage of the great design freedom offered by the domain of immersive data exploration. By leveraging the symmetry of the room setup and proxies alongside a custom *180° resetting controller*, our system can convey any EPC with only a single mesh-ball prop, two funnel props and a single machine prop. The *180°* resetting controller is similar to the adapted *Stop and Reset* technique used by Simeone et al. [2020]. It activates, for example, when the user stands in front of a platform’s output socket and drops the mesh ball at an outgoing tube to send an information packet to the following platform. When teleporting to this platform, the resetting controller quickly fades the user’s view to black, teleports the user’s position to the target platform’s input socket, rotates the view of the user by 180° and fades the view back in. As a result, the user can pick up the mesh ball just released at the previous platform’s output from the new platform’s input and continue the experience by turning around. These resets effectively mirror the real-to-virtual registration and enhance the flexibility of the proxy-based haptic feedback by enabling the reuse of proxies across the IVE. Hence, in summary, we use resetting both as a technique to enable the exploration of the IVE through walking, and at the same time as a technique to establish proxy *Colocation* as proposed by Suhail et al. [2017].

3.3 Evaluation of Immersive Process Model Exploration

To investigate the potential offered by an experience-focused VR interface for abstract data exploration, we compare our proposed *Immersive Process Model Exploration* system to a traditional interface that allows users to explore the 2D graph representation of processes on a tablet device. To study how our immersive interface affects the performance of users in understanding and learning new processes, we base our evaluation on the established *conceptual model understandability reference framework* by Houy et al. [2012]. Specifically, our study collects data on all three of the framework's main understandability dimensions (objective effectiveness, objective efficiency, and subjective effectiveness) to assess how well users understand a process model when explored with the different interfaces. Moreover, our user study is designed to validate and evaluate the proposed system and to gain insights into its benefits and drawbacks.

In particular, within the scope of our user study, the suitability of the multi-sensory *Immersive Process Model Exploration* system for mediating processes unknown to the user was tested. Our study scenario is motivated by the use case of familiarizing a new employee with a company process as part of the onboarding procedure – a scenario where content is to be communicated in a motivating and interesting way. To reflect this scenario, it was of particular interest to investigate how well users with little or no previous experience with EPCs and the domain of the test process can learn and understand a new process flow.

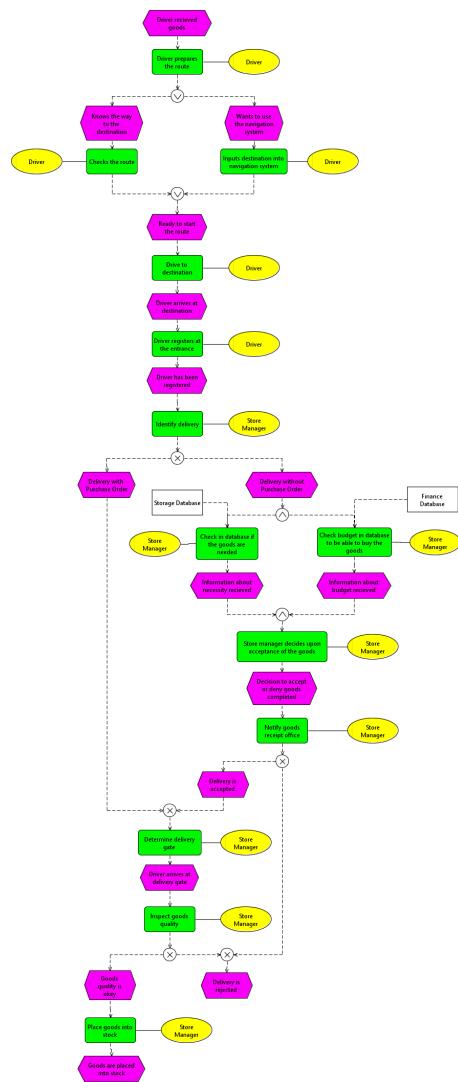


Figure 3.11: 2D layout of the process used in our study (generated using the *bflow** Toolbox [Böhme et al., 2010]). It describes the delivery of goods to a store (an extended and slightly modified version of the test process by Recker and Dreiling [2011]). © 2020 IEEE.

Following **RQ 1**, the evaluation of our system aims

1. to compare our novel 3D VR interface (*Controller* condition and *Props* condition) to a traditional 2D interface (*iPad* condition)
2. to compare the mode with active vibrotactile feedback (*Controller* condition) to the mode with proxy-based haptic feedback (*Props* condition)

In this context, we hypothesized that:

- H1** Learning an EPC with our VR interface will require more time than with a traditional 2D interface due to the interactive experience involved.
- H2** Learning an EPC with our VR interface yields better learning results than learning with a 2D interface, due to multiple senses being involved.
- H3** EPC exploration with our VR interface offers an enhanced user experience compared to traditional 2D EPC interfaces, as it is designed to spark the interest of the user.
- H4** Proxy-based haptic feedback increases immersion when exploring EPCs in VR compared to vibrotactile controller feedback, as supported by prior research [Insko, 2001].

To investigate these hypotheses, our evaluation comprised three conditions in total. The *Controller* and *Props* conditions both represent the *Immersive Process Model Exploration* system and are implemented as described before, providing vibrotactile feedback with the controllers and passive proxy-based haptic feedback in VR, respectively. In addition, a traditional *iPad* condition (shown in Figure 3.12) served as a control condition in our experiment. For this, we displayed a 2D EPC visualization generated by the open-source EPC modeling application *bflow* Toolbox*⁵⁵ [Böhme et al., 2010] on an Apple *iPad*. We chose a 2D representation on a mobile device since tablets represent a state-of-the-art exploration method which allows to inspect arbitrarily large EPCs with an interface fixed in size. In this *iPad* condition, users could freely explore the 2D graph using standard multitouch interactions such as scrolling and zooming, which additionally renders this kind of interface more flexible than paper while providing a more comfortable form factor and reading experience than desktop monitors.

The process used in our study is an extended and slightly modified version of a test process used in related research [Recker and Dreiling, 2011]. It depicts the process of delivering goods to a store, starting with the delivery driver and ending with the acceptance or rejection of the goods by the store manager. Figure 3.11 displays the 2D visualization of the full test EPC as it was shown in the *iPad* condition. The experiment was approved by the ethical review board of our faculty and took place in our lab.

⁵⁵*bflow* Toolbox*. <https://bit.ly/3CxOyTu>



Figure 3.12: The three conditions tested in our user study. ©2020 IEEE.

3.3.1 Participants

A total of 29 volunteer participants recruited with flyers on the local campus took part in the study. We only included participants who had normal or corrected-to-normal vision, and who confirmed having neither a hearing impairment nor a haptic perception disorder that could affect their VR experience. Out of the 29 complete data sets, 27 of them (16 male, 11 female) were included in the final data analysis, while the data of 2 participants had to be excluded from analysis as the participants did not fulfill these aforementioned requirements for participation. The average age of the participants was 25 years (min. 22 years, max. 34 years); 2 participants were left handed, while 25 were right handed. Apart from 1 participant, all participants were inexperienced with EPCs and the domain of the test process. Moreover, 20 participants had very little or no experience with VR, while 7 were somewhat or very experienced in VR.

3.3.2 Apparatus

The experimental setup can be seen in Figure 3.12. For the *iPad* condition, an Apple iPad (iPad Wi-Fi 128 GB) was used, while for the *Props* and *Controller* conditions, an HTC Vive Pro⁵⁶ VR system was set up. It consists of a Lighthouse Tracking System, an HTC Vive Pro VR HMD, two HTC Vive Pro Controllers⁵³ and additional HTC Vive Trackers⁵⁴ for tracking physical props. Our VR application was developed using the Unity 3D engine¹⁷. The passive haptic proxies used in the *Props* condition were assembled using readily available materials such as wood, styrofoam, plastic gutters and a toy ball, as can be seen from Figure 3.12.

⁵⁶HTC Vive Pro System. <https://bit.ly/3GDx2j4>

3.3.3 Procedure

Each participant was assigned to one of the three tested conditions (*iPad*, *Controller*, *Props*). For the *Props* condition, the experimenter calibrated the setup by spatially registering the physical props (two funnel props, one machine prop) with the virtual objects in the scene (input and output sockets, connector interface and function machine). For the *Controller* condition, the experimenter cleared the tracking space of any physical objects and for the *iPad* condition, a table and a chair were prepared for the participant. After signing a consent form, a short introduction about the concept of EPCs was read by the participant, followed by a short tutorial on the respective method used to explore the process model.

When the introduction was completed, the task of the participants was to freely explore the test process with the assigned interface and to inform the experimenter as soon as they felt that they had understood the process. Participants were not required to visit every process platform in VR. The experimenter observed the exploration and recorded the time it took until the participant indicated having understood the process. Upon this indication, the participant stopped using the respective interface and answered a series of questionnaires on a laptop. The study took approximately 90 minutes per participant and each participant received a compensation of € 10 for their time.

3.3.4 Design

The study was designed as a one-factorial between-subjects experiment with the factor being the EPC exploration interface. The three conditions *iPad*, *Controller*, and *Props* were experienced by 9 participants each. Participants answered a set of questionnaires after exploring the EPC (in the order given below), to assess the following dependent variables:

1. three central dimensions of the *conceptual model understandability reference framework* by Houy et al. [Houy et al., 2012]:
 - (a) *Objective efficiency*, given by the time measured from the beginning of the process exploration until the participant indicated to have understood the process. (H1)
 - (b) *Subjective effectiveness*, given by the response to a corresponding 7-point Likert scale question. (H2)
 - (c) *Objective effectiveness*, given by the number of correct answers to 12 comprehension checkbox questions about the test process. Our questions were based on the questions used by Recker and Dreiling [2011] in their previous work on process model understandability. (H2)
2. task load (measured by the NASA TLX [Hart and Staveland, 1988]) (H3)
3. immersion (measured by the SUS [Slater et al., 1994]) (H4)
4. user experience (measured by the UEQ [Laugwitz et al., 2008]) (H3)
5. qualitative responses from the participants, gathered through answers to experiment-specific questions and debriefing comments.

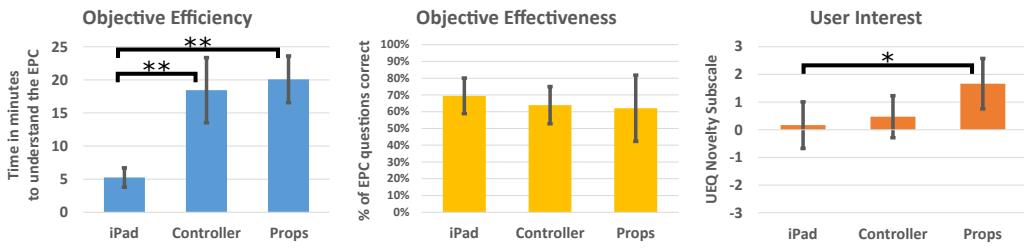


Figure 3.13: From left to right: Between-condition comparison of the *objective efficiency* (i.e., time in minutes participants took to understand the process [Houy et al., 2012]), *objective effectiveness* (i.e., participant performance in answering process model understandability questions [Recker and Dreiling, 2011]), and user interest (measured by the UEQ novelty subscale [Laugwitz et al., 2008; Schrepp, 2019]). Brackets indicate statistically significant differences ($p' < .05(*)$; $p' < .01(**)$). Error bars show 95% confidence intervals. All charts © 2020 IEEE.

3.3.5 Results

We compared the three EPC exploration interfaces by conducting a series of statistical tests on the measurements of the dependent variables. Our significance level was set to $\alpha = .05$ and we conducted non-parametric Kruskal-Wallis tests with post-hoc Dunn-Bonferroni tests and Bonferroni-Holm correction (p') where applicable, to test for significant differences between conditions. In the following, we only describe the most relevant and significant results of our analysis.

The results for *objective efficiency*, i.e., the average time in minutes that participants took to understand the model, showed, that the exploration time of the 2D process graph with an *iPad* was significantly shorter than the exploration time in VR with *Controller* ($Z = 3.362, p' \leq .003, r = .79$) and *Props* ($Z = 3.868, p' \leq .003, r = .91$). However, our tests did not indicate any significant differences between the three interfaces concerning *subjective effectiveness* and *objective effectiveness*, i.e., the participants' performance in understanding the EPC and answering the understandability questions. Concerning *task load* (NASA TLX) [Hart and Staveland, 1988], again no significant differences were found. Figure 3.13 shows the results for *objective efficiency* (left) and *objective effectiveness* (center).

A central aspect of our investigation was to evaluate the effects of our VR interface on the user experience, measured by the UEQ [Laugwitz et al., 2008]. When analyzing the respective subscales we found the *UEQ novelty subscale* to be rated significantly higher for the *Props* condition than for the traditional 2D *iPad* condition ($Z = -2.560, p' \leq .03, r = .60$). Figure 3.13 (right) visualizes the corresponding results on the respective scale from -3 to 3. This subscale encompasses four questionnaire items and measures a hedonic quality aspect of an interface. It is used to assess if a system is perceived as "*innovative and creative*" and whether it "*catch[es] the interest of users*" [Schrepp, 2019]. As such, it is a crucial measure for our experience-focused system. Other subscales of the UEQ did not yield statistically significant results. To better understand the qualities

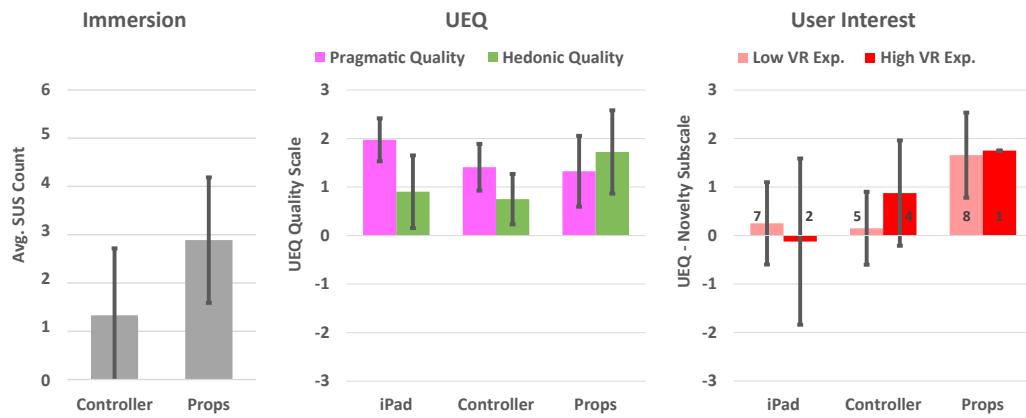


Figure 3.14: Left: Immersion results given by the SUS Count [Slater et al., 1994]. Center: Comparison of the *pragmatic quality* and *hedonic quality* as assessed by the UEQ. Right: Comparison of the *user interest* as assessed by the UEQ novelty subscale between participants with low VR experience (self-assessments ≤ 3 on 1-to-7 Likert scale; $n = 20$) and participants with high VR experience (self-assessments > 3 ; $n = 7$). Labels show respective participant count; error bars show 95% confidence intervals (except rightmost bar). All charts © 2020 IEEE.

of the tested interfaces, we also analyzed the scores for *pragmatic quality* and *hedonic quality* provided by the UEQ (see central chart in Figure 3.14). The *iPad* interface showed the highest ratings for pragmatic quality ($M = 1.97$, $SD = .58$), which supports the results for objective efficiency, but our tests did not show the pragmatic quality of the *Controller* ($M = 1.41$, $SD = .63$) and *Props* ($M = 1.32$, $SD = .95$) VR interfaces to be significantly different. Concerning hedonic quality, the VR *Props* condition scored highest ($M = 1.72$, $SD = 1.12$), supporting the UEQ novelty subscale results. As for pragmatic quality, however, corresponding tests did not show a statistically significant difference from the hedonic quality of the *Controller* condition ($M = .75$, $SD = .68$) or the *iPad* condition ($M = .90$, $SD = .97$).

Our evaluation of the system's immersion was based on the well-established SUS presence questionnaire [Slater et al., 1994]. To test if the type of haptic feedback affected immersion, we compared both the *SUS Count* and *SUS Mean* measures between the two VR conditions *Controller* and *Props* with non-parametric Mann-Whitney-U tests. *SUS Mean* ($M = 4.96$, $SD = .85$) and *SUS Count* ($M = 2.89$, $SD = 1.69$) immersion scores were higher for the passive haptic *Props* condition compared to the *SUS Mean* ($M = 4.11$, $SD = 1.03$) and *SUS Count* ($M = 1.33$, $SD = 1.80$) scores of the vibrotactile feedback *Controller*. However, differences were not statistically significant concerning both *SUS Mean* ($U = 59.5$, $p = .092$) and *SUS Count* ($U = 61$, $p = .063$). Table 3.1 summarizes the comparisons and Figure 3.14 (left) visualizes the *SUS Count* results.

In a concluding questionnaire and debriefing, we asked participants to reflect on the positive and negative aspects of the experienced interface and asked for any sickness symptoms experienced during the study. The very low overall

Measure	Range		iPad	ControllerProps
Objective Efficiency	minutes	<i>M</i> <i>SD</i>	5.24 1.91	18.44 6.40
Objective Effectiveness	%	<i>M</i> <i>SD</i>	69.44 13.82	62.04 14.43
Subjective Effectiveness	1 to 7	<i>M</i> <i>SD</i>	6.56 .527	5.44 .667
NASA TLX	0 to 100	<i>M</i> <i>SD</i>	43.62 12.84	38.14 14.97
UEQ Novelty (Interest)	-3 to 3	<i>M</i> <i>SD</i>	.17 1.10	.47 .99
UEQ Pragmatic Quality	-3 to 3	<i>M</i> <i>SD</i>	1.97 .58	1.41 .63
UEQ Hedonic Quality	-3 to 3	<i>M</i> <i>SD</i>	.90 .97	.75 .68
SUS Mean	1 to 7	<i>M</i> <i>SD</i>	- -	4.11 1.03
SUS Count	0 to 6	<i>M</i> <i>SD</i>	- -	2.89 1.80

Table 3.1: Comparison of the *iPad*, *Controller*, and *Props* conditions. © 2020 IEEE.

sickness rating ($M = 1.41$, $SD = 0.93$) out of a 1-to-7 Likert scale self assessment confirmed the absence of cybersickness issues. The qualitative feedback of the participants is discussed in the following.

3.4 Discussion of Immersive Process Model Exploration

Through our study we gained insights into the benefits and drawbacks of our proposed *Immersive Process Model Exploration* interface, of the role of proxy-based haptics, and of the conventional 2D approach in comparison. Furthermore, we discovered a central tradeoff when immersive VR and conventional 2D interfaces are considered against each other.

3.4.1 User Opinions: Benefits and Drawbacks of 2D and VR

From the observations in the context of our study, our results, and the qualitative feedback of our participants we could identify several important benefits and drawbacks of the tested approaches. We summarize them in the following, referring to comments from our participants.

The traditional 2D interface was appreciated by our participants for being “*easy to hold and carry around and [...] good for showing and interacting with other people to discuss the EPC*”. Moreover, a participant commented “*I do not have to learn some new interaction techniques as zooming etc. works as expected*”. While the interface was described as “*easily understandable*”, others found that it “*gets confusing if the graph is very wide*” and “*if you look at multiple [...] even more complex EPCs you might start to get lost [...]*”. Finally, one participant wrote “*memorizing something works*

better for me if I actually interact a little bit with the things I have to memorize, instead of only reading” – highlighting a need that our proposed interactive VR interface aims to satisfy.

Participants described the immersive data visualization as “*clear*” and “*mak[ing] it easier to go through the steps afterwards again when needed*”. The *Controller* condition was perceived as “*a very useful tool*” for making new employees familiar with new processes in a company by one participant and was further characterized as “*new and different to other learning experiences*” and thus as being “*more attractive*”. Commenting on our gamification of the EPC exploration, one participant wrote that “*since you can't go further, when you ignore the operators, you quickly learn your mistake and can fix it – that makes it more memorable*”, which supports our experience design. Many participants also reported to connect a personal experience with the walkthrough – thinking of logical branching within the graph more as “*being in different places*” and “*making decisions that have different consequences*” (comments translated to English). This circumstance was summarized by one participant stating “*also it is more fun than staring at a piece of paper with the graphical representation as in the introduction*”. The possibility to interact with the process representation is one of the central differences of our VR system compared to 2D interfaces and was received well by most participants. In the *Props* condition, for example, one participant commented that it is “*easier to remember EPC because of physical interaction, more senses are involved in [the] experience [...]*”.

Concerning the limitations of the *Immersive Process Model Exploration* system, drawbacks mentioned by our participants on the one hand encompassed general limitations of today's VR systems, such as uncomfortable HMDs (“*VR helmet is too heavy and it gets too hot inside*”). On the other hand, however, some participants also pointed out limitations specific to the implemented VR EPC interface. Supporting the results for *objective efficiency*, one participant noted that with a VR interface it “*takes time to have a look on all events [and] functions*”. Moreover, one participant criticized the recurring interactions implemented in our system, commenting that “*tasks are so standardised that you can get to the end of the process without thinking about the actual content*”. Another drawback that was mentioned is the limited mobility of our setup and the currently limited possibility for collaboration with others. Finally, one participant described the 180° remapping as “*confusing (but necessary)*”.

3.4.2 When to Choose VR and When to Choose 2D?

The results of our user study make an important and central tradeoff apparent, which is to be considered when deciding whether interactive VR should be used to learn an EPC, or whether sticking to a conventional 2D representation as a graph is more suitable. We could show that understanding an EPC of the size of our test process can be completed significantly faster with a traditional 2D interface than in VR (**H1**) (see left chart in Figure 3.13). At the same time, however, the users' interest is significantly lower compared to using our proposed

immersive interface with proxy-based haptic feedback, as can be seen from the right chart in Figure 3.13 (H3). Depending on the scenario, an interface that does not catch the interest of users might lead to them being demotivated and not paying attention to the communicated information. This would be detrimental and counterproductive, for example, in educational scenarios, when an employee is familiarized with changes in a company process relevant to his work, or when results of a process optimization are presented to customers. Visual analysis of the right chart in Figure 3.14 further provides indication that the observed increase in user interest is not only a novelty effect of experiencing VR in general. While low participant counts disqualify a more detailed statistical analysis of the impact of VR experience on user interest, a trend of increased user interest in the *Props* condition is visible in Figure 3.14, independent of the participants' VR experience. The plots for non-experienced and experienced VR users both closely resemble the general tendency observed in Figure 3.13 (right). That being said, we also like to stress that our question for previous VR experience did not explicitly probe the participants' prior experience with haptic proxy interaction, leaving the impact of a potential novelty effect related to the use of passive haptics unclear and to be explored in future work.

Concerning the EPC learning performance, our study results could not show H2. However, it is interesting to note that concerning the actual performance of the users in understanding the depicted process, we could not detect any significant differences with our study of $N = 27$ participants between 2D and VR, as apparent from Figure 3.13 (center). Instead, we found very similar model understandability performances across conditions. While this does not prove the absence of performance differences and further investigation with higher participant counts is required in future work, the fact that differences did not become statistically striking with $N = 27$ provides initial support for the assumption that an immersive VR interface featuring proxy-based haptic feedback can be a suitable alternative to conventional 2D EPC interfaces in certain situations. It is this tradeoff between *efficiency* and *user interest* that represents the central finding of our evaluation.

While users could fall back to conventional 2D interfaces for time-critical EPC tasks, an immersive VR interface could be the first choice for less time-critical application scenarios such as presentations to customers, training and onboarding of employees, education, communication and related scenarios, to leverage the improved user experience. While not statistically significant, our results concerning pragmatic and hedonic quality as assessed by the UEQ and shown in Figure 3.14 provide further support for this, as do the qualitative comments of our participants in debriefing after the study.

3.4.3 The Impact of Passive Haptic Proxies

Our analysis of immersion did not yield statistically significant results and consequently we could not show H4, but visual analysis of the corresponding plot

on the left in Figure 3.14 indicates a tendency. The average SUS Count of participants in the *Props* condition was more than double the corresponding value in the *Controller* condition. Based on this result, we assume that proxy-based haptic feedback can increase the user's sense of presence in the tested application scenario – an assumption also supported by the findings of previous research [Insko, 2001; Hoffman, 1998]. Referring to the observed immersion ratings, we suggest to implement proxy-based haptic feedback when experience-focused VR is used as a data analysis tool, instead of a solution that is based solely on controllers, in order to maximize the system's immersion.

3.5 Conclusion & Contribution to the Research Questions

The investigation of process models is an important aspect in many professional domains, and 2D graph-like representations are the currently established standard interface for their exploration. As a result, with today's standard interfaces, only the visual perceptual channel of the user is involved. In this chapter, we explored a novel approach to explore and communicate business process model knowledge based on the principle of multisensory VR and investigated in the context of **RQ 1** how proxy-based haptic feedback can support this endeavor.

Guided by **RQ 1**, our contribution in this chapter is twofold: Firstly, we introduce the novel concept of *Immersive Process Model Exploration* and an implementation of this concept. By this, we bring forward novel solutions that help solving the *Similarity* and *Colocation* challenges. With our proposed system, we demonstrate how IVEs, haptic proxies, and interactions can be designed that provide an immersive, memorable, and interactive experience when exploring graph-based process model data in VR. To realize our solution, we build upon a combination of existing techniques, such as passive haptics, visual spatialization, gamification, and different locomotion and remapping techniques. Secondly, we present the results of a user study with $N = 27$ participants revealing the pivotal role of proxy-based haptics when it comes to retaining user interest in the process exploration.

For studying the potential of haptic proxies in this novel application domain, we implemented two levels of *Haptic Interactions* to increase the sense of presence when interfacing with the process components: vibrotactile feedback conveyed through standard VR controllers as a baseline, and proxy-based haptics. Our proposed system leverages physical proxies to augment the interactions with the information bits flowing through the process – highlighting the applicability and versatility of proxy-based haptic feedback. In this context, the utilization of the marble run principle using symmetrical, uneven, funnel-shaped or gutter-shaped props in combination with a tracked spherical prop is, to the best of our knowledge, a novel contribution to the field of proxy design. Specifically, our system demonstrates how to solve the challenge of *Similarity* by exploiting proxy symmetry and the great design freedom offered by the domain of abstract data visualization. Moreover, taking advantage of only a limited set of proxies

in combination with a 180° resetting controller, our implementation achieves a compression of arbitrarily large process models into a limited physical space while ensuring proxy *Colocation*. As such, the proposed interaction concept represents an evolution of classical approaches to passive haptics by addressing scalability and reusability issues.

In our concluding user study with $N = 27$ participants, we eventually compared the effect of our proposed VR interface on model understandability and user experience to a traditional 2D interface on a tablet device. Our results indicate a central tradeoff between *efficiency* and *user interest*, but did not indicate significant differences in model understandability performance across the tested conditions *iPad*, *Controller* and *Props*. Instead, and crucial to **RQ 1**, our evaluation could show VR in combination with proxy-based haptic feedback to significantly increase user interest compared to a standard process model exploration interface – at the cost of increased exploration time. Based on these results, we assume that multisensory and experience-focused data exploration interfaces in VR that leverage proxy-based haptic feedback, such as the presented system, can be suitable alternatives to established 2D interfaces in certain situations. We imagine such interfaces to be of particular value for less time-critical applications such as customer presentations, training, communication, education, and related scenarios with a focus on user interest.

Overall, this chapter highlighted the ease with which physical proxies enable haptic VR experiences that feature multiple haptic modalities and that let users perceive both the kinesthetic and tactile aspects of virtual interactions. We could demonstrate that with haptic proxies, the fast, low-cost, and low-complexity realization of a domain-specific system is achievable. Realizing a system such as the *Immersive Process Model Exploration* system while following the paradigm of AHF, in contrast, would likely take longer and would require significantly more costly and more complex equipment. Furthermore, comparable AHF-solutions would likely fail to provide the same level of multimodality achieved with our proposed proxy-based system. Hence, this chapter emphasizes the value that VR in conjunction with proxy-based haptic feedback bears for a wide range of application areas. As such, our work provides the motivating grounds for our research presented in the following three parts, which aims to improve the concept of using haptic proxies by enhancing the flexibility of props in order to push towards the realization of an ultimate display.

Part III

Enhancing Proxy-Based Haptics The *Physical Approach*



... a *dynamic* duo

Research Question 2	<i>Improvement</i>
<p data-bbox="652 833 965 871">This part addresses RQ 2:</p> <p data-bbox="406 878 1219 968">How can the gap in the haptics continuum be filled with a concept that enhances the flexibility of proxy-based haptic feedback and enables improved kinesthetic perceptions in VR with only minimal actuation?</p>	

Chapter 4

Dynamic Passive Haptic Feedback

In Part II, we focused on the benefits of haptic proxies and showcased their **Applicability** in the context of **RQ 1**. The findings of previous research on PHF reviewed in section 2.3 and the results of chapter 3 underline physical proxy objects to constitute a valuable tool for realizing multimodal haptic feedback for VR – a quality today’s AHF technology still widely fails to achieve [Wang et al., 2020a]. Passive haptic proxies, such as those employed in the system presented in Part II, provide tangibility and embody virtual objects while being low-cost, low-complexity, and mostly free of safety risks. However, following the paradigm of PHF, each new VR application that employs proxies, and each change of a virtual object inside existing applications, implies a manual adaptation of existing props, or the design and fabrication of new proxies. In other words, passive haptic proxies share one common disadvantage: the inflexibility of their feedback and the resulting lack of generality and *Versatility* [Münder et al., 2022]. This central drawback limits the set of virtual objects a proxy can appropriately represent if no compensatory techniques are applied and thus hinders the reuse of proxies across different applications and use cases.

Starting with this part, the following three parts of this thesis turn towards this central limitation of PHF and focus on the **Improvement** of proxy-based haptics. Specifically, our research in the following chapters pushes towards Sutherland’s vision of the ultimate display. We do so by contributing concepts, prototypes, techniques, and results that eventually grant the VR system more *control* over the proxy-based haptic feedback perceived by the user to evolve PHF to a technique where “*the computer can control the existence of matter*” [Sutherland, 1965]. Across Parts III and IV, we approach the problem of proxy inflexibility first with techniques operating in the physical environment, before we attend to solutions operating in the virtual environment. Finally, in Part V, we will investigate the potential of joining both physical and virtual strategies through

combined approaches. Throughout these parts, we propose novel concepts and solutions, and improve existing techniques that abstract away from individual application scenarios and are applicable in a wide variety of use cases. By this, we contribute important building blocks that increase the flexibility of haptic proxies and eventually reduce the number of proxy objects required to realize immersive VR systems that deliver multimodal haptic feedback.

To organize these efforts, our contributions concentrate on the two central challenges of proxy-based haptics introduced in section 2.4. As outlined in Figure 1.4 in the introduction of this thesis, we start our **Improvement** of proxy-based haptics in Part III by following **RQ 2**. Starting with this chapter, we first concentrate on the challenge of *Similarity* and contribute to the field of real-time physical techniques by introducing a novel category of mixed haptic feedback alongside two practical instances thereof.

This chapter is based on the following publication. Parts of the text in this chapter, as well as the presented concepts have been published previously therein:

Zenner, A. and Krüger, A. (2017). Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1285–1294. © 2017 IEEE. Final published version available in the IEEE Xplore® Digital Library. DOI: [10.1109/TVCG.2017.2656978](https://doi.org/10.1109/TVCG.2017.2656978)

4.1 Introduction

As our review of VR haptics in Part I made apparent, two gaps exist in the current research landscape that both appear relevant to advance the field of proxy-based haptics. **RQ 2** considers both of them. With the feedback concept that we will introduce in this chapter, we present a solution to both of the research gaps recapped in the following and thus to **RQ 2**.

4.1.1 The Gap in the Active-Passive Haptics Continuum

The first open research problem addressed by **RQ 2** is of conceptual nature and concerns the Active-Passive Haptics Continuum illustrated in Figure 1.3 and reviewed in chapter 2. Physical approaches to VR haptics can be classified within this continuum spanning from PHF to AHF, with techniques that combine actuation and proxies being classified as mixed haptic feedback [Jeon and Choi, 2009]. Conceptually, the haptics continuum is closely related to the Reality-Virtuality Continuum by Milgram and Kishino [1994] introduced in section 2.1, and likewise features sub-categories like Augmented Reality (AR) and Augmented Virtuality (AV). While the Reality-Virtuality Continuum classically distinguishes between AR and AV based on the degree of virtuality experienced by the user visually, haptic AR and haptic AV can be distinguished by the degree of actuation involved.

Within the field of mixed haptic feedback for VR, the concept of Encountered-Type Haptic Feedback (ETHF) reviewed in subsection 2.5.2 has already been established by previous research. Given its similarity to AHF, ETHF is commonly classified as haptic Augmented Virtuality (AV) when projected into the Active-Passive Haptics Continuum [Jeon and Choi, 2009; Lee et al., 2020], covering primarily the continuum's active half. A conceptual analogue that covers the passive half of the continuum, however, has been left undefined by previous research – motivating Part III of this thesis.

4.1.2 The Lack of Varying Kinesthetic Proxy Sensations

The second open research gap addressed by **RQ 2** is of technical nature. In conventional PHF setups, each proxy provides multisensory feedback including kinesthetic sensations such as those of mass, mass distribution, and inertia. Being passive, however, traditional proxies cannot easily render different kinesthetic impressions without being adapted manually (e.g., by weighting the props during or in between VR experiences). Apart from some early concepts discussed in section 2.5 [Schneider et al., 2005; Niiyama et al., 2014], practical solutions to vary kinesthetic sensations in proxy-based VR, i.e., solutions that vary impressions of weight, weight distribution, resistance, and inertia, are still lacking.

Apart from manually adapting a proxy's kinesthetic qualities, the most common approach for the simulation of different virtual weights and resistance effects is to simulate acting forces by rendering them to the user with techniques similar to those employed for AHF. VR systems that employ ETHF, for example, can simulate different kinesthetic properties of the proxies through robotic actuation. As a consequence, however, these solutions depend on powerful actuation mechanisms (e.g., sufficiently actuated robots) that *actively counteract forces applied by the user* – an approach that rather aligns with the paradigm of AHF than that of proxy usage. Hence, the second part of **RQ 2** aims at developing novel feedback concepts that can convey varying kinesthetic properties of virtual objects with a single proxy and are rooted in the paradigm of proxy-based haptics, i.e., do not require significant actuation.

4.1.3 Motivation

It is with the concept proposed in this chapter that we fill the conceptual gap inside the Active-Passive Haptics Continuum. Following **RQ 2**, we propose a class of mixed haptic feedback characterized by the primary use of proxy-based haptic impressions in combination with minimal actuation. By this, we complete the continuum and ease communication among researchers. Moreover, our concept enables novel haptic feedback techniques that deliver varying kinesthetic impressions of proxies as we will demonstrate in chapters 5 and 6.

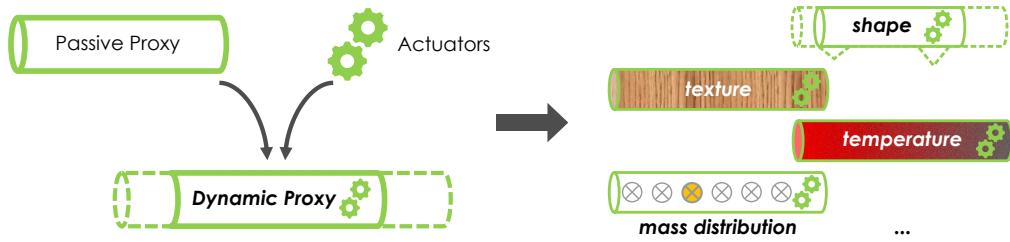


Figure 4.1: Dynamic Passive Haptic Feedback (DPHF) combines proxy objects with simple actuation to create dynamic proxies that can change their passive haptic qualities to simulate various virtual objects.

4.2 Definition

We introduce a novel class of haptic feedback for VR, called

Dynamic Passive Haptic Feedback (DPHF)

which mixes aspects of Active Haptic Feedback (AHF) and Passive Haptic Feedback (PHF) as sketched in Figure 4.1. Specifically, we propose to equip passive haptic proxies, devices, and environments, with actuating elements known from AHF systems to build *hybrids*, *i.e.*, *dynamic proxies*, that use actuators to change their passive haptic properties (*e.g.*, their size, shape, weight, weight distribution, texture, temperature, position, orientation, function, *etc.*), without exerting noticeable active forces on the user.

4.3 Discussion

By combining proxy objects with simple actuation, DPHF empowers VR systems to change the proxy's haptic qualities as illustrated in Figure 4.1. The concept combines the strengths of PHF (*i.e.*, high-quality, multimodal haptic feedback) with those of AHF (*i.e.*, flexibility) and uses actuation only to change the proxy's properties and not to actively render forces to the user that would simulate a quality of the virtual object. By this, DPHF promotes the idea of proxy-based



Figure 4.2: Updated version of Figure 1.3. By introducing *Dynamic Passive Haptic Feedback (DPHF)*, we filled in the conceptual gap in the Active-Passive Haptics Continuum.

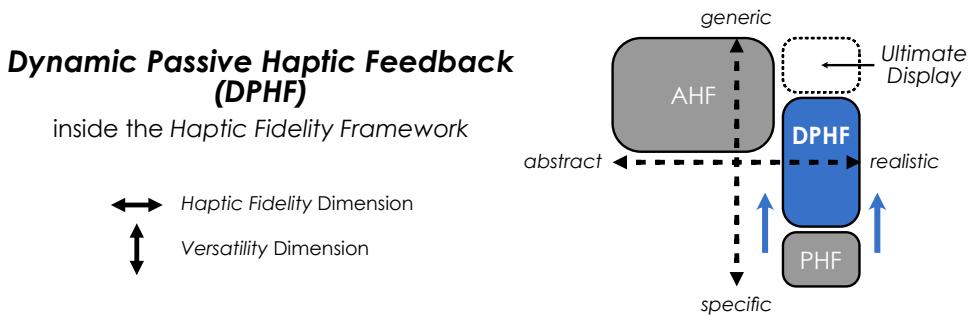


Figure 4.3: Classification of *Dynamic Passive Haptic Feedback (DPHF)* inside the *Haptic Fidelity Framework* by Münder et al. [2022]. By increasing the *Versatility* of haptic proxies, we advance towards the realization of an ultimate display.

sensations instead of synthetic force renderings and focuses on the passive half of the haptics continuum as indicated in Figure 4.2. The type and number of haptic properties a DPHF proxy could change is thereby not limited. Instead, DPHF proxies can be imagined that would change, for example, their shape, temperature, texture, etc., or any combination of those and other properties to establish *Similarity* with different virtual objects.

The concept of DPHF aims to increase the *Versatility* of haptic proxy objects, devices, and environments, while maintaining the high *Haptic Fidelity* [Münder et al., 2022] provided by physical proxies as sketched in Figure 4.3. DPHF systems only require knowledge about when to switch between different proxy configurations in order to simulate different virtual objects. Hence, VR systems that employ DPHF usually can spare complex physical simulations of acting forces as commonly required by solutions based on AHF or ETHF. As a consequence, DPHF, in contrast to ETHF, only requires very little knowledge about the virtual world at runtime and can thus be classified as a form of haptic AR. Table 4.1 summarizes the characteristics of DPHF analog to the overview of AHF and PHF provided in Table 2.2. Given these characteristics, DPHF fills the gap addressed by **RQ 2** as illustrated in Figure 4.2.

Our review in subsection 2.5.2 shows that since the introduction of DPHF in 2017, significant research efforts have been taken by the community to investigate a variety of novel feedback concepts based on dynamic proxies. These works cover different haptic dimensions by changing, for example, a proxy's stiffness [Murray et al., 2018; Stellmacher, 2020], shape [Zhao et al., 2017; Teng et al., 2018; Gonzalez et al., 2021a], temperature [Han et al., 2020], or texture [Whitmire et al., 2018; Degraen et al., 2020]. In the following two chapters, we will add to this research by introducing the haptic devices *Shift:y* and *Drag:on*. Both devices are designed to improve the perception of kinesthetic qualities in VR and can be classified as DPHF proxies since they fulfill the definition outlined above.

4.4 Conclusion & Contribution to the Research Questions

In this chapter, we started our **Improvement** of proxy-based haptics and contributed to the first part of **RQ 2** by introducing the novel mixed haptic feedback concept of *Dynamic Passive Haptic Feedback (DPHF)*. DPHF promotes the use of dynamic (instead of passive) haptic proxies, which leverage basic, low-power, and low-complexity actuation to adapt the passive haptic properties of props on-the-fly. By this, DPHF proxies can establish *Similarity* with a large number of virtual objects, enhancing the flexibility of proxy-based approaches at the expense of a slight increase in complexity. Dynamic adaptations of a prop can be of different nature and might concern only a single physical dimension, or combinations thereof. Yet, even when adapting only a single haptically perceivable property, a dynamic proxy is more flexible than an identical passive prop and might just offer enough variation in feedback to induce convincing visual-haptic illusions.

Dynamic Passive Haptic Feedback	
Generality	rather generalized
Interaction	direct with the hands
Feedback (Multimodality)	multimodal
Feedback (Variability)	variable
Supported Virtual Objects	static, interactable, (animated)
Costs (Equipment)	low
Costs (Computation)	low
Workspace	unlimited
Safety Risks	low
Tracking	partly required
Actuators	required
Embedded Electronics	required
Power Supply	required
Physical Properties of IVE	partly required at runtime
Body of Research	most since 2017
Classification¹	haptic augmented reality
Haptic Fidelity²	high
Versatility²	medium

Table 4.1: A high-level characterization of *Dynamic Passive Haptic Feedback* in comparison to AHF and PHF as outlined in Table 2.2. ¹ refers to the *Haptic Reality-Virtuality Continuum* by Jeon and Choi [2009]. ² refers to the *Haptic Fidelity Framework* by Münder et al. [2022].

Chapter 5

Shifty – Dynamic Passive Haptics

Based on Weight Shift

With the class of Dynamic Passive Haptic Feedback being introduced and the first part of **RQ 2** being covered, in this chapter, we turn towards the second part of **RQ 2**. Our **Improvement** of proxy-based haptics proposed in the following addresses the challenge of simulating virtual objects and interactions that differ in terms of their kinesthetic qualities. Specifically, we take advantage of DPHF and introduce a novel feedback concept based on a dynamic proxy designed to vary its physical properties. In addition, we present an implementation of this concept in the form factor of a handheld haptic VR controller and evaluate it in two user experiments.



Video Link⁵⁷

A video⁵⁷ about the work presented in this chapter is available online. This chapter is based on the following publication. Images and parts of the text in this chapter, as well as the presented concepts, implementations, and results have been published previously therein:

Zenner, A. and Krüger, A. (2017). Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1285–1294. © 2017 IEEE. Final published version available in the IEEE Xplore® Digital Library. DOI: [10.1109/TVCG.2017.2656978](https://doi.org/10.1109/TVCG.2017.2656978)

⁵⁷Shifty Video. <https://bit.ly/3x6F9Ae>

5.1 Introduction

When interacting with objects in the real world in our daily life, we constantly perceive a broad range of haptic cues that help us understand the object's physical qualities. Such haptic perceptions are essential for a safe, precise, and effective interaction with our surrounding. Among the most prominent cues that we get from interactions, such as touching and lifting an object, are the perception of tactile aspects as well as kinesthetic information about shape, weight, weight distribution, and material. When simulating virtual objects haptically with props, these properties can be replicated during the proxy fabrication stage. However, once a passive proxy matching a virtual object is built, changing those properties that contribute to our kinesthetic perception of the object, such as its weight, shape, or material, is difficult and usually results in the fabrication of an entirely new proxy. As a consequence, the question arises how dynamic proxies following the DPHF paradigm can be constructed that can change their kinesthetic properties so as to embody various virtual objects that differ, for example, in their virtual weight.

While sophisticated AHF systems specialized in the simulation of virtual forces have been developed in the past, the haptic solutions that come with today's consumer VR systems still rely entirely on handheld controllers (such as the HTC Vive Pro Controller⁵³, or the Oculus Quest 2 Controller⁵⁸), which are designed with mechanical simplicity in mind. These controllers serve as multipurpose proxy objects and are to embody all the virtual objects the user interacts with – even across applications and use cases. As such, handheld VR controllers constitute a challenging field for proxy-based haptics, calling for proxy designs that stand out in terms of their reusability and flexibility.

Yet, when inspecting the current state of the field, it becomes apparent, that the standard controllers found today do not appropriately meet this requirement. In terms of haptics, they act as proxies for virtual tools, weapons, and other items, but offer only a fixed shape, mass, and material impression. Moreover, concerning their weight, they are designed with minimum fatigue in mind and thus lightweight. As a result, their haptic simulation of large and heavy virtual objects turns out to be especially unrealistic. Picking up and moving virtual objects that visually imply a significant weight results in unrealistic, balloon-like impressions and unexpected object handling, which reminds users of the fact that they only experience a deficient simulation and increases the risk of BIPs.

As outlined in section 2.2, VR controllers usually aim to substitute kinesthetic impressions, like virtual weight, with basic vibrotactile patterns to compensate for their inability to convey true kinesthetic cues [Jerald, 2015; pp. 304 ff.]. For this, the devices rely primarily on built-in vibration motors. The vibrations they produce typically vary in vibration strength and frequency, but fail to appropriately stimulate the kinesthetic receptors located at deeper levels inside our body,

⁵⁸Oculus Quest 2 Controller. <https://ocul.us/3nqultH>

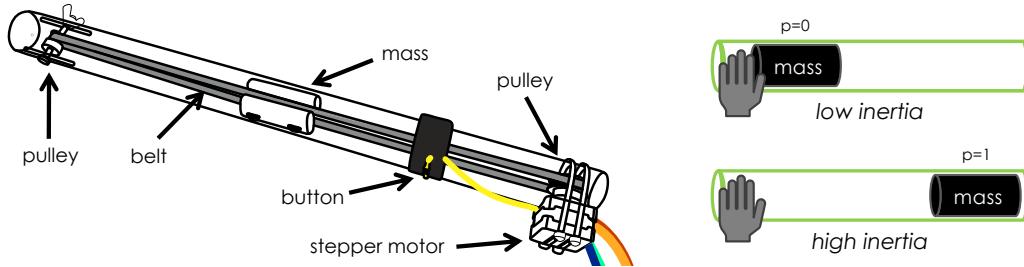


Figure 5.1: Sketch of *Shifty* with its main components (© 2017 IEEE) and inertial states.

such as in our muscles and tendons in the arms. This limitation of today’s VR systems highlights the need for research on novel feedback concepts compatible with the form factor of handheld VR controllers yet able to convey different kinesthetic cues [Lim et al., 2021]. With the work in this and in the following chapter, we contribute two proxy designs that fulfill these requirements.

5.2 Concept of *Shifty*

In the following, we introduce the DPHF proxy *Shifty*. The goal of *Shifty* is to enhance the perception of kinesthetic aspects of virtual objects and interactions. To achieve this, *Shifty* leverages the characteristics of human perception as discovered by previous research on dynamic touch (introduced in subsubsection 2.2.3) [Chan, 1995; Turvey, 1996; Kingma et al., 2004].

Shifty is designed as a rod-shaped proxy object. In contrast to traditional *passive* proxies, *Shifty* can *alter* its physical properties through an actuating motor, controlled by the VR system to haptically represent a large set of different virtual objects. The user holds *Shifty* with one hand at the grip end and by *shifting an internal weight* between the grip and top end, it can translate its Center of Mass (CoM) and change its rotational inertia. Figure 5.1 shows a concept sketch. Shifting an internal weight towards the top of the object and thus away from the rotational axes passing through the user’s wrist results in an increase of the rotational resistance. This effect does not change the absolute weight of the proxy, but results in the user having to apply stronger forces in order to move the object. Specifically, taking advantage of the physics of inertia, *Shifty* increases (decreases) its moment of inertia I by shifting mass inside the proxy away (towards) the hand holding it at one end. At the same time, *Shifty* also increases (decreases) its static moment $M = m \cdot d$ by increasing (decreasing) the distance d of the CoM to the point of rotation located at the user’s wrist.

We leverage these effects and claim that the slow and continuous change in *Shifty*’s inertial configuration can, when synchronized with appropriate visual and auditory feedback, change the user’s perception of the linked virtual object. *Shifty* is expected to be able to perceptually induce illusions of different object shapes (e.g., different object lengths according to Turvey [1996] and Kingma et al.



Figure 5.2: A user interacting with the DPHF proxy *Shifty*. *Shifty* can change its internal weight distribution to adapt its passive haptic feedback. By this, *Shifty* is designed to enhance the perception of virtual objects and their kinesthetic qualities. © 2017 IEEE.

[2004] due to changes in I and M) and object weights (according to Kingma et al. [2004] due to changes in M). More specifically, we claim that by increasing both I and M of the proxy, users believably perceive the virtual object they interact with as becoming heavier, thicker, or longer, depending on the visual feedback. Conversely, by decreasing I and M , we expect the virtual object to be perceived as becoming lighter, shorter, or thinner.

Equipped with a pushbutton and means for spatial tracking, *Shifty* serves as a basic VR controller that enables users to interact with the IVE through the most common interaction patterns. One of the most relevant thereof being picking up other virtual objects by pushing and holding the button on the device. In this regard, we further claim that the perceived realism can be increased significantly by changing the inertial state of the VR controller when picking up virtual objects. We will investigate these claims in two experiments.

5.3 Implementation of *Shifty*

Shifty consists of a lightweight body and remotely controlled internal mechanics that linearly displace a mass along this body. The position of the mass is defined as $p \in [0, 1]$ where $p = 0$ means the mass is at the grip end and $p = 1$ means it is located at the top end of the proxy. Grasping the proxy at the grip end with one hand, the user can use a pushbutton beneath the index finger to trigger actions in the IVE. To be compatible with large-scale VR systems and walking techniques, *Shifty* was designed to be mobile. Hence, the power supply and controlling electronics of the prototype used in the evaluation of this chapter are all built into a small backpack and the actual proxy is connected to this backpack via cables. The cable length allows for unrestricted arm movement and *Shifty* communicates wirelessly with the VR system.

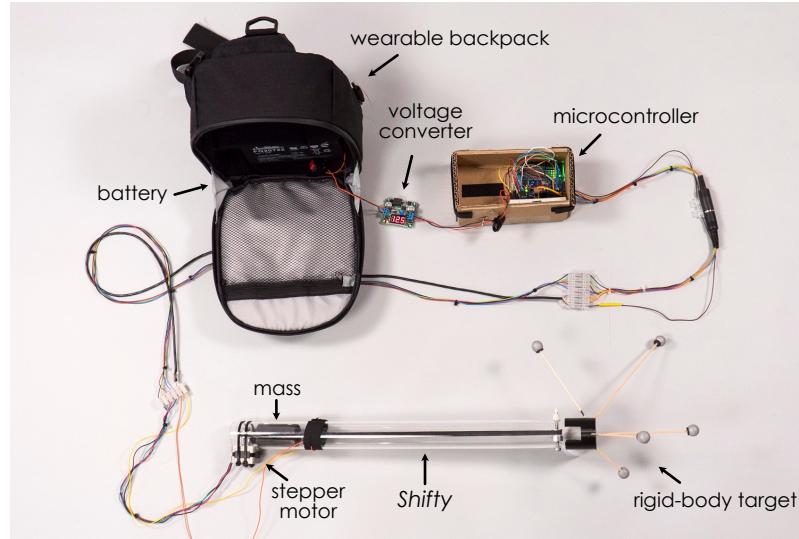


Figure 5.3: *Shifty* components: A backpack containing the battery and electronics, cables connecting backpack and proxy, and *Shifty* with rigid-body tracking target.

5.3.1 Hardware

The prototype of *Shifty* is depicted in Figure 5.2. A lightweight plexiglass pipe (length = 505mm, wall thickness = 2mm, diameter = 40mm) is used for the body and a low-cost and lightweight NEMA-14 type stepper motor is fixed with two ties at the grip end of the pipe. The motor actuates an internal belt system and as an inherent by-product produces slight audible and haptic noise when active. Inside the pipe, an aluminum pulley is fixed on the motor's axis. At the top end, two slots of 55mm length and 5mm width are cut into the pipe. Passing through these two slots, 450mm away from the motor axis, a bolt with a diameter of 5mm is fixed with a wing nut. Inside the pipe, this top bolt carries an aluminum pulley on small bearings, allowing it to spin without too much friction. Both pulleys and the toothed belt are widely available parts typically used in 3D printers. The internal weight is a custom-designed 3D-printed object of 60mm length and 33mm width. It contains four chambers filled with lead and has recesses for the belt. It is fixed on the belt and carries four bearings as wheels to minimize friction at the pipe's wall. Figure 5.4 shows the assembled prototype. In total, the proxy weighs 440g including a moving weight of 127g, which gives a *moving weight : total weight* ratio of $\frac{127g}{440g} \approx 0.29 = 29\%$.

When the internal weight is shifted completely towards the grip end (for $p = 0$), the proxy's CoM is located 13.6cm from the pipe's bottom end. Shifting the weight with a step-resolution of $0.39 \frac{mm}{step}$ over the complete range of 36.5cm takes around 2.8s with a speed of $\approx 0.13 \frac{m}{s}$. If shifted maximally towards the top (for $p = 1$), the CoM is at a distance of 24.5cm from the pipe's bottom end. For each $p \in [0, 1]$, the CoM is thus interpolated between these two locations, covering a range of approximately 11cm.

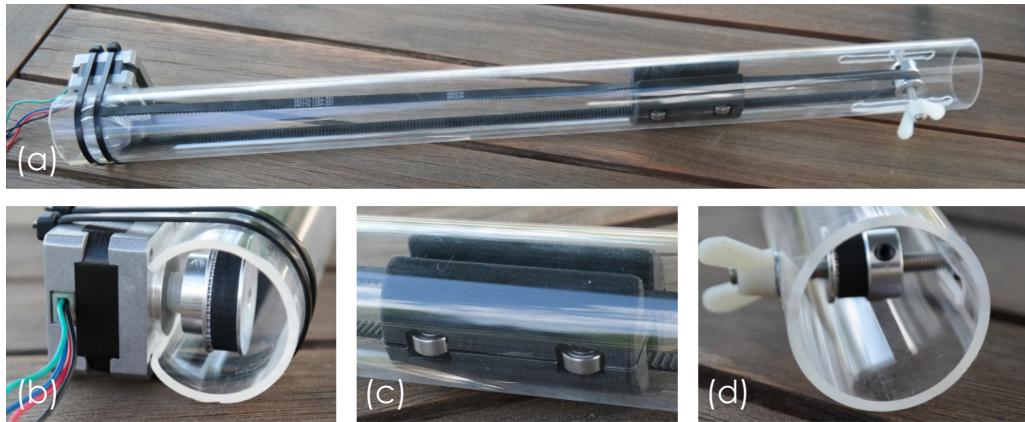


Figure 5.4: *Shifty* prototype: (a) the assembled proxy without the button and tracking target, (b) NEMA-14 type stepper motor and grip end pulley with the belt, (c) 3D-printed weight filled with lead and bearings, (d) top end pulley. All images ©2017 IEEE.

For user input, we fixed a small pushbutton, such as used in car keys, on the outside of the proxy with a stretchable band and velcro fastener. This way, we could adjust the location of the button on the proxy to the user's hand size. The proxy controller software runs on an Arduino microcontroller placed in a small backpack with a motor shield stacked on it. A 12V rechargeable battery in the backpack serves as the power supply. To track the prototype, we use an OptiTrack system⁵⁹. A custom-designed rigid-body target made out of a 3D-printed plug is attached to the top end of the pipe. It holds five wooden sticks with fixed reflective markers.

5.3.2 Software



GitHub Link⁶⁰

The C++ software controls the stepper and handles communication with the VR system. For this, the Arduino either connects to an existing WiFi network or opens a dedicated WiFi hotspot. Commands are sent by client applications using a custom C# Application Programming Interface (API) via network to the Arduino. Conversely, button events are communicated via TCP back to connected clients. The API offers a set of functions to move the weight in the proxy and to retrieve predictions of the transformation time.

All relevant resources, such as a detailed list of parts, relevant 3D model data for 3D printing, source code, as well as a construction manual can be found online in an open-source repository⁶⁰ allowing interested readers to recreate *Shifty*.

⁵⁹OptiTrack Motion Capture Systems. <https://bit.ly/3FwKCmN>

⁶⁰*Shifty* on GitHub. <https://github.com/AndreZenner/shifty>

5.4 Evaluation of *Shifty*

To evaluate the concept of weight shifting DPHF, we conducted two experiments with our *Shifty* prototype.

5.4.1 Experiment 1: Change in Object Length and Thickness

In a first experiment, we investigated how a DPHF proxy like *Shifty* can enhance the perception of virtual objects that change in shape. For this, we put users in an IVE and let them interact with two objects changing their length and thickness, respectively, and assessed their impressions of realism, enjoyment, and exertion during the interaction. In the physical environment, participants interacted alternately with our weight-shifting proxy and with our proxy holding the internal weight stationary at the grip end, mimicking an identical but passive haptic prop.

Based on this setup, we compare the perception of the virtual objects when using *Shifty*, i.e., DPHF, with the perception achieved through PHF. We hypothesize:

H1 DPHF will increase the perceived realism and enjoyment compared to PHF.

H2 DPHF will result in greater exertion compared to PHF.

Participants

12 volunteer participants took part in the first experiment (5 female, 7 male, avg. 28 years, between 21 years and 37 years old). 5 of them wore glasses or contact lenses during the experiment and 9 of them were right-handed, while 3 were left-handed. We also asked participants how regularly they play 3D video games on a scale from 1 (= never) to 7 (= regularly). The results showed that all types of gaming behavior were represented and the average score was 3.5. Moreover participants were asked to rate their previous experience with VR technology on a scale from 1 (= never used) to 7 (= regular use), and the average score of all participants was 1.6 with answers between 1 and 3.

Apparatus

The experiment was carried out in our lab with the *Shifty* prototype. We used a laptop to record the participants' answers and to run the OptiTrack Motive tracking software to track *Shifty*'s rigid-body target. Moreover, an HTC Vive¹⁴ HMD was used to track the participant. The HMD was additionally equipped with three infrared reflecting markers used to translate between the Vive's and OptiTrack's coordinate systems. To provide auditory feedback and to minimize the perceived noise of *Shifty*'s motor, all participants wore over-ear headphones. The IVE and the experiment were implemented using the Unity 5.3 engine¹⁷,

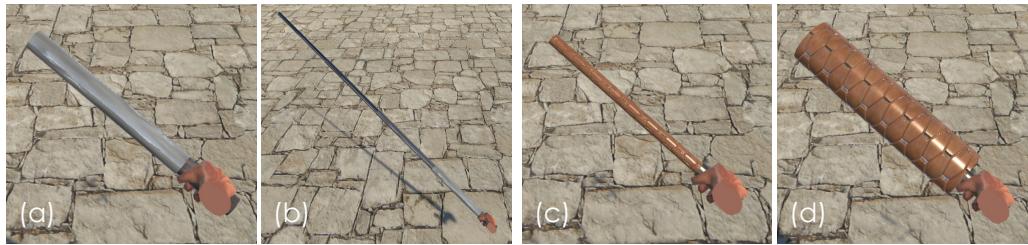


Figure 5.5: The objects of Experiment 1 changing in length and thickness: (a) the telescope object at 50cm length and (b) at 200cm length, (c) the rod-shaped object in the thinnest configuration and (d) in the thickest configuration. All images © 2017 IEEE.

which ran on a powerful desktop computer. The data of the tracked rigid bodies was streamed with low latency via a custom middleware and a local network from the laptop to the rendering machine. *Shifty* connected via WiFi to this local network and communicated with the Unity engine via the proxy’s API. During the experiment, participants stood in the center of the tracked space. While they could move freely, they did not need to walk around in the IVE.

Figure 5.2 shows a user interacting with *Shifty* and Figure 5.6 depicts our setup. For the study, instead of a 3-camera OptiTrack V120:Trio, a 6-camera OptiTrack rig was used to ensure robust tracking and a large capture volume.

Procedure

Each participant was informed about the course of the experiment and their tasks in the IVE. *Shifty* was not shown to the participants before the experiment.

The experiment itself consisted of two phases. In each phase, *Shifty* was used once changing its weight distribution as introduced, and once always holding its internal weight at the grip end. By holding the weight stationary, it served as a conventional PHF proxy.

The first phase of the experiment was concerned with the user’s perception of a virtual object continuously changing in length. Here, participants interacted with a virtual telescope that could smoothly extend and retract. The virtual telescope changed its length in four steps. Participants saw a floating virtual cube marked with a “+” as well as a cube marked with a “-” symbol. By intersecting one of the two cubes, the virtual telescope extended or retracted one step, respectively. In total, the virtual length changed from 50cm to 200cm in four equidistant steps. Figure 5.5 depicts the completely retracted (a) and extended (b) telescope. The actual task of the participants was to stepwise extend the telescope from 50cm to 200cm and to freely swing and wield the telescope at each step. After that, participants had to stepwise retract it again. Finally, they were asked questions about the perceived realism, exertion, enjoyment, and their personal preference. As this procedure was performed once with DPHF and once with PHF, the only difference between both runs was whether the passive haptic feedback changed



Figure 5.6: A user interacting with *Shifty* in our experiment setup. The user wears an HMD and headphones. The rigid-body target on the proxy is tracked by a motion capture system. © 2017 IEEE.

or whether it stayed fixed. To exclude ordering effects, the order of DPHF and PHF runs alternated between participants.

The second phase was designed equivalently. Here, however, the virtual object did not change its length, but its thickness above the grip. It sized up in four steps from 200cm^3 to 3000cm^3 . Figure 5.5 shows the thinnest (c) and thickest (d) state of the virtual object. The second phase was also conducted once with changing haptic feedback and once with constant passive feedback. At the end, the same set of questions was asked here again.

Design

The first experiment was designed as a within-subjects experiment. For each of the two independent phases, a Latin square for $n = 2$ was used six times to counterbalance the order of the haptic feedback modes. The independent variable was the type of haptic feedback used. We tested two conditions: *DPHF* against *PHF*, or in other words we tested *using a proxy with changing weight distribution* against *using a proxy with constant weight distribution* as a baseline condition. The dependent variables were the perceived realism, combined mental and physical exertion, and enjoyment, all assessed as self-reported absolute values on a 7-point Likert scale and as direct comparisons between both types of feedback. Additionally, we also asked for the participants' personal preference in the form of a direct comparison between DPHF and PHF. A final question asked participants to rate on a 7-point Likert scale how much they perceived the physical proxy to really change its length or thickness with the virtual object after experiencing the DPHF feedback.

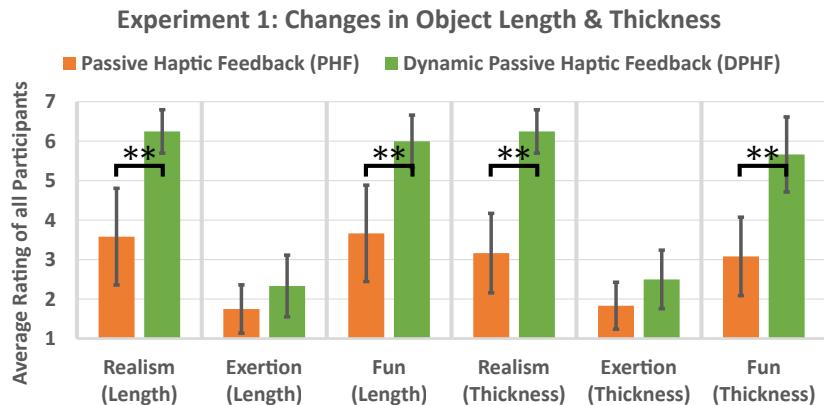


Figure 5.7: Average ratings for perceived realism, exertion, and enjoyment of PHF (orange) and DPHF with *Shifty* (green). The three comparisons on the left relate to the virtual object changing in length, while the three comparisons on the right relate to changes in virtual thickness. Brackets indicate statistically significant differences ($p < .05$ (*); $p < .01$ (**)). Error bars show 95% confidence intervals. © 2017 IEEE.

Results

The results of Experiment 1 are presented in the following. We outline the findings concerning changes in virtual object length first, before presenting the results for changing thickness.

Changing Length The results of the absolute ratings for the perceived realism, exertion (mental and physical demand combined), and enjoyment are summarized in Figure 5.7. Using Wilcoxon signed-rank tests, we compare the average ratings of each dependent variable for DPHF, i.e., using weight shift, against PHF, i.e., using a proxy with fixed weight distribution.

Regarding realism, participants were asked to rate how realistic they perceived the interaction with the object to be on a 7-point Likert scale from 1 (= very unrealistic) to 7 (= very realistic). According to a Wilcoxon signed-rank test ($Z = -2.814$, $p = .002$, $r = .81$), the difference in the perceived realism between the two conditions DPHF ($Mdn = 6.50$, $M = 6.25$, $SD = 0.87$) and PHF ($Mdn = 3.50$, $M = 3.58$, $SD = 1.93$) is significant on a significance level of $\alpha = 0.05$. The results for enjoyment on a scale from 1 (= none) to 7 (= very much) were very similar. On the same significance level, a Wilcoxon signed-rank test ($Z = -2.781$, $p = .004$, $r = .80$) found the enjoyment ratings to differ significantly between the DPHF condition ($Mdn = 6.00$, $M = 6.00$, $SD = 1.04$) using weight-shifts and the PHF condition with a fixed weight distribution ($Mdn = 4.00$, $M = 3.67$, $SD = 1.92$). With the increasing realism, the DPHF results for mental and physical exertion on a scale from 1 (= not at all) to 7 (= very exertive) ($Mdn = 2.00$, $M = 2.33$, $SD = 1.23$) also increased compared to the PHF exertion ratings ($Mdn = 1.50$, $M = 1.75$, $SD = 0.96$). However, a Wilcoxon signed-rank test ($Z = -1.511$, $p = .250$, $r = .44$) did not indicate this difference to be significant.

Besides rating the dependent measures, each of the 12 participants was asked to directly compare the conditions DPHF and PHF with respect to the factors realism, exertion, enjoyment, and personal preference. Here, participants consistently favored DPHF over PHF with regard to the realism of the haptic feedback (DPHF = 11, both equal = 0, PHF = 1), fun (DPHF = 11, both equal = 1, PHF = 0), and personal preference (DPHF = 11, both equal = 0, PHF = 1). Consistent with the absolute ratings, 7 out of 12 participants perceived DPHF as more physically and mentally demanding while no one perceived PHF as requiring more exertion. 5 participants perceived PHF and DPHF as equally exertive.

When the participants were asked how strongly they felt that the object in their hand really changed its length, when in fact the proxy shifted its internal weight, a strong feeling was recorded. Participants could rate on a scale from 1 (= not at all) to 7 (= very strong feeling). The obtained average score ($Mdn = 6.50$, $M = 6.08$, $SD = 1.08$) was high.

Changing Thickness The results obtained for the perception of the virtual object changing in thickness are similar to those obtained for changing length. A summary of the absolute ratings can be seen in Figure 5.7 as well.

As for changing length, the difference between the perceived realism of the DPHF condition ($Mdn = 6.00$, $M = 6.25$, $SD = 0.87$) and the PHF condition ($Mdn = 4.00$, $M = 3.17$, $SD = 1.59$) was statistically significant according to a Wilcoxon signed-rank test ($Z = -2.952$, $p = .001$, $r = .85$) for $\alpha = 0.05$. The results on enjoyment also matched the results of the first phase. DPHF was rated to be significantly more fun to interact with ($Mdn = 6.00$, $M = 5.67$, $SD = 1.50$) than PHF ($Mdn = 3.50$, $M = 3.08$, $SD = 1.56$), as the results of a Wilcoxon signed-rank test show ($Z = -2.915$, $p = .002$, $r = .84$). Although again slightly lower than the results of DPHF ($Mdn = 2.00$, $M = 2.50$, $SD = 1.17$), the exertion ratings for PHF ($Mdn = 2.00$, $M = 1.83$, $SD = 0.94$) were not found to differ significantly according to a Wilcoxon signed-rank test ($Z = -1.725$, $p = .156$, $r = .50$).

The similarity to the results concerning length-changing objects continues when it comes to the direct comparison of DPHF and PHF. Participants again consistently favored DPHF over PHF regarding the realism of the haptic feedback (DPHF = 11, both equal = 0, PHF = 1), fun (DPHF = 10, both equal = 0, PHF = 2), and personal preference (DPHF = 10, both equal = 0, PHF = 2). 8 participants perceived DPHF as more physically and mentally demanding while 2 participants stated PHF required more exertion and 2 others found PHF and DPHF to be equally exertive.

When participants rated how strongly they felt the object in their hand changing its thickness, when in fact the proxy shifted its weight, a slightly lower average value ($Mdn = 6.00$, $M = 5.58$, $SD = 1.62$) was obtained. The difference from the rating in the first phase, however, was not found to be significant according to a Wilcoxon signed-rank test ($Z = -0.641$, $p = .586$, $r = .18$).

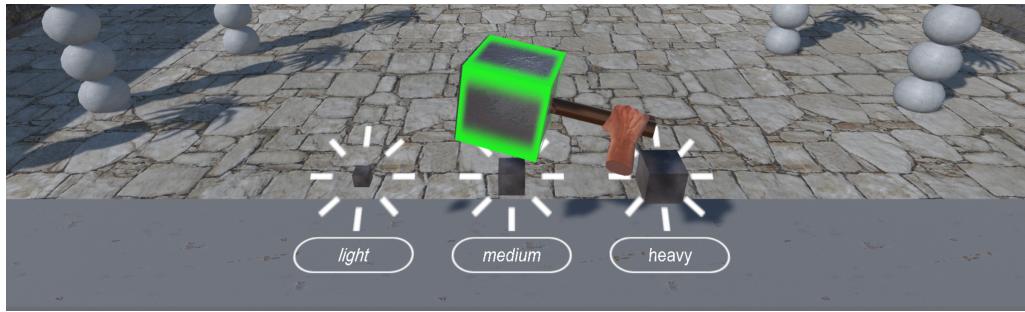


Figure 5.8: Holding the heavy cube into the target area in Experiment 2. © 2017 IEEE.

5.4.2 Experiment 2: Picking Up Virtual Objects

While Experiment 1 investigated how *Shifty* can enhance the perception of objects that gradually change over time, Experiment 2 is concerned with instantaneous events. In particular, we investigate how the DPHF of *Shifty* can enhance one of the most crucial interactions in VR: picking up a virtual object. In most VR applications this requires users to bring the VR controller close to the virtual object and to then press a button snapping the object to the controller. By holding the button pressed, the virtual object stays attached, and releasing the button drops the object. In simulations or games, the controller is hereby often visualized as a virtual hand or a virtual object similar in shape to the physical device.

Using *Shifty* as a VR controller, we leverage its weight-shifting capabilities to enhance the feeling of virtual weight. In our experiment, participants pick up a light-, a medium-, and a heavy-looking object from a virtual inventory as depicted in Figure 5.8. With the picked-up object they then solve a simple docking task by holding it into a target area. As we want to compare *Shifty*'s DPHF to the PHF provided by conventional VR controllers, we again mimic an equivalent PHF controller with *Shifty*. Similar to the first experiment, we compare a PHF baseline condition, in which *Shifty* is holding the weight fixed at the grip, against five different DPHF conditions involving *Shifty*'s dynamic feedback. Compared to each other, these five conditions differ only in their visual and auditory feedback. By testing five different visualizations of the pick-up process, we try to find one that can compensate for the visual-haptic mismatch arising when *Shifty*'s weight is shifted for up to 2.8s during pick-up. For this we again assess the participants' perceived realism, enjoyment, personal preference, and exertion. Additionally, we assess how disturbing the different conditions are with respect to immersion. Our hypotheses can be summarized as follows:

- H3** DPHF will increase the perceived realism, enjoyment, exertion, and preference compared to PHF.
- H4** DPHF with enhanced visualizations during pick-up will result in greater perceived realism and enjoyment, while being perceived as less disturbing and less exertive compared to DPHF with a standard visualization.

	Condition	Haptics	Sync. Visuals	Sync. Audio
<i>Base</i>	Baseline	PHF	None	None
<i>Hapt</i>	Haptic-Only	DPHF	None	None
<i>Prog</i>	Progress Bar	DPHF	Progress Bar	<i>Whoosh</i>
<i>Scal</i>	Scaling	DPHF	Scaling Up	<i>Whoosh</i>
<i>Trans</i>	Transparency	DPHF	Becoming Opaque	<i>Whoosh</i>
<i>Mask</i>	Masking	DPHF	Smoke Mask	<i>Whoosh</i>

Table 5.1: Overview of the conditions in Experiment 2. © 2017 IEEE.

Participants

12 volunteer participants took part in the second experiment (3 female, 9 male, avg. 27 years, between 21 years and 37 years old). Half of the participants wore contact lenses or glasses and 10 were right-handed while 2 were left-handed. Compared to the first experiment, the average gaming experience was slightly lower, with a score of 2.9, and the VR experience was slightly higher with an average rating of 1.8. Here, answers between 1 and 4 on the same 7-point Likert scale as in the first experiment were recorded.

Apparatus

The setup used for Experiment 2 was equivalent to that of Experiment 1.

Procedure

As in the first experiment, *Shifty* was not shown to the participants before the second experiment and participants were briefed about the course of the study. Each participant experienced six different conditions in succession. To account for ordering effects, the order of these conditions was counterbalanced using a Latin square. For each condition, the participant's task was to pick up the light, the medium, and the heavy virtual object, and to hold it into a highlighted target area for a duration of 1s. This ensured that the participants had comparative experiences for each condition. Starting with the completion of the second condition, questions were asked after the completion of all following conditions. In these questions, participants directly compared the last two experienced conditions. Here, we asked participants to state in which of the last two conditions the interaction with the objects after picking them up felt more realistic and in which the pick-up interaction was perceived as less disturbing regarding immersion. We further asked which condition was less exertive, which one was more enjoyable, and finally, which one they would personally prefer. It was also valid to rate both as equal. After all conditions were experienced, participants were asked for their personal overall favorite condition.

All six conditions are summarized in Table 5.1. The PHF baseline condition (*Base*) represents the current state of VR controller interaction. It does not involve a

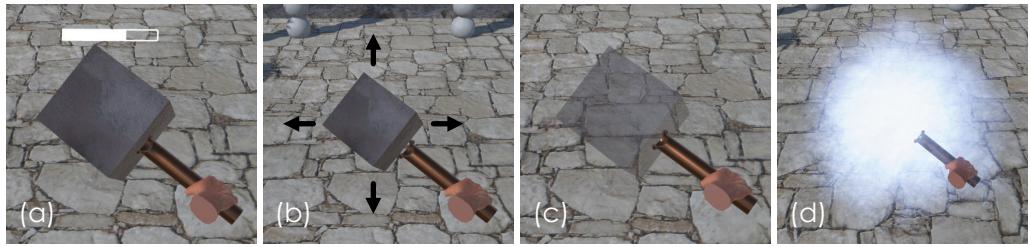


Figure 5.9: The four different visualizations tested in Experiment 2: (a) The progress bar (*Prog*), (b) the scaling animation (*Scal*), (c) the transparency transformation (*Trans*), and (d) the smoke mask (*Mask*). All images © 2017 IEEE.

change in the kinesthetic haptic feedback nor does it involve any special visual or auditory feedback. A second condition did not involve any special visual or auditory feedback either, but did use *Shifty*'s DPHF. We refer to this condition as the haptic-only (*Hapt*) condition in the following. Besides *Base* and *Hapt*, four further conditions were tested. Like *Hapt*, they all employed *Shifty*'s DPHF. In all DPHF conditions, *Shifty* adapted its mass distribution based on the virtual weight of the object that was picked up. Specifically, *Shifty*'s internal weight moved to $p = 1$ when picking up the heavy object, to $p = 0.5$ when picking up the medium object, and to $p = 0.1$ for the light object. Each of these four animation conditions involves the same auditory feedback combined with a different synchronized visual animation effect. The auditory feedback in all animation conditions is a *whoosh*-like sound that was played synchronously with the shifting weight. Visually, one condition displayed a progress bar showing the progress of *Shifty*'s weight shift when objects were picked up. We call this condition the progress condition (*Prog*). A second animation condition, the scaling condition (*Scal*), made objects scale up from the inventory-icon size to the object's actual size when picked up, synchronized with *Shifty*'s weight shift. A third condition, the transparency condition (*Trans*), transformed the object's transparency from transparent to opaque, and the masking condition (*Mask*) visually masked the picked-up object by displaying a thick smoke field around the object. The smoke only disappeared when *Shifty*'s weight shift was finished. Figure 5.9 shows screenshots of these four visualizations.

Design

The second experiment was also designed as a within-subjects experiment. With six different feedback conditions, a Latin square for $n = 6$ was used twice to counterbalance the order of the conditions. The independent variable was the combination of haptic, visual, and auditory feedback. Six combinations were tested: the conditions *Base*, *Hapt*, *Prog*, *Scal*, *Trans* and *Mask* as introduced. The dependent variables were the obtained measures regarding the perceived realism during interaction with the objects, the disturbing influence on immersion during pick-up, the exertion, the enjoyment, and the personal preference.

Realism			Least Disturbing			Exertion			Fun			Preference		
#	Cond.	Σ	#	Cond.	Σ	#	Cond.	Σ	#	Cond.	Σ	#	Cond.	Σ
1	Scal	29	1	Scal	25	1	Base	11	1	Mask	30	1	Mask	30
2	Trans	27	2	Mask	22	2	Trans	17	2	Scal	26	1	Scal	30
3	Prog	23	3	Base	21	3	Scal	21	3	Prog	22	2	Prog	18
4	Mask	20	4	Trans	18	4	Hapt	23	4	Trans	17	3	Trans	17
5	Hapt	18	5	Prog	17	4	Mask	23	5	Hapt	16	3	Hapt	17
6	Base	3	5	Hapt	17	5	Prog	25	6	Base	9	4	Base	8

Table 5.2: Final ranking tables for all measures in Experiment 2. © 2017 IEEE.

The measures were obtained as a set of direct comparisons. Since participants always compared the last two experienced conditions, they were asked each of the five comparison questions five times. For the evaluation, we define a direct comparison of two conditions as a *match* played between these two conditions. Due to the Latin square design and the amount of participants, each condition played four times against each other condition, twice experienced before the compared condition and twice after. As a result a complete ranking table is computed for each dependent variable, i.e., each of the five questions. When a condition wins a direct comparison question, its score for the considered measure is increased by 2 points, while the losing condition's score stays the same. If two conditions are rated as equal, both score 1 point for the compared measure. After all 60 comparisons were recorded, all points scored by a condition were summed up and a final ranking table was computed for each measure. In this way, general tendencies towards a certain condition can be identified as favored conditions are more likely to win more comparisons than less-favored conditions. Since the second experiment focuses on the qualitative assessment of user preferences, we will discuss the resulting ranking, providing insights into which conditions are generally preferred or considered worse with respect to the five dependent measures. A more in-depth investigation and analysis is left to future work.

Results

The results of the second experiment are depicted in Table 5.2 summarizing the final ranking tables for all five dependent measures. The scores are computed as described in the previous section, adding 2 points when a match against another condition was won, 0 points if lost and 1 point for each condition, if two conditions were rated as equal.

The scaling condition *Scal* scored best concerning the perceived realism of the interaction after picking up a virtual object. *Scal* was also classified as impairing immersion least during the pick-up interaction and, together with the masking condition *Mask*, scored best concerning the participants' preference. Regarding the enjoyment during the interaction, *Mask* scored highest. Concerning exertion, the baseline condition *Base* scored best, i.e., was rated as being least exertive.

In the concluding question participants had to state their overall favorite con-

dition. Here, 7 out of 12 participants chose the scaling condition *Scal* as their favorite, 2 chose the progress bar condition *Prog* and 2 others chose the haptic-only condition *Hapt*. 1 participant voted for the masking condition *Mask*.

5.5 Discussion of *Shifty*

Our experiments showcased how weight-shifting DPHF can enhance the perception of virtual interactions and the kinesthetic properties of virtual objects. In the following, we discuss the findings of our two experiments, before we review the relationship of *Shifty*'s actuation and its location inside the Active-Passive Haptics continuum.

5.5.1 Discussion of Experiment 1

The first experiment was designed to demonstrate how the DPHF proxy *Shifty* can be used to enhance the perception of kinesthetic object qualities in VR. At the same time, the experiment evaluated *Shifty*'s performance compared to an equivalent passive haptic proxy that does not change its kinesthetic feedback during runtime. By transferring our previous results on discrete weight distributions and their influence on VR proxy interaction to the continuous level [Zenner, 2016], we assessed how well *Shifty* is suited to provide haptic feedback for virtual objects changing in shape.

The results show that for virtual objects changing in length and thickness, the haptic feedback provided by *Shifty* is significantly more realistic and is enjoyed more by users, compared to an equivalent proxy with fixed weight distribution – verifying **H1**. Participants liked that the kinesthetic feedback adapts to changes in the virtual world. Positive comments by the participants support these results. In a direct comparison of dynamic and fixed kinesthetic feedback, *Shifty*'s dynamic feedback was generally favored by participants. One participant even commented that “*without the motor, it wasn't any fun, especially not if experienced after [the condition with] the motor*”.

Of course, the increased realism of *Shifty* comes at some cost: the interaction with longer, thicker, or heavier objects increases the physical demand and might lead to fatigue. Our results did not show a significant increase in the user's exertion ratings, but that is likely to change when users interact for a longer period of time. Hence, **H2** was only partially supported by our results. Additionally, we would like to note that as a slight amount of audible and haptic noise could not be completely prevented, it cannot be entirely ruled out that these side effects contributed to the participants' experience to a minor degree as well. However, based on our observations, the results of previous investigations [Zenner, 2016], given the experienced shifts, and the participants' comments, the changing kinesthetic feedback was the primary and most significant factor enhancing the experience.

Besides objects changing in shape, we think that *Shift*y's feedback is also suitable to enhance the perception of virtual objects that change their weight or their weight distribution, that are filled up or emptied, or that even change their material. We believe that by leveraging appropriate visualizations and the phenomenon of visual dominance, the perception of many types of object changes can, when visualized in synchronization with *Shift*y's weight shift, be enhanced.

5.5.2 Discussion of Experiment 2

The second experiment shows how *Shift*y can be used to enhance the haptic perception of arbitrarily formed virtual objects with different weights – tackling the drawback of current VR controllers with which all virtual objects feel the same regarding their kinesthetic feedback. The results of the second experiment show that a DPHF proxy designed like *Shift*y increases the perceived realism while interacting with virtual objects of different size and weight. By changing its weight distribution, *Shift*y provides compelling and adapting passive haptic impressions that enhance the perception of virtual weight. In summary, participants clearly favored weight-shifting DPHF over a conventional PHF proxy with fixed weight distribution.

The Effect of DPHF

To evaluate the effect of DPHF, we compared the results of our PHF baseline condition *Base* and our DPHF condition *Hapt*. Differences in the results of *Base* and *Hapt* can be attributed to the differences in the haptic feedback, as this is the only way *Base* and *Hapt* differ. Table 5.2 shows that considering the perceived realism during the interaction after picking up, enjoyment, and the participants' personal preference, the DPHF condition *Hapt* outperformed the PHF condition *Base*, which supports **H3**. This means that users have more fun and prefer interacting with objects that change their kinesthetic feedback compared to purely passive props. Moreover, they perceive the interaction with virtual objects as more realistic when using *Shift*y. In further conformance with **H3**, *Base* was rated as requiring the least exertion.

Moreover, *Base* had less negative influence on the immersion than *Hapt*. This is plausible, as in *Base*, the process of picking up is a very instantaneous action without enduring haptic change. This, in general, is not disturbing to the user, as no noticeable mismatch is involved, despite the general lack of haptic adjustment to the virtual object's weight. In *Hapt*, the user sees the same quick pick-up as in *Base*. Thus the user expects the process to be over as soon as the visual feedback suggests so. But as the weight still moves to its target position for up to 2.8s, users can be irritated by the visual-haptic mismatch during this time. The lack of visual cues that help the user to understand the change in haptic feedback or its progress brings a risk of BIPs. *Hapt* is thus ranked worst for this measure.

The Effect of Audio-Visual Animations

As the previous section summarized the effect of adding the capability of changing its passive haptic feedback to a PHF proxy, we describe in the following how the perception of the user changes when the pick-up process is additionally animated visually and auditorily. For this, we compare the DPHF conditions *Hapt*, *Prog*, *Scal*, *Trans* and *Mask*.

Regarding the perceived realism, *Hapt* is outperformed by all conditions involving auditory and visual animations, which supports **H4**. This shows that the perceived realism can further be increased by animating the object during the pick-up process. This animation should be synchronized with the physical adjustment of the feedback. Furthermore, the results show that animations physically describing the haptic change, like *Scal* and *Trans*, yield the highest perceived realism. We suspect this to be the case as these animations minimize the perceived visual-haptic mismatch during the shift. Scaling an object or making an object become more dense provides a plausible explanation for the changes in *Shifty*'s haptic feedback. Nonetheless, animations that are less related to the haptic change like *Prog* and *Mask* still improve the realism compared to *Hapt*. They still make users aware of an ongoing change in the haptic feedback and allow them to estimate its duration.

Investigating the negative influence of remaining visual-haptic mismatches onto immersion during the shift, we see that all animations score at least as good as *Hapt*. Most score better than *Hapt* and some even better than *Base*, indicating further support for **H4**. *Trans* and *Prog* fall behind *Base* in the ranking, as some participants perceived progress bars in general as disturbing and stressful and some were slightly distracted by the transparent objects. *Mask* and *Scal* scored better than *Base*. In general, *Scal* was noted to be the most natural and suitable animation by some participants as the haptic feedback matched the visual effect of growing objects. The smoke masking effect in *Mask* was perceived differently. While some could explain the effect as the output of the virtual rod in the hand, others could not relate the effect to the object or interaction at all. Thus for some users, immersion was well sustained as the effect did fit into the virtual world. Others, however, were rather distracted by it. In the general case, it certainly depends on the application and scenario. The masking effect should match the context and should be explainable. In games, for example, one can think of effects that match the setting of the game. In summary, DPHF comes with the risk of reducing immersion when the haptic change is not synchronized with the visual or auditory channel. However, the results show that a good and plausible visual and auditory animation matching the change in haptic feedback can lower this risk or even improve immersion.

Adding more realistic inertial feedback also means more physical demand. Here, not simulating an object's weight, as in *Base*, is certainly the least demanding way. The physical demand of DPHF was slightly higher but equivalent for all DPHF conditions. Thus regarding the exertion measure, the mental demand

makes the difference. As *Mask* and especially *Prog* were perceived by some as rather stressful or distracting, they are ranked no better than *Hapt*. *Trans* and *Scal* were considered more suitable and scored better than *Hapt*. This shows that explainable and suitable animations can ease the interaction by decreasing mental demand – a finding that is also in line with **H4**.

Considering enjoyment, the most spectacular animation, the smoke masking *Mask*, clearly leads the score, followed by the most realistic condition *Scal*. Both were generally considered interesting and fun, even by those participants who could not relate the masking effect to the interaction or the virtual object. Conditions *Prog* and *Trans* scored slightly less, but higher than *Hapt* as we expected in the context of **H4**. This ranking emphasizes that more noticeable and less subtle effects in combination with DPHF can increase the entertainment factor.

Finally, considering personal preferences, 2 groups among the DPHF conditions can be identified: the generally preferred conditions *Mask* and *Scal* both scoring highest and the remaining conditions *Prog*, *Trans* and *Hapt*. The condition rated most realistic, *Scal*, thus seems to be as popular as *Mask*, the masking animation rated most entertaining. Less popular is the progress bar in *Prog*, which was described as rather annoying, and the subtle transparency animation *Trans*. When asked for their overall favorite condition, participants clearly preferred realism over the entertaining factor with more than half of the participants choosing *Scal*.

5.5.3 From DPHF to AHF

Our experiments showed that participants favored DPHF over PHF. The distinguishing factor between *Shifty* and a conventional passive prop is *Shifty*'s capability of changing its internal mass distribution via simple actuation. When holding the internal weight at some position $p \in [0, 1]$, *Shifty* is a classical passive haptic proxy object. It then provides the user with passive kinesthetic feedback and the actuator does not exert noticeable forces. However, when using an actuator to change a proxy's physical property, a continuous transition between active and passive haptic feedback becomes apparent.

Imagine the virtual object in the user's hand changes its length over an interval of $t \approx 3s$. The proxy's weight would then be translated to the corresponding target position $p' \in [0, 1], p \neq p'$ in this interval t to change the passive haptic properties of the proxy. The average corresponding shift speed would then be $v = \frac{|p' - p|}{t}$. For slow speeds v , users perceive the desired change in the object's haptic feedback but no noticeable forces are exerted on them actively. However, for more instantaneous or even discrete tasks like picking up a virtual object, the theoretical change interval $t \rightarrow 0s$ and thus $v \rightarrow \infty$. Besides obvious mechanical problems that would arise, high translation speeds would transform the DPHF proxy into an AHF device due to the arising repulsion forces.

To avoid such undesired active forces, changes of the passive properties, even if theoretically instantaneous, have to be realized by DPHF proxies in an appropri-



Figure 5.10: Illustration of the impact of *Shifty*'s transformation speed on its location inside the Active-Passive Haptics continuum. © 2017 IEEE.

ate amount of time $t > 0s$. This prevents noticeable active forces as a side effect of the transformation. The process of picking up an object is a prominent example of such an instantaneous task. Our second experiment helps us to understand how such instantaneous actions can be conveyed when implemented through DPHF transformations that take non-negligible time $t > 0s$. Specifically, we found complementing sensory input, such as visual and auditory animations, to be promising approaches that can reduce the impact on the user's VR experience and compensate for visual-haptic mismatch during the transformation.

As these considerations reveal, the transition between AHF, DPHF, and PHF is flowing with *Shifty* following the pattern illustrated in Figure 5.10. We conclude that for DPHF proxy objects, an increasing transformation speed can shift their feedback closer towards the AHF end of the continuum.

5.5.4 Limitations

While our experiments show that *Shifty* can enhance the perception of mass and inertia well, our prototype still has some limitations. Currently, *Shifty* produces slight vibration and noise as a by-product of the weight shift and when heavily shaken. In future iterations, this could be further reduced by damping the motor and the internal weight. However, most of the participants perceived this effect as complementary feedback similar to the vibration feedback employed by commercial VR controllers and only some found it slightly disturbing. Despite that, *Shifty* does not change its actual size and the grip grasped by the user always feels the same in terms of texture and shape. Moreover, *Shifty* cannot simulate arbitrary forces or weights. Our experiments showed, though, that the range of inertia that *Shifty* is capable to produce suffices to enhance the interaction with typical everyday objects inside the IVE. Finally, when considering the physics of *Shifty*, a special case exists: When holding the proxy in an upright position, the effective lever arm vanishes and the user only perceives inertia when trying to move the proxy. However, as this special case could be prevented using redirection techniques, it does not restrict *Shifty*'s areas of application.

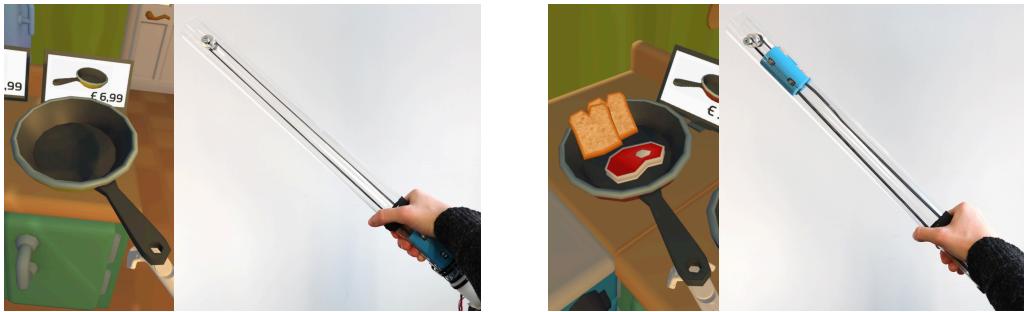


Figure 5.11: *Shift* conveying the feeling of different pans in the virtual shopping scenario. Left: A lightweight pan. Right: A heavy pan.

5.6 Applying *Shift* in Different Application Domains

To demonstrate how a VR controller like *Shift* can enhance VR applications in different domains, we developed a demonstration application. This application integrates the *Shift* prototype in three different VR experiences. Each experience is a reference to a different application domain and showcases how the domain could profit from a haptic VR controller that can render different kinesthetic properties. The demonstration is regularly shown to visitors at the German Research Center for Artificial Intelligence (DFKI)⁶¹ and the Innovative Retail Laboratory (IRL)⁶². It integrates *Shift* in the following scenarios:

Virtual Shopping Simulation The virtual shopping scenario demonstrates the virtual pan store outlined in the introduction of this thesis in chapter 1. Users experiencing this demo find themselves in a virtual kitchen with the three different pans shown in Figure 1.2 on display. Inspired by the vision of VR online shops, each of the pans represents a different product that features a different price and different physical characteristics. *Shift* is used in this scenario as a VR controller that enables users to pick up and wield the pans, conveying the differences in weight and balance of the pans through its weight shifting feedback as illustrated in Figure 5.11. Moreover, the demo application encompasses a mini game, which allows users to use the virtual pans for cooking. The virtual kitchen holds two machines that dispense virtual ingredients. By holding the pan beneath the dispensers, users can feel the differences in weight as the pans fill with virtual ingredients, allowing them to playfully assess the pan's handling. A video⁶³ about this demonstration can be found online.



Video Link⁶³

Virtual Maintenance Simulation In a second demonstration scenario, *Shift* is used to enhance a virtual maintenance simulation, inspired by VR applications in

⁶¹DFKI Homepage. <https://bit.ly/3FBV8t0>

⁶²*Shift* Demo on the IRL Homepage. <https://bit.ly/3oPtdc>

⁶³*Shift* Demo Video. <https://bit.ly/32bGTNa>



Figure 5.12: Two more applications that take advantage of *Shifty* to convey the kinesthetic qualities of virtual objects. Left: The maintenance simulation featuring various tools. Right: The shelf restocking simulation featuring retail products of different heaviness.

the industrial sector. Here, users are immersed in a virtual factory environment and are to repair a virtual machine. To do so, different tools are available, such as a virtual hammer, a drill, and electrical components. The DPHF of *Shifty* is used here to simulate differences in heaviness, balance, and handling of the tools similar to how the feeling of different pans is conveyed in the virtual shopping scenario. The left of Figure 5.12 shows a screenshot.

Virtual Shelf Stocking Simulation The third example application in our demonstration simulates a retail store in VR and lets the user experience products that differ in size and shape. In reference to our experimental evaluation of *Shifty* outlined above, users here can interact with objects that can change in length and thickness, such as a virtual toy sword or a virtual piece of cheese, and feel their differences. Moreover, this scenario implements a shelf stocking simulation. Users find themselves in front of a supermarket shelf that has been messed up, with items lying on the floor and misplaced items inside the display. The application can be used to train the correct restocking of the shelf while providing kinesthetic cues. By this, users can, for example, feel that heavy items are to be placed in lower shelf compartments, while lightweight objects can easily be placed into compartments at or above head height. The scenario is depicted on the right in Figure 5.12.

5.7 Conclusion & Contribution to the Research Questions

In this chapter, we lift the concept of DPHF from the theoretical stage to practice by presenting our first DPHF proxy object: *Shifty*. Furthermore, we present our first solution to the second part of **RQ 2** by introducing a novel concept for proxy-based kinesthetic feedback that only requires minimal actuation. This concept is based on the idea of dynamically adapting a proxy's *weight distribution* to take advantage of dynamic touch [Turvey, 1996], i.e., the human perception of inertia. Specifically, with *Shifty* we propose a VR proxy (or controller) that

has a basic build-in actuator to modify the location of an internal mass, allowing the proxy to change its passive haptic feedback (here: the kinesthetic qualities perceived by the user when wielding the object) in order to establish perceptual *Similarity* with a variety of virtual objects. *Shifty* combines the advantages of PHF and AHF as it is ungrounded and made out of cheap and widely available materials like a passive prop, but offers enhanced flexibility and generality due to its actuation mechanism. As such, we contributed to **RQ 2** by demonstrating that DPHF allows for the realization of kinesthetic haptic feedback while only leveraging low-cost and low-power actuation – in contrast to AHF approaches usually used to simulate properties like virtual weight and shape.

In two experiments, we further added to **RQ 2** when we showed that *Shifty* enhances the flexibility of proxy-based haptics as it can haptically represent a large set of virtual objects more realistically than an equivalent PHF prop. We confirmed that *Shifty* enhances the perception of virtual objects that change in length or thickness. Moreover, we found that the perceived realism when interacting with *Shifty* is significantly higher, that users have significantly more fun, and that users generally prefer interacting with a DPHF proxy compared to an equivalent PHF prop. These results, and the fact that we found support for all four of our hypotheses **H1** to **H4** highlights that weight-shifting DPHF is a particularly versatile category of haptic feedback for VR.

Leveraging the effect of visual dominance, we could produce a variety of different visual-kinesthetic illusions based on dynamic touch (introduced in section 2.2), such as those of interacting with objects of different length, thickness, or heaviness. We achieved these perceptual illusions by combining the ever-same haptic stimuli (i.e., varying states of inertia) with different visualizations (i.e., virtual animations that show objects change in length, or thickness, or objects being picked up). Depending on the visualization perceived by the user, we could control the kind of illusion conveyed, such as a length-, thickness-, or weight-illusion. We see this approach, illustrated in Figure 5.13, as an important tool on the path towards the ultimate display, making proxy-based haptics more flexible and believable. As such, we see our verification of this approach in conjunction with a weight-shifting proxy presented in this chapter as a central contribution to **RQ 2**. Based on our results, we are confident that with weight-shifting proxies like *Shifty*, we could also enhance the perception of object properties beyond shape and weight. It is easy to imagine visual-haptic illusions that would target, for example, the perception of virtual materials, the content of virtual objects, or virtual weight distributions and states of balance. In line with this, we think that interactions like filling or emptying a virtual object could likewise be enhanced.

While the first experiment focused on rendering objects that gradually change their haptic properties, the second experiment considered events that imply instantaneous kinesthetic changes. For this, we used *Shifty* as a VR controller and investigated how virtual objects of different weight could be picked up with it. We found that theoretically instantaneous haptic changes cannot unrestrictedly be simulated by fast transformations of the proxy's properties due to arising

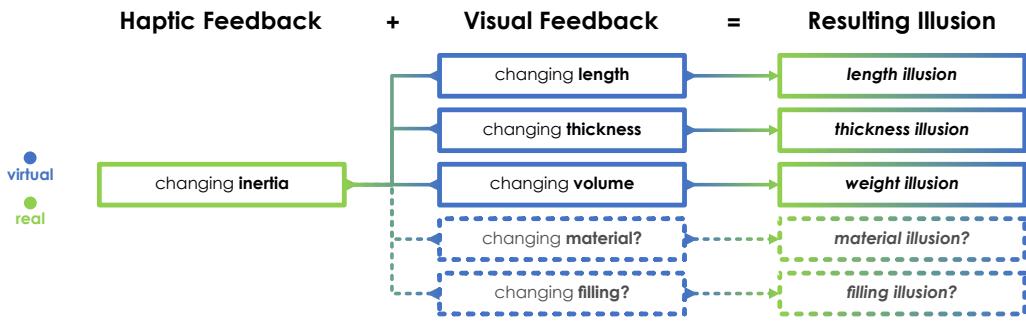


Figure 5.13: Combining different visualizations with weight-shifting DPHF allows for a variety of visual-haptic illusions.

repulsion forces. This led to our finding that the speed of the transformation (i.e., weight shift) determines the location of the DPHF prop inside the Active-Passive Haptics continuum, as illustrated in Figure 5.10. As a consequence, we implemented *Shifty* so that its mass moves with a speed that does not exert noticeable active forces on the user. To compensate for the arising visual-haptic mismatch during the shift, we further investigated visual and auditory animations. We found that appropriate visualizations matching the perceived haptic change, even if abstract, can decrease the negative impact on immersion and increase the perceived realism and enjoyment. In our experiment, scaling a picked-up virtual object while *Shifty*'s weight moves was perceived as most realistic and was generally favored by participants.

As outlined by Schneider et al. [2005] and verified by our experiments, weight-shifting DPHF is compatible with the form factor of VR controllers and has the potential to tackle the *Similarity* challenge for many types of virtual objects. The power of leveraging weight shift to simulate the most diverse virtual objects has also been realized by other researchers in this field and spawned a range of follow-up works. Our own follow-up work is presented in the next chapter while that of other researchers is summarized in detail in subsection 2.5.2. Based on our work on *Shifty*, for example, Shigeyama et al. [2019] recently transferred the concept of weight-shifting from 1D to 2D with *Transcalibur* in order to simulate virtual objects of different shapes. Furthermore, with their device *SWISH*, Sagheb et al. [2019] applied weight-shifting feedback in 3D to simulate the feel of virtual fluids moving inside a virtual container.

Chapter 6

Drag:on – Dynamic Passive Haptics Based on Drag and Weight Shift

In the preceding chapter, we introduced our first DPHF-based proxy and investigated the concept of leveraging weight shifts to convey different kinesthetic impressions in VR. In this chapter, we will continue our **Improvement** of proxy-based haptics by exploring the concept of DPHF further and evolving the idea presented in the previous chapter. Continuing our investigation of the second part of **RQ 2**, we will extend the idea of weight-shifting feedback and pick up the concept of visual-kinesthetic illusions again with a second, novel feedback mechanism that is also capable of conveying different kinesthetic cues through a single dynamic proxy. This chapter will present the concept and implementation of a corresponding prototype, which, again, was realized in the form factor of a handheld VR controller. Furthermore, we present the results of a user evaluation that compares the presented DPHF-based feedback to conventional PHF and to the vibrotactile feedback supported by today's commercial VR controllers.



Video Link⁶⁴

A video⁶⁴ about the work presented in this chapter is available online. This chapter is based on the following two publications. Images and parts of the text in this chapter, as well as the presented concepts, implementations, and results have been published previously therein:

Zenner, A. and Krüger, A. (2019a). Drag:on – A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'19, pages 1–12. ACM. © 2019 André Zenner and Antonio Krüger. Final published version available in the ACM Digital Library.
DOI: [10.1145/3290605.3300441](https://doi.org/10.1145/3290605.3300441)

⁶⁴Drag:on Video. <https://bit.ly/3DX8OP6>

Zenner, A., Degraen, D., Daiber, F., and Krüger, A. (2020a). Demonstration of Drag:on – A VR Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'20, pages 1–4. ACM. © 2020 André Zenner and co-authors. Final published version available in the ACM Digital Library. DOI: [10.1145/3334480.3383145](https://doi.org/10.1145/3334480.3383145)

Certain aspects of the feedback concept presented in this chapter are protected by the following German patent:

Zenner, A. and Krüger, A. (2021). Hand-Virtual Reality/Augmented Reality-Steuengerät, Virtual Reality/Augmented Reality-System mit demselben sowie Verfahren zur Simulation der Haptik. *German Patent 10 2019 105 854*. June 05, 2021.

6.1 Introduction

The work in this chapter is motivated by the same gap in the landscape of VR haptics that already drove our research on *Shifty* in the previous chapter. Following **RQ 2**, we continue our search for kinesthetic feedback that does not require significant actuation and at the same time is compatible with the form factor of ungrounded, handheld VR controllers. Similar to our work on *Shifty*, we look for a concept that has the potential to complement the vibrotactile feedback commonly built into consumer VR controllers, which fails to provide different kinesthetic impressions such as the feeling of weight, resistance, or inertia – haptic impressions that we expect and rely on when interacting with our environment.

In this chapter, we propose a solution that extends the concept of weight-shifting feedback and visual-kinesthetic illusions by combining it with a second novel feedback approach. In particular, we present a novel concept for providing haptics, introducing a combination of air resistance and weight shift as a means of generating haptic feedback in VR. We contribute a shape-changing VR controller called *Drag:on*, which leverages the airflow that occurs at the controller during VR interaction to provide a range of different haptic sensations. For this, *Drag:on* can self-transform while the user interacts with it in VR, increasing or decreasing its surface area to adapt its air resistance profile and its mass distribution. Similar to *Shifty*, the *Drag:on* controller is a general-purpose DPHF proxy that uses simple built-in actuators only to change its physical configuration. By this, the device adapts to the virtual interaction yielding different passive haptic impressions when moved through the air.

In the following, we will introduce the underlying haptic feedback concept and our low-cost and mechanically simple prototype implementation. After that, we present a user evaluation in which we studied how *Drag:on*'s haptic feedback can enhance the perception of virtual interactions in five VR scenarios.



Figure 6.1: *Drag:on* leverages the airflow occurring during controller movements.

6.2 Concept of *Drag:on*

During many VR interactions, users swing, drag, throw, or rotate virtual objects. While doing so, they perform both rotational and translational motions with the corresponding haptic proxy (e.g., the VR controller) as illustrated in Figures 6.1 and 6.3. As a consequence of such movements, an airflow starts forming around the device as the user pushes its resisting surface through the air. The feedback concept of *Drag:on* proposed in this chapter takes advantage of exactly this phenomenon.

The central idea underlying *Drag:on*'s DPHF is to adjust the proxy's surface area to produce different kinesthetic sensations of resistance when it is physically moved. By this, instead of simulating constant forces in a 1-to-1 manner, *Drag:on* leverages the motions of the user to provide resistance impressions that vary with velocity. In line with the concept of *Shifty*, our proposed concept of *Drag:on* utilizes immersive visualizations and the dominant impact of vision in order to bridge visual-haptic mismatches and to induce a variety of visual-kinesthetic illusions based on air resistance, following the paradigm outlined in Figure 5.13.

Inspired by concepts of shape-changing interfaces [Alexander et al., 2018], we propose to implement our concept with proxies that transform their shape in order to adjust their surface area. Depending on the implementation of the shape change, this form of proxy adjustment lends itself to additionally include secondary physical effects in the illusion. To make use of such in this chapter, we opted for a proxy design with foldable surfaces, i.e., fans – a form factor that has already been focus of research on foldable displays [Lee et al., 2008]. We base our decision to use foldable structures on two main considerations: (1) fan-based designs are mechanically simple, low-cost and easy to replicate, and (2) in addition to drag, they allow to leverage inertial changes, which we showed in the previous chapter to be suitable for providing kinesthetic cues.



Figure 6.2: The five states of our shape-changing haptic VR controller *Drag:on* investigated in this chapter. (a) The device state S_{closed} with minimal surface area. When increasing its surface area symmetrically as shown in (b) (S_{half}) and (c) (S_{full}), the controller adapts its drag and mass distribution to provide different haptic sensations during VR interaction. If opened asymmetrically on the (d) left (S_{left}) or (e) right (S_{right}) side, torque is induced when moving the controller.

Specifically, in our proposed design, the surface area of a VR controller is increased or decreased by opening or closing two fans symmetrically on the left and right side of the device, as depicted in Figure 6.2 (a) to (c). Changing the shape of the controller in this way affects its drag coefficient and allows the VR system to configure different device states that will result in different kinesthetic perceptions during interaction. Moreover, we propose to control the surface areas on both sides of the controller individually. By this, it is possible to increase the area only on one side of the device, as shown in Figure 6.2 (d) and (e), in order to induce torque. This torque rotates the controller inside the user’s wrist during motion and allows for rendering asymmetric forces and differences in resistance, as we will illustrate in our experiments at the end of this chapter. Assuming a device with two adjusting fans, such as the prototype we will introduce in the following section, the physical state of the proxy can be described by a tuple $S = (open_{left}, open_{right})$, given by the percentage of opening of the left and right fan. While such a proxy can take any state $S \in [0, 100] \times [0, 100]$, we focus our investigation on the five states defined in Table 6.2 and shown in Figure 6.2.

Apart from drag, our proposed design also takes advantage of changes in the mass distribution of the proxy when opening or closing its fans, affecting its inertial response when *rolling* or *swinging* it. By opening the fans as shown in Figure 6.3 (a), for example, the moment of inertia I_{roll} , i.e., the rotational resistance when *rolling* the device about the longitudinal axis (indicated in red), increases as mass is moved away from the axis. This supports and amplifies the resistance feedback felt when rolling the controller with the wrist, in addition to the increased drag.

The design proposed here, however, also comes with two minor drawbacks. Firstly, when considering *swinging* the controller as shown in Figure 6.3 (b), opened fans lead to mass being moved towards the swing-axis (indicated in red) passing through the user’s shoulder. By this, the corresponding rotational inertia I_{swing} is slightly reduced, acting against the effect of increased resistance. This characteristic is not a practical limitation, though, as the relative change of I_{swing} when opening or closing the fans is much lower than the relative change

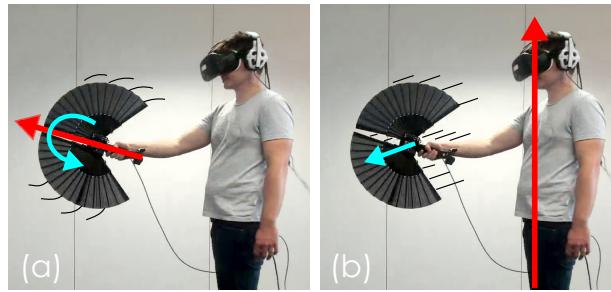


Figure 6.3: The two types of interaction investigated in this chapter. Movement direction (blue), rotational axis (red), and motion (black lines) are highlighted. (a) Illustration of rotational movements (rolling the controller). (b) Illustration of translational movements (swinging the controller).

of I_{roll} . As a simplifying assumption, we thus regard swinging the controller as translation in the following. Secondly, in our presented implementation of *Drag:on*, opening the fans leads to a reduction of the rigidity of the controller's overall structure, which in turn leads to the controller slightly bending in the airflow at higher motion speeds. While such bent shapes make the device more aerodynamic, the increase in surface area which is caused by the shape change remains the factor dominating the felt drag force in our design. Our evaluation will show that in practice, the drag feedback of our controller can produce the desired impressions despite these two counteracting effects.

It seems noteworthy to highlight that existing air-based haptic feedback (as reviewed in section 2.3) usually relies on powered propellers with high energy requirements or air jet actuation, or requires compressed air or air compressors to render forces. In contrast, our entirely novel approach presented in this chapter works without any of these, leveraging solely the ubiquitous airflow that occurs at the controller during VR interactions.

6.3 Implementation of *Drag:on*

In the following we present *Drag:on*: the simple, low-cost and easily reproducible implementation we used for studying the presented concept. Our device is not a definitive implementation, but rather one of many imaginable designs.

All relevant resources, such as a detailed list of parts, relevant 3D model data for 3D printing, source code, as well as a construction manual can be found online in an open-source repository⁶⁵ allowing interested readers to recreate *Drag:on*.



GitHub Link⁶⁵

⁶⁵*Drag:on* on GitHub. <https://github.com/AndreZenner/dragon>

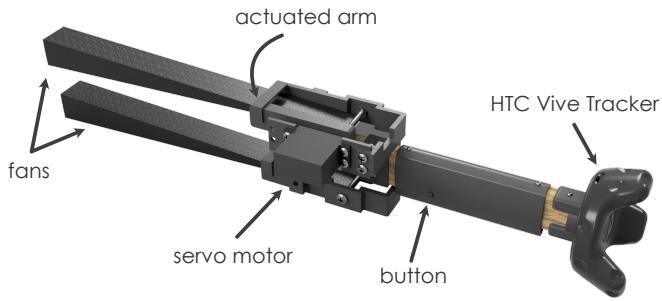


Figure 6.4: 3D rendering of the *Drag:on* prototype.

6.3.1 Hardware

The 3D rendering in Figure 6.4 shows the main components of the *Drag:on* device. The controller consists of a wooden base with a screwed-on custom-designed 3D-printed mount to attach the HTC Vive Tracker⁵⁴. In addition, a 3D-printed grip holds a small pushbutton attached with a rubber band. The location of this button can be adjusted to account for the handedness of the user. The actuation mechanism depicted in Figure 6.5 (b) is located at the top end of the controller. On both sides of the controller, we fixed an MG996R servo motor using custom 3D-printed parts. Each servo actuates a 3D-printed arm attached to the topmost layer of a commercially available flamenco hand fan. The fans are 31cm long and made out of wood and fabric, as can be seen in Figure 6.2. The bottommost layer of the fan is rigidly attached to a 3D-printed support structure, pointing away from the user. By actuating the servo, the arm opens or closes the fan. Figure 6.5 (b) shows a servo and an actuated arm opening a fan. To allow for unconstrained movements, the maximum opening angle of the right fan is slightly limited to leave enough space for the user's arm (see Table 6.1). Moreover, to ensure comfortable interaction, we designed the prototype to concentrate its mass close to the user's hand, minimizing its overall moment of inertia.

Figure 6.5 (a) shows the final prototype with its main system components. The device is connected to a controller box containing an Arduino Nano microcontroller and the necessary circuits. An external power adapter connects to this box to provide 7.6V to the motors. The Arduino interfaces with the VR system via USB serial communication (115200 baud). Table 6.1 summarizes the technical

Fan Angle	min.	5°
	max. left	152.5°
	max. right	132.5°
Area	min. (S_{closed})	320cm ²
	max. (S_{full})	2400cm ²
Time ($S_{closed} \rightarrow S_{full}$)	total	570ms
Power Consumption	idle	0.23W
	peak	6.84W
Length	fan	31cm
	total	54cm
Weight	fans	2 × 75g
	total	598g

Table 6.1: *Drag:on*: Technical Data



Figure 6.5: Left: *Drag:on* is connected to a box holding the microcontroller and circuits. The box connects to the computer via USB. Motor power is provided by an external power adapter. Right: Servo motor and connected arm opening a fan.

data of the *Drag:on* prototype with the HTC Vive Tracker attached. The surface area referred to in Table 6.1 and Table 6.2 is the area of its orthographic projection on a plane parallel to the fans, i.e., the area visible in Figure 6.2. *Drag:on* can increase its surface area by up to $\frac{A(S_{full}) - A(S_{closed})}{A(S_{closed})} = \frac{2400\text{cm}^2 - 320\text{cm}^2}{320\text{cm}^2} = 650\%$ in 570ms.

6.3.2 Software

The software stack of *Drag:on* involves two central components, depicted in green in the architecture overview in Figure 6.6. The C++ software controlling the device runs on the Arduino Nano. It forwards button state changes to the VR system and controls the servo motors upon reception of transformation commands. The Arduino uses a simple custom protocol to communicate with the VR system on the computer via USB serial connection. The second main component is the C# interface script for the Unity 3D engine. This script handles serial communication with the controller and implements convenient functions to control the state of the device. With these functions, the IVE logic can send transformation commands to the *Drag:on* and receive button state changes.

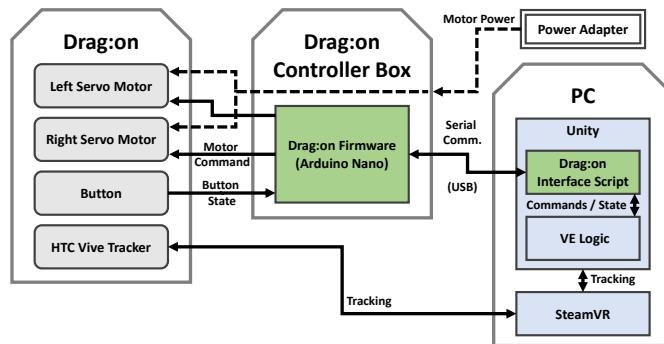


Figure 6.6: Overview of the software architecture of the *Drag:on* system. Sensors and actuators are colored in gray, software in blue, and the two main software components of *Drag:on* are highlighted in green.

State	open _{left} (%)	open _{right} (%)	Area	Figure 6.2
S_{closed}	0	0	320cm^2	(a)
S_{half}	50	50	1320cm^2	(b)
S_{full}	100	100	2400cm^2	(c)
S_{left}	100	0	1410cm^2	(d)
S_{right}	0	100	1250cm^2	(e)

Table 6.2: *Drag:on*: Investigated States

6.4 Evaluation of *Drag:on*

To evaluate our second DPHF concept, we conducted a user experiment with the *Drag:on* prototype. Our study is organized in two parts and studies how a haptic VR controller providing combined drag and inertial DPHF can enhance the user’s perception of various objects, interactions, and IVEs.

The first part of our study investigated how the DPHF of our controller is perceived in three different interactive VR scenarios. To this end, we compared different states of *Drag:on* to test if they can provide distinguishable levels of haptic feedback. The second part studied the two types of interaction introduced in Figure 6.3, i.e., rotating and translating the controller, in two additional VR scenarios individually. Here, in addition to comparing different *Drag:on* states, we compared the DPHF of *Drag:on* to a PHF baseline and the vibrotactile feedback of standard HTC VIVE Controllers. The experiment was approved by the ethical review board of our faculty.

6.4.1 Participants

Our study was conducted with $N = 18$ (4 female, 14 male) volunteer participants aged between 21 and 33 years ($Mdn = 27$, $SD = 3$). All participants had normal or corrected-to-normal vision and 15 were right-handed.

6.4.2 Apparatus

The experiment took place in a quiet lab environment and was carried out with our *Drag:on* prototype and an HTC Vive¹⁴ HMD, Trackers⁵⁴, and Controllers⁵³. Software-wise, the study was implemented with the Unity 3D engine¹⁷. Participants stood in the center of the tracking area and had enough space to freely swing the controller. To dampen the sound of the servos, participants wore over-ear headphones with which they could hear the interactions and background sounds of the IVE. To further exclude effects due to users perceiving servo noise or vibrations, we implemented an obfuscation mechanism to create random transformation noise each time *Drag:on* was supposed to change state. Whenever instructed to transform to a target state S , *Drag:on* first transformed to a random state $S' \in [0, 100] \times [0, 100]$ before transforming to S . This effectively doubled the

transformation time to up to 1140ms. In the study, however, this did not introduce significant delays as *Drag:on* only transformed in between interactions.

6.4.3 Procedure

Before starting the experiment, participants were briefed by reading through a prepared document explaining the five VR scenarios encountered during the study. They were intentionally not informed about the controllers interacted with and the *Drag:on* prototype was hidden from them until the end of the study. At the beginning of each scenario, participants could become familiar with the interaction and their task by performing a short training trial. We recorded their responses only after completion of the training trial. Upon completion of the last scenario, participants filled in the SUS presence questionnaire [Slater et al., 1994], a demographic questionnaire, and additional post-study questions. The experiment ended with a verbal debriefing and took ca. 95min per participant.

Part 1: Comparing *Drag:on* States

The first part of the study compared the *Drag:on* states introduced in Table 6.2 in the three VR scenarios *Scale*, *Material* and *Flow*. In all these scenarios, participants were immersed in a virtual factory environment, holding the *Drag:on* prototype in their dominant hand and a secondary HTC Vive Controller in their other hand. When a trial started, *Drag:on* transformed to the state associated with that trial (independent variable), with active obfuscation. *Drag:on* then remained in this state until the beginning of the next trial. In each trial, participants interacted with a virtual object (*Scale*, *Material*) or environment (*Flow*). The task of the participants in each scenario was to freely explore the haptic response of the object in their hand (*Scale*, *Material*) or the environment (*Flow*) to get a feel for it. They were free to do so by swinging the controller (in all three scenarios) or rotating it (only in *Scale* and *Material*). Their task was then to adjust the VR visualization of the virtual object or environment until it matched their haptic impression best. For this, in all three scenarios, a simple UI was displayed on the secondary controller that allowed participants to adapt the visualization of the objects or environment

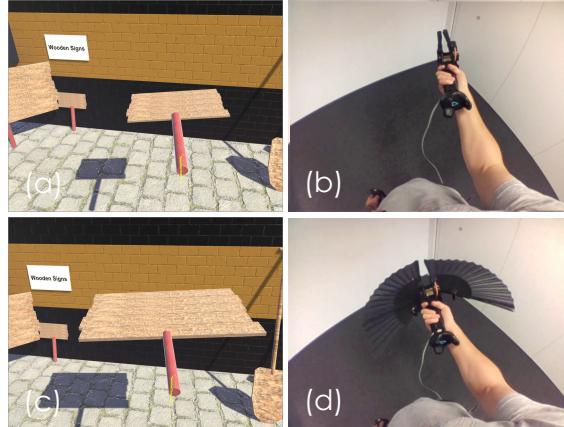


Figure 6.7: *Scale* scenario: (a) The avg. scale (1.33) associated with (b) the state S_{closed} is significantly smaller than (c) the avg. scale (2.38) associated with (d) S_{full} .

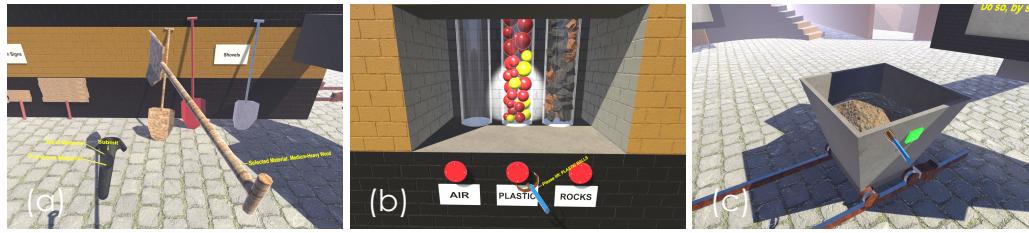


Figure 6.8: Material, Ratchet and Wagon scenarios: The participant (a) selected a wooden *Material* for the shovel, (b) is supposed to lift the plastic balls with the corresponding *Ratchet*, and (c) is about to move a half-filled *Wagon*.

interacted with. By pressing a button on the controller, participants could record their best-matching configurations (dependent variables).

Scenario 1: Scale The *Scale* scenario (**S**) compared the states S_{closed} , S_{half} , and S_{full} , and investigated our hypothesis:

H-S

Different *Drag:on* states are associated with different object sizes.

For this, participants interacted with a virtual wooden sign as shown in Figure 6.7. With the UI on the secondary controller, they could scale the sign up or down in the scaling range [1, 3], and record their selected best-matching scale.

Scenario 2: Material The *Material* scenario (**M**) compared S_{closed} , S_{half} and S_{full} to test our hypothesis:

H-M

Different *Drag:on* states are associated with different object materials.

In this scenario, participants interacted with a virtual shovel as shown in Figure 6.8 (a) and could change its material. Using the UI on the secondary controller, they could select and record their best-matching material from a set of three materials that visually implied different weights: lightweight plastic, medium-heavy wood, and heavy metal.

Scenario 3: Flow The *Flow* scenario (**F**) explored how asymmetric drag-based haptics, especially in comparison to symmetric feedback, can enhance the perception of environmental elements like virtual gas flows. For this, we explored S_{left} , S_{full} and S_{right} to test hypothesis:

H-F

Different *Drag:on* states are associated with different gas flow distributions.

To investigate this, participants interacted with a virtual paddle and faced an upper and lower gas stream, released through two pipes in front of them inside the IVE, as shown in Figure 6.9 (a) and (c). Their task was to swing the paddle horizontally through both gas streams towards the pipes (as illustrated by the arrow) to feel which of the streams is stronger, or if both are equally strong. Using the UI on the secondary controller, they could adjust the visualization of the streams and configure the relative stream strengths that they perceived as best-matching. For this, they could distribute a total power of 100% between the two streams. The relative upper gas flow strength of their selected best-matching configuration was recorded. We accounted for the handedness of participants by adapting the IVE accordingly.

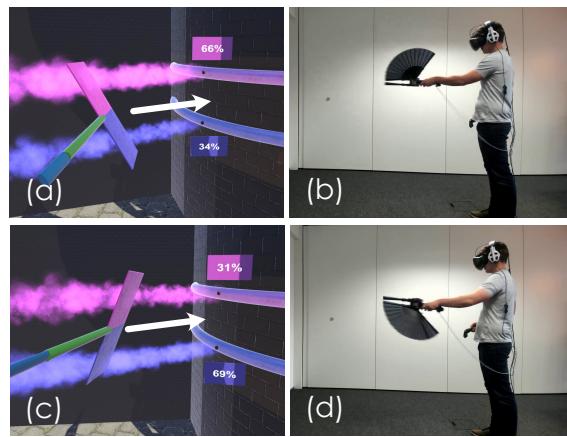


Figure 6.9: Flow scenario: (a) The avg. relative upper gas flow (66%) associated with (b) S_{right} is significantly stronger than (c) the avg. relative upper flow (31%) associated with (d) S_{left} .

Part 2: Comparing Haptic Feedback Techniques

The second part of our study investigated rotational (*Ratchet* scenario) and translational motions (*Wagon* scenario).

Scenario 4: Ratchet The *Ratchet* scenario studied how *Drag:on* can render mechanical resistance felt when turning ratchets (i.e., virtual dials). Participants stood in front of the three ratchets in Figure 6.8 (b) and three glass containers filled with air, plastic balls, and rocks, respectively. By turning the ratchet beneath a container with a rotational movement as shown in Figure 6.3 (a), the corresponding content could be lifted up. The task of the participants was to lift each material three times and the material to lift next was indicated by a spotlight.

Scenario 5: Wagon The *Wagon* scenario investigated how *Drag:on* can render the weight felt when moving virtual objects inside the IVE. Here, participants had to move a virtual wagon as shown in Figure 6.8 (c) along rails from right to left and back again by grasping, swinging, and releasing it with the controller as shown in Figure 6.3 (b). The wagon was visually either empty, half-filled, or completely filled with sand, and each fill state was experienced three times.

Scenario 4 & Scenario 5 The last two scenarios compared the haptic perception of S_{closed} , S_{half} and S_{full} as in Part 1, and additionally compared *Drag:on*'s DPHF to a PHF baseline and the vibrotactile feedback of HTC Vive Controllers. Each of the scenarios was experienced once with each haptic technique (DPHF, PHF, VIVE). In DPHF conditions, the different ratchets and fill states of the wagons were mapped to the tested states (air/empty $\rightarrow S_{closed}$, plastic/half-filled $\rightarrow S_{half}$, rocks/full $\rightarrow S_{full}$) and *Drag:on* transformed to them when grasping the ratchet or wagon. In PHF conditions, participants also interacted with *Drag:on*, which here only transformed for obfuscation and always returned to S_{closed} for each ratchet and wagon — providing the feedback of an equivalent *passive* prop during interaction. In VIVE conditions, users interacted with an HTC Vive Controller⁵³ instead of *Drag:on*, providing different vibration patterns, implemented with the SteamVR Interaction System for Unity (SteamVR haptic racks: air/empty \rightarrow [64 pulses, each 1ms], plastic/half-filled \rightarrow [128 pulses, each 2.5ms], rocks/full \rightarrow [256 pulses, each 4ms]). The visual-haptic feedback combination represents the independent variable.

After each interaction, participants were asked about the resistance (*Ratchet*) or weight (*Wagon*) experienced during interaction on a 1-to-7 Likert scale (1 = very low resistance/very lightweight; 7 = very high resistance/very heavy). When completing a scenario with a haptic feedback technique, participants also rated the haptic realism (1 = not at all realistic; 7 = highly realistic). Perceived resistance, weight, and haptic realism represent the dependent variables.

For the *Ratchet* (**R**) and *Wagon* (**W**) scenarios, we tested for each haptic technique (DPHF, PHF, VIVE) the hypothesis:

H-<Scenario>-<Haptic Technique>

Users perceive different resistances/weights of the ratchets/wagons.

We further hypothesized for both scenarios:

H-<Scenario>-Range

The range of resistances/weights conveyed with DPHF is greater than with PHF and VIVE.

H-<Scenario>-Realism

Users perceive the DPHF rendering of resistances/weights as more realistic than PHF and VIVE.

6.4.4 Design

The study was a within-subjects experiment. The order of scenarios in Part 1 (*Scale*, *Material*, *Flow*) was counterbalanced by a 6×3 Latin square [Williams, 1949], with the nine trials in each scenario counterbalanced by a 18×9 Latin square. In Part 2, 9 participants experienced the *Wagon* scenario after *Ratchet*, while for all others, this was reversed. The order of the three haptic conditions

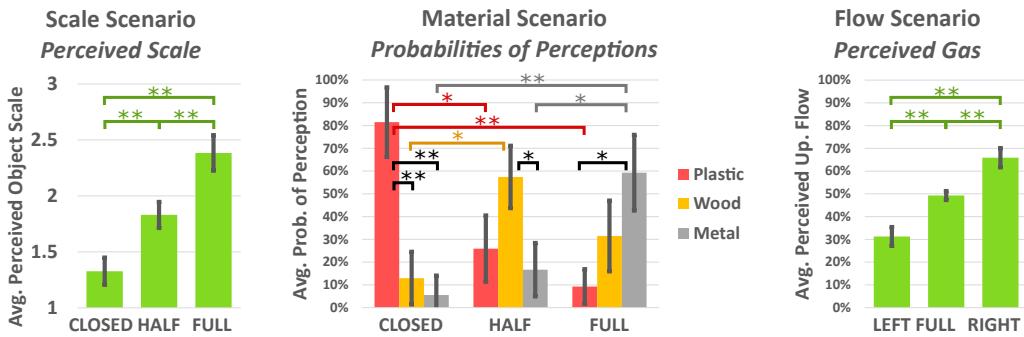


Figure 6.10: From left to right: *Scale* chart plotting perceived virtual object scale. *Material* chart showing perception probabilities for the different materials. *Flow* chart illustrating perceived relative strength of the upper gas stream. Tested *Drag:on* states on the x-axis. Brackets indicate statistically significant differences ($p < .05$ (*); $p < .01$ (**)). Error bars show 95% confidence intervals.

(DPHF, PHF, VIVE) tested in *Wagon* and *Ratchet* was counterbalanced by a 6×3 Latin square. Within each haptic condition, participants performed nine trials (three levels of feedback, each 3x), counterbalanced by a 18×9 Latin square.

6.4.5 Results

We first summarize the results of Part 1 of our experiment, and then outline the results of Part 2.

Results of Part 1

We investigate the effect of the tested *Drag:on* states on the dependent variables of each scenario. For multiple comparisons, we performed non-parametric Friedman tests with pairwise post-hoc Wilcoxon signed-rank tests and applied a Bonferroni-Holm correction. Significant results of pairwise tests are indicated in the referenced charts ($\alpha = .05$).

Figure 6.10 shows the main results of the *Scale*, *Material* and *Flow* scenarios. Friedman tests found significant effects of *Drag:on* state on perceived object scale ($\chi^2(2) = 32.11, p < .001$), and on perceived relative upper gas flow strength ($\chi^2(2) = 36, p < .001$). Post-hoc analysis revealed a significant difference in mean perceived object scale and upper gas flow strength for all pairwise comparisons of S_{closed} , S_{half} and S_{full} in *Scale*, and of S_{left} , S_{full} and S_{right} in *Flow* (all $p < .001$). To evaluate the *Material* scenario, we computed the average probability of selecting a material as “best-matching” for each state and tested within each state for differences, as well as across states individually for each material. Friedman tests confirmed material probabilities to differ significantly within each state, as well as across states (all $p \leq .009$). The probabilities and the results of the pairwise comparisons can be seen in the second chart in Figure 6.10.

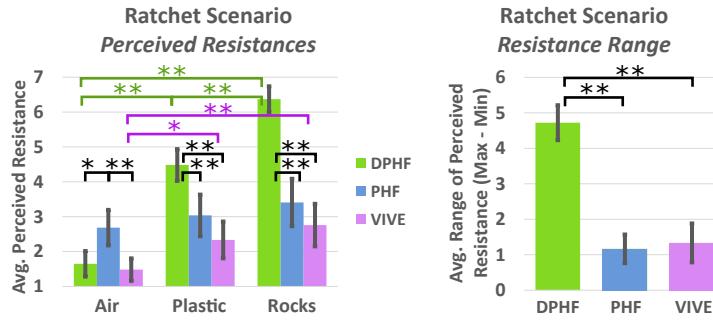


Figure 6.11: Perceived *Ratchet* resistances and ranges in Part 2 of our study. Brackets indicate pairwise significant differences ($p < .05$ (*); $p < .01$ (**)). Error bars show 95% confidence intervals.

Results of Part 2

To evaluate Part 2, we applied the same test procedures as in Part 1. Friedman tests showed perceived resistance and weight to vary significantly (1) with the visual-haptic impression of the ratchets and wagons for *DPHF* (both $p < .001$), *PHF* (both $p \leq .036$) and *VIVE* (both $p \leq .007$); and (2) with the haptic technique for air/empty, plastic/half-filled and rocks/full (all $p < .001$). Significant results of the pairwise comparisons are indicated in Figures 6.11 and 6.13 presenting the results of both scenarios. The chart in Figure 6.12 depicts a comparison of the haptic realism achieved with the different feedback techniques. Friedman tests also found a significant effect of feedback technique on haptic realism and range of feedback provided (shown in Figures 6.11 and 6.13) for *Ratchet* and *Wagon* (all $p < .001$).

Results of Post-Study Questionnaires & User Feedback

Post-study SUS counts ($M = 1.78$, $SD = 1.47$) and means ($M = 4.47$, $SD = .85$) verified the VR system and the IVE to be generally immersive. In the post-study questionnaires, we also asked if participants felt sick during their time in the IVE (1 = not at all; 7 = I felt very sick). The obtained post-study sickness ratings confirmed the absence of sickness issues ($M = 1.33$, $SD = .58$).

In debriefing, participants described *Drag:on* and its feedback as varied, suitable for many different applications, comfortable, and “*feel[ing]* much more real than

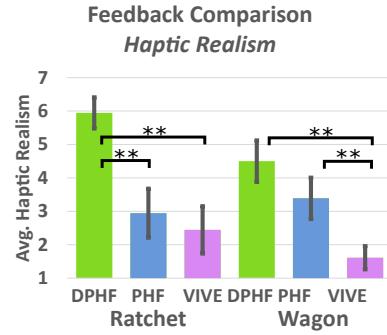


Figure 6.12: The haptic realism experienced in the *DPHF*, *PHF*, and *VIVE* conditions in Part 2 of our study ($p < .05$ (*); $p < .01$ (**)). Error bars show 95% confidence intervals.

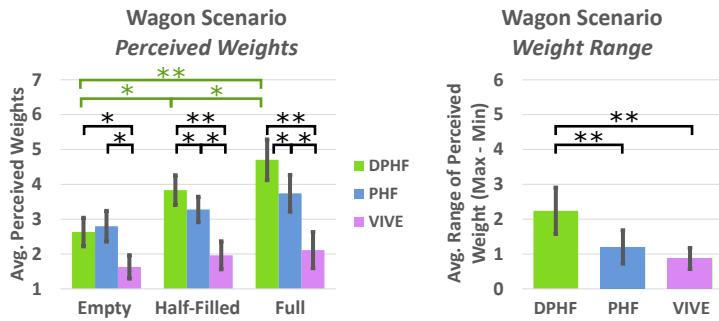


Figure 6.13: Perceived Wagon weights and ranges in Part 2. Brackets indicate pairwise significant differences ($p < .05$ (*); $p < .01$ (**)). Error bars show 95% confidence intervals.

the standard controllers”. When introduced to the concept and prototype, some participants were surprised about how the controller looked and how it worked, and stated that they did not expect it to leverage air resistance.

6.5 Discussion of *Drag:on*

In the following, we discuss our findings in relation to **RQ 2** and our hypotheses. First, we will conclude the results of Part 1, before we turn to the discussion of Part 2. We then use the insights gained in the experiments to derive recommendations and elaborate on potential application areas for *Drag:on*’s haptic feedback. Finally, we will discuss the limitations of our approach.

6.5.1 Discussion of Part 1

The results of Part 1 show that the haptic responses of the tested *Drag:on* states are distinguishable. *Drag:on* can successfully convey different object scales (here, of wooden signs — as illustrated in Figure 6.7) and the (dis)equilibrium of environmental effects like gas streams, given corresponding visual feedback and following the concept of visual-haptic illusions. The comparison of S_{left} , S_{right} and S_{full} shown in Figure 6.10 suggests that asymmetric states are suitable to convey relative resistance differences, especially in conjunction with symmetric states representing the absence of such. Our findings thus confirm **H-S** and **H-F**. Concerning material perceptions, with each of the tested *Drag:on* states, a different material was associated most often. The results indicate that S_{closed} is suitable to convey relatively lightweight materials like plastic, S_{full} is associated with rather heavy materials like wood or metal, and S_{half} can be used to render materials of intermediate weight (like plastic or wood). Comparing, for example, the results for S_{closed} and S_{full} , it can be seen that different states are indeed associated with different materials – confirming also **H-M**.

6.5.2 Discussion of Part 2

The results of Part 2 show that *Drag:on*'s DPHF could also render distinguishable levels of resistance and weight, confirming hypotheses **H- $\{R,W\}$ -DPHF**. The corresponding ranges rendered by *DPHF* were significantly greater than those of *PHF* and *VIVE*, confirming also **H- $\{R,W\}$ -Range**. Moreover, the haptic realism of *DPHF* was significantly higher than that of *PHF* and *VIVE* when rendering ratchet resistances, corroborating **H-R-Realism**. As *PHF* was not found to convey significantly different resistances, nor weights, and as *VIVE* did not yield significantly different sensations of weight, **H- $\{R,W\}$ -PHF** and **H-W-VIVE** were not supported by our results. For *VIVE*, ratchet resistances of air differed significantly from those of plastic and rocks. However, the difference between plastic and rocks could not be communicated with the *VIVE* technique, which delivered a significantly smaller range of resistances than *DPHF*. **H-R-VIVE** was thus only partially confirmed. The same applies to **H-W-Realism** as in *Wagon*, the perceived realism of *DPHF* and *PHF* did not differ significantly.

From the results we conclude that mainly low resistances and weights were perceived with *VIVE*. *PHF* provided slightly higher resistances and weights, but due to its passive nature did not adapt to different materials or fill states. Different perceptions across *PHF*, although not significant, were likely caused by the visualization. In contrast, with *DPHF*, significantly different resistances and weights could be rendered, which significantly increased the haptic realism of the VR experiences compared to using standard VR controllers.

6.5.3 Recommendations

We condense our findings, observations and experiences into a set of basic recommendations for drag and inertia-based haptic VR controllers such as *Drag:on*.

1. To convey haptic impressions, interactions should be designed so as to cause controller movements.
2. To minimize real-virtual discrepancy, virtual objects should align with the fan plane (i.e., so that the plane through their largest surface area coincides with the plane parallel to the fans through the controller).
3. When moved, controller states with
 - min. surface area A and low rotational inertia I_{roll} (e.g., S_{closed}) are suitable to render *small*, *lightweight*, or *empty* objects, or *low mechanical resistances*.
 - max. A and high I_{roll} (e.g., S_{full}) are suitable to render *large*, *heavy*, or *filled* objects, or *high resistances*.
 - intermediate A and I_{roll} (e.g., S_{half}) can render *intermediate states of size, heaviness, filling, or resistance*.

- asymmetric drag properties (like S_{left} , S_{right}) are suitable to render *relative differences in resistance*, resulting in torque felt while swinging the controller.
- symmetric drag properties (like S_{closed} , S_{half} , S_{full}) can be used in contrast to asymmetric states to render the *absence of relative differences*.

4. *Rotational movements* as in Figure 6.3 (a) are suitable to convey a *broad range of resistances*, and *high absolute resistances*, as drag and inertial feedback act in concert.
5. *Translational movements* as in Figure 6.3 (b) are suitable to convey *relations of resistance* through torque and *low absolute resistances*.

6.5.4 Application Areas

Besides the investigated VR interactions, we believe the DPHF concept underlying *Drag:on* can enhance many more VR scenarios. As also suggested by our participants, we imagine *Drag:on* to enhance the realism of VR sport experiences (e.g., curling, racket sports, or golf), and other physical interactions like rowing, swimming, or diving in VR. Its feedback might also suit to simulate the resistances felt when handling tools like screwdrivers, hammers, or axes, or resistances expected during everyday interactions like stirring a pot. Besides realistic scenarios, *Drag:on* could also enhance the feel of unrealistic IVEs. In games, different device states could render the feel of swords or the dense atmospheres of distant planets. Holding a *Drag:on* controller in each hand, participants also suggested to simulate the feeling of being a flying bird. In a commercial controller design, we imagine the user to mount custom fans and weights that either ship with the application, or can be self-fabricated, optimizing the experience.

6.5.5 Limitations

The design of *Drag:on* also comes with certain limitations and drawbacks. As per design, *Drag:on* only provides distinguishable haptic impressions when moved by the user. In our study, participants were instructed to move *Drag:on* naturally as in an actual application. Even though we observed different speeds, our results show that natural interactions suffice to perceive the desired effects. It is noteworthy that these encouraging results were obtained although the feedback curve of *Drag:on* did not match the exact resistance profiles that would have been encountered during the tested interactions in reality. Yet, it is advisable to keep the velocity dependence of *Drag:on*'s feedback in mind when designing corresponding VR interactions.

A mechanical limitation of our device is *Drag:on*'s fixed orientation of the fan plane. When moving the controller parallel to this plane, the drag effect vanishes. While this leads to realistic feedback for rather flat objects, it might be unrealistic

for other shapes. This can, however, be improved in future device iterations by adding an actuator to rotate the top end of the controller. The fans could thereby rotate dynamically around the roll-axis of the device to optimize their angle of attack. Integrating such an actuator would additionally enable decoupling the drag felt when swinging from the resistance felt when rotating. Using a motor to compensate for rotations of the device about its roll-axis, resistance could be felt only during translational movements. Vice versa, the fan plane could be rotated to always coincide with the translation direction to convey resistance only when rolling the device.

Other limitations of our prototype include its relatively high weight, audible and vibrotactile noise as a byproduct of the transformation, and the physical space requirements (e.g., when used in small rooms or during bi-manual interaction). Beyond that, users can perceive the airflow during certain interactions and might perceive the weight imbalance of certain device states. Moreover, its transformation time of *570ms* might still be too slow for some VR interactions. Most of these limitations, however, can be addressed in future device iterations by considering alternative form factors (e.g., origami [Fuchs et al., 2018]) or device designs that adapt drag independent from inertia (e.g., variable fan perforation), or by using lighter materials, faster motors, dampening, or optimized size-to-weight ratios. Beyond that, more advanced fan control could dynamically adjust the device size for collision avoidance, or compensate for velocity disparities between different users.

6.6 Conclusion & Contribution to the Research Questions

In this chapter, we presented the concept and implementation of *Drag:on*, our second novel haptic proxy in the form factor of a handheld VR controller that can provide kinesthetic sensations in IVEs based on DPHF. *Drag:on* dynamically resizes its surface area to leverage the airflow occurring at the proxy during interaction and to adapt its rotational inertia. Following up on our work on *Shift*, with *Drag:on*, we introduced a second mechanism for ungrounded kinesthetic feedback in VR that only requires minimal actuation, and by this we continued to show how the flexibility of proxy-based haptics can be improved. Specifically, our approach combines DPHF based on *drag* and weight shift and allows for implementations characterized by the sole use of low-cost and 3D-printed parts. As such, the concept of *Drag:on*, as well as its prototypical implementation contribute to the second part of **RQ 2**.

In a user study comprising five VR scenarios, we studied how users perceive *Drag:on*'s DPHF and how it compares to the PHF provided by an equivalent passive prop, and to the vibrotactile AHF provided by state-of-the-art HTC Vive Controllers. By this, we evaluated *Drag:on*'s capabilities to convincingly represent a variety of different virtual objects and interactions that feature different kinesthetic characteristics. In particular, we explored rotational and translational

controller movements and showed that *Drag:on* delivers distinguishable levels of haptic feedback. As we found support for most of our hypotheses, including **H-S**, **H-M**, and **H-F**, we demonstrated in Part 1 of our study that our concept and our prototypical implementation provide suitable haptic feedback for virtual objects differing in scale or material, and even for perceiving relative differences in the strength of virtual gas streams. In Part 2, we further confirmed **H-R-DPHF** and **H-W-DPHF** and showed that *Drag:on* also improves the perception of resistances felt when turning virtual ratchets and of the weight felt when moving virtual wagons. All these results indicate the significant potential of drag- and inertia-based DPHF to tackle the *Similarity* challenge by establishing perceptual *Similarity* between a single shape-changing proxy and various virtual objects. Moreover, they highlight that the concept of visual-kinesthetic illusions discussed in section 5.7 also is applicable to drag-based DPHF. As such, our work on *Drag:on* represents another step towards more flexible proxy-based haptics in line with **RQ 2**.

We also found that *Drag:on*'s haptic feedback significantly increased the haptic realism compared to standard VR controllers since we found **H-R-Realism** to be supported and **H-W-Realism** to be partially supported. Moreover, DPHF yielded significantly greater ranges of conveyed resistance and weight than PHF and the vibrotactile AHF of the standard VR controllers, confirming **H-R-Range** and **H-W-Range**. All these observations further emphasize the concept of DPHF to be an adequate solution for the first part of **RQ 2**, capable of tackling the *Similarity* challenge. Finally, we compiled our findings and observations in a set of basic recommendations for the application of drag- and weight-shift-based DPHF and the design of suitable VR experiences. Our results encourage future research to uncover the full potential of haptic feedback for VR based on air resistance and weight shift.

Chapter 7

Concluding Remarks

We conclude Part III with the observation that the two DPHF concepts of changing mass distribution (i.e., inertial changes) and changing shape (i.e., aerodynamic changes) are conceptually orthogonal and could be combined in a single proxy. As the two feedback approaches take effect in different phases of an object's movement cycle, their combination promises to cover all the relevant aspects of virtual object movement. We summarize this relationship in the conceptual sketch in Figure 7.1. When statically holding an object, the *static moment* $M = m \cdot d$, determined by the object's mass m and its weight distribution (with d being the distance between grasp location and CoM), dominates the perceived kinesthetic aspects of the virtual object interacted with [Kingma et al., 2002]. Once the user accelerates the object, in addition to the static moment, also the object's *moment of inertia* I and its aerodynamic properties like *drag* take effect. As a result, the user is perceiving different rotational resistances when wielding the object during interaction [Turvey, 1996; Kingma et al., 2004]. These perceivable resistances during rotation, in turn, can result in visual-haptic illusions when combined with appropriate visualizations in VR, as we could show with our experiments in chapters 5 and 6. After accelerating an object for a while during interaction, for example when swinging the VR controller, it will reach critical velocities at which sufficiently strong air flows have formed around the proxy. Taking advantage of these, further kinesthetic sensations can be produced by varying the prop's aerodynamic properties. Through increasing or decreasing its surface area, for example, varying resisting forces and torques can be induced that are likewise suitable for triggering visual-haptic illusions as we demonstrated in chapter 6. Finally, towards the end of an object's motion, the user will decelerate the object again and the effect of *drag* starts to vanish. At this point, the perception of kinesthetic properties is again controlled primarily by the proxy's inertial state, controlled along with its mass distribution following the concept of DPHF.

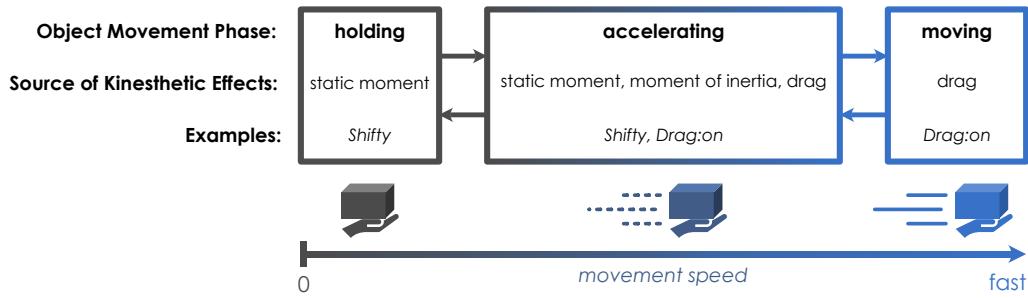
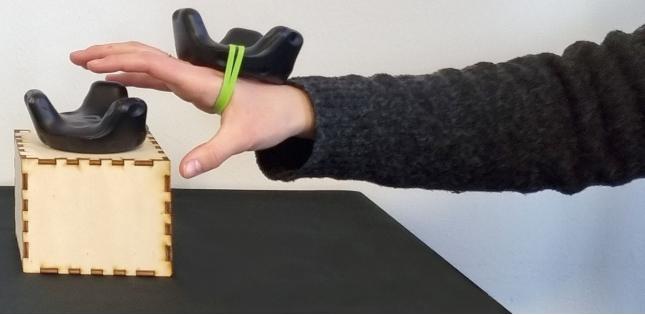


Figure 7.1: The DPHF concepts for ungrounded kinesthetic feedback introduced in Part III with *Shifty* and *Drag:on* cover the entire object movement cycle.

Shifty demonstrated the potential that inertia-based feedback bears for the two phases of *statically holding* and *accelerating* an object. In addition, *Drag:on* showcased that when the user is *moving* the proxy through the air at common movement speeds in the third phase, drag-based feedback is a suitable means to introduce kinesthetic effects. While our *Drag:on* prototype already combined weight-shifting feedback with aerodynamic adaptations for interactions in which users roll the proxy, further research on the combination of both concepts needs to be conducted in the future. With our findings in chapters 5 and 6 in place, future work can now study the combination of weight-shifting and drag-changing DPHF in detail and across the whole movement cycle. In particular, we expect psychophysical research to be of great importance on the path towards holistic feedback solutions, complementing the results summarized by Lim et al. [2021]. Corresponding future experiments could result in the formulation of perception-based models that could drive future DPHF rendering algorithms and the design of novel DPHF proxies and VR controllers.

Part IV

Enhancing Proxy-Based Haptics The *Virtual* Approach



... *seeing* is believing

Research Question 3

Improvement

This part addresses **RQ 3:**

What limitations and potentials does human perception imply for the technique of Body Warping-based Hand Redirection (HR)?

Chapter 8

Continuous Hand Redirection

Part III presented our research on improving proxy-based haptics by taking advantage of techniques operating in the physical environment. In this context, we introduced a novel class of haptic proxies capable of adapting their properties in order to solve the *Similarity* challenge. Yet, as outlined in section 2.4, fulfilling the requirement of *Similarity* is not enough. Instead, the realization of proxy-based haptic feedback also requires the VR system to ensure *Colocation* between proxies and their virtual counterparts.



Video Link⁶⁶

Part IV of this thesis will present fundamental research on techniques that promise to provide solutions for exactly this challenge of *Colocation*. By continuing our research towards an **Improvement** of proxy-based haptics, in this part, we will turn to the real-time virtual technique of hand redirection introduced in chapter 2, specifically to Body Warping-based Hand Redirection (HR), and investigate **RQ 3**. Being a basic building block for techniques like haptic retargeting and redirected touching, results that broaden our understanding of how HR is perceived, and how it can be realized, bear great potential to impact all the techniques that build upon it. Specifically, our research in this part is motivated by three major gaps that we identified in the current state of research, each of which is addressed in its own chapter in the following.

Firstly, our review of related work in chapter 2 revealed that the knowledge of how HR is perceived by users in common scenarios is incomplete. In particular, the degree to which the technique, which relies on intentional mismatches of visual and proprioceptive feedback, can be applied without users noticing it, is still understudied. Previous research only considered either:

- (a) non-conservative scenarios (e.g., HR while the user is playing a game), which might drastically overestimate the range of unnoticeable redirection as considerable distractions are involved [Burns et al., 2006]

- (b) HR while the user is in contact with a physical proxy (e.g., following contours on a shape display [Abtahi and Follmer, 2018] or manipulating parts of a proxy [Feick et al., 2021]), which, according to the findings by Lee et al. [2015], can be assumed to result in enlarged thresholds compared to interactions that lack any haptic signals during redirection
- (c) less common types of hand visualizations [Lee et al., 2015] and HR algorithms, which, for example, apply constant or purely gain-based hand offsets [Burns and Brooks, 2006; Benda et al., 2020; Esmaeili et al., 2020]

Consequently, research still lacks knowledge about Conservative Detection Thresholds (CDTs) that apply even in worst-case scenarios and hold for common scenarios like haptic retargeting, where HR continuously increases the hand offset as the user reaches in mid-air for a target. To fill this gap, in this chapter, we will conduct a corresponding psychophysical investigation of CDTs.

Secondly, past research did not yet exhaustively explore the potential of the perceptual phenomenon of change blindness for HR, which, however, has been successfully exploited to enhance RDW [Suma et al., 2011a; Langbehn et al., 2018b]. To advance the field in this direction, we introduce a novel HR technique in chapter 9 that leverages human eye blinks. Moreover, we conduct a psychophysical investigation to assess its potential and compare it to a state-of-the-art HR method.

Thirdly, while previous research has brought up a range of different hand redirection techniques leveraging body warping (see Table 2.3), world warping, and hybrid warping, up to now, a unified entry point for VR developers and researchers to the domain of hand redirection is missing. While the *Redirected Walking Toolkit* by Azmandian et al. [2016a] closes a similar gap for the domain of RDW, a comparable resource does not exist for hand redirection. To meet this need, we will conclude Part IV by introducing the open-source *Hand Redirection Toolkit (HaRT)* in chapter 10, which is to serve as an accessible framework for novices and experts in the field of hand redirection.

A video⁶⁶ about the work presented in this chapter is available online. This chapter is based on the following publication. Images and parts of the text in this chapter, as well as the presented concepts, implementations, and results have been published previously therein:

Zenner, A. and Krüger, A. (2019b). Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'19, pages 47–55. IEEE. © 2019 IEEE. Final published version available in the IEEE Xplore® Digital Library. DOI: [10.1109/VR.2019.8798143](https://doi.org/10.1109/VR.2019.8798143)

⁶⁶Continuous Hand Redirection Detection Thresholds Video. <https://bit.ly/3loXfJK>

8.1 Introduction

As we have reviewed in chapter 2, previous research introduced several virtual (i.e., software-based) techniques that are applied in real-time to tackle the challenges of proxy-based haptics. In order to achieve *Colocation*, for example, the technique of haptic retargeting has proven itself as highly useful [Azmandian et al., 2016b], and the related approach of redirected touching even achieves to establish perceptual *Similarity* between virtual objects and differently shaped physical props [Kohli, 2013a]. Moreover, techniques like pseudo-haptics rely on visual-proprioceptive illusions in order to convey haptic qualities like virtual weight during the interaction with a virtual object or proxy [Lécuyer, 2009].

All these techniques have in common that they rely on the concept of visual-haptic illusions and the particularities of the human perceptual system. Specifically, the techniques build on the particular strategies that our perception applies in order to cope with the imperfection of our body's biological sensors. In this context, the concept of multisensory integration [Ernst and Banks, 2002] plays a central role as it describes how the human brain integrates stimuli from different senses into a coherent percept, even if the individual sensory inputs are slightly discrepant. As described in section 2.2, in many cases, such sensory mismatches are solved in favor of vision, which means that what we *see* can largely determine what we *perceive*. It is this effect of visual dominance [Gibson, 1933] that represents the basis of Body Warping-based Hand Redirection (HR) techniques, which, in turn, form the foundation of the most common implementations of haptic retargeting and redirected touching.

The idea behind HR is to refrain from a 1-to-1 mapping from real to virtual space. Instead, common implementations of HR displace the virtual hand seen by the user inside the IVE from the position of their physical hand, as reviewed in detail in chapter 2. By this, HR algorithms utilize the fact that the VR system is in full control of what the user sees to introduce a visual-proprioceptive mismatch. As a result of the visual dominance effect, the user will perceive their hand to be where it is *shown*, instead of where it actually is located in the physical world if the mismatch is not too large. When displacing the virtual hand while the user reaches out to touch a proxy, for example, the user will compensate for the displacement during the reach and adapt the path of their real hand accordingly – effectively redirecting the physical hand to an offset destination.

By employing this trick to let users touch misaligned proxy objects, HR can greatly enhance the generality and reusability of proxy-based haptic feedback. However, for common scenarios it remains unclear, how much the user's hand can be redirected before the redirection becomes *too large*. If users consciously notice the mismatch, the VR system risks BIPs [Jerald, 2015; p. 48], which can destroy the believability of the experience. To prevent this, it is essential to know about the Detection Thresholds (DTs) of common HR techniques [Kohli, 2009]. In particular, as described by Kohli [2013a], there is a need for a formal investigation of DTs

for HR in VR. Knowledge about the detectability will allow for suitable proxy mappings and for the design of appropriate virtual and physical environments that make use of haptic retargeting and redirected touching to improve the flexibility of proxy-based haptics.

In this chapter, we derive corresponding estimates and determine the order of magnitude at which HR can go *unnoticed* when reaching for a virtual target under redirection. Our investigation is motivated by the haptic retargeting use case and aims to investigate a common HR scenario [Azmandian et al., 2016b; Cheng et al., 2017b]. To establish Conservative Detection Thresholds (CDTs) for desktop-scale HR, we conduct a psychophysical experiment. Our thresholds are to be classified as conservative since we explicitly explain to participants how the evaluated HR techniques work, and instruct them to pay careful attention to detecting corresponding hand offsets. Moreover, in contrast to related work, we focus our investigation on a common HR approach that *continuously* increases the hand offset during the reaching movement and displays only a realistic virtual hand visualization [Dewez et al., 2021]. Furthermore, we study the case of mid-air approach, i.e., a hand movement towards a target location, without any haptic signal present at the user's hand during redirection. Our study explores horizontal, vertical, and gain-based redirection in three different conservative interaction scenarios that differ in the user's distraction from the redirection. Combining these results, we formulate general recommendations regarding the amount of HR that can be applied unnoticeably even in worst-case desktop-scale redirection scenarios.

8.2 Continuous Hand Redirection

The HR approach investigated in this chapter, referred to as *Continuous Hand Redirection*, represents one of the most common HR approaches. It maps a single point in the physical space to an offset point in virtual space and gradually displaces the virtual hand rendering as the physical hand closes in towards the physical target position. Table 2.3 classifies our redirection approach in comparison to related work. For our investigation in this chapter, we chose such a simple HR technique as it is comparably easy to implement, versatile, and characterized by a low complexity. Our warping algorithm, which computes in each frame the amount of offset that the virtual hand is shifted from the tracked location of the real hand, is based on a geometric calculation and allows us to study the CDTs of such HR in a controlled fashion. Following the thoughts of Kohli [2009], we investigate the detectability of HR along different directions individually, splitting up 3D redirection in three intuitive dimensions. The following sections introduce the corresponding warping algorithms that proved well suited for our study, as they allow for horizontal and vertical angular redirection, as well as gain-based displacement of the virtual hand.

Algorithm 1 Pseudocode of the Rotational Warp Algorithm

Input: real hand position H_p , warp origin O , unit forward vector \hat{f} , unit redirection vector \hat{r} , redirection angle α

Output: virtual hand position H_v

```

1: procedure ROTATIONALWARP( $H_p, O, \hat{f}, \hat{r}, \alpha$ )
2:    $\hat{h} = \hat{f} \times \hat{r}$                                       $\triangleright$  compute unit height vector
3:    $height = (H_p - O) \cdot \hat{h}$                        $\triangleright$  save height
4:    $\vec{p_{proj}} = H_p - height \cdot \hat{h}$             $\triangleright$  project on redirection plane
5:    $d_{proj,r} = \vec{p_{proj}} - O$                        $\triangleright$  unwarped offset in plane
6:    $\alpha_r = atan2(\vec{d_{proj,r}} \cdot \hat{r}, \vec{d_{proj,r}} \cdot \hat{f})$      $\triangleright$  angle rel. to  $\hat{f}$  &  $O$ 
7:    $\alpha_v = \alpha_r + \alpha$                              $\triangleright$  adding angular offset
8:    $\vec{d_{proj,v}} = \sin(\alpha_v) \cdot |d_{proj,r}| \cdot \hat{r} + \cos(\alpha_v) \cdot |d_{proj,r}| \cdot \hat{f}$      $\triangleright$  warped offset in plane
9:    $H_v = O + \vec{d_{proj,v}} + height \cdot \hat{h}$          $\triangleright$  final warped position
10:  return  $H_v$ 
11: end procedure

```

8.2.1 Horizontal Warping

The first investigated type of HR horizontally offsets the virtual hand (H_v) by a warp angle α as the real hand (H_p) moves away from a warp origin (O). To compute the warped position, the physical hand is projected onto a horizontal plane, its angle relative to a forward direction and the warp origin is incremented by α and then projected back into 3D space.

To implement this behavior, we propose a generic rotational warp algorithm allowing for displacements in arbitrary planes defined by a unit forward vector \hat{f} and an orthogonal unit redirection vector \hat{r} indicating the direction of positive displacement. Further inputs are the location of the warp origin O (i.e., the physical hand position when the warp starts) and the redirection angle α . Algorithm 1 shows the pseudocode of the rotational warp algorithm. For horizontal displacement as investigated here, the algorithm is instantiated with $\hat{f} = +\vec{z}$ (z-axis) and $\hat{r} = +\vec{x}$ (x-axis). Figure 8.1 shows the effect of the horizontal warp algorithm.

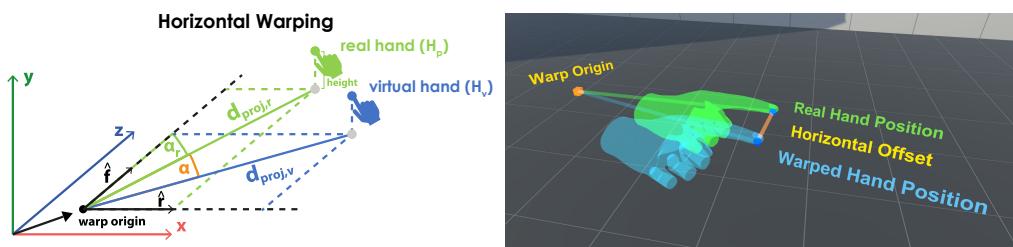


Figure 8.1: Illustration of horizontal warping. The warp origin is shown as a yellow dot, the displaced hand is visualized in blue, and the real hand in green. Study participants only saw the warped hand rendered with a realistic texture. © 2019 IEEE.

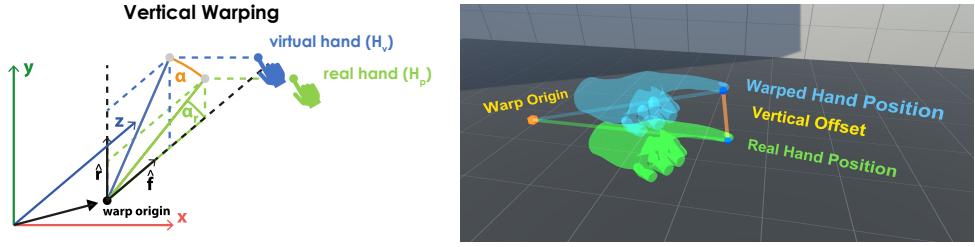


Figure 8.2: Illustration of vertical warping. © 2019 IEEE.

8.2.2 Vertical Warping

Vertical HR offsets the virtual hand up or down as the real hand moves away from the warp origin. For this, Algorithm 1 is instantiated with $\hat{f} = +\vec{z}$ (z -axis) and $\hat{r} = +\vec{y}$ (y -axis). Figure 8.2 shows the effect of the vertical displacement.

8.2.3 Gain Warping

The third investigated HR algorithm scales the distance of the real hand from warp origin. It computes the distance vector \vec{d}_r to the unwarped position of the real hand and applies a gain factor g , effectively decreasing (if $0 < g < 1$) or increasing (if $g > 1$) the speed of the hand moving away from O . Algorithm 2 sketches the pseudocode. The effect of the warp is illustrated in Figure 8.3.

8.3 Evaluation of *Continuous Hand Redirection*

To study the detectability of *Continuous Hand Redirection*, we assumed a desktop VR setting, with the user being seated and interacting in the limited space in front. This setup resembles common scenarios for haptic retargeting [Azmandian et al., 2016b] and redirected touching [Kohli, 2013a]. We conducted an experiment investigating the three individual redirection dimensions (horizontal warping, vertical warping, and gain warping along the depth axis), each in three different scenarios. In the experiment, participants were immersed in a simple IVE with their hand being tracked. In nine conditions, they repeatedly performed a simple

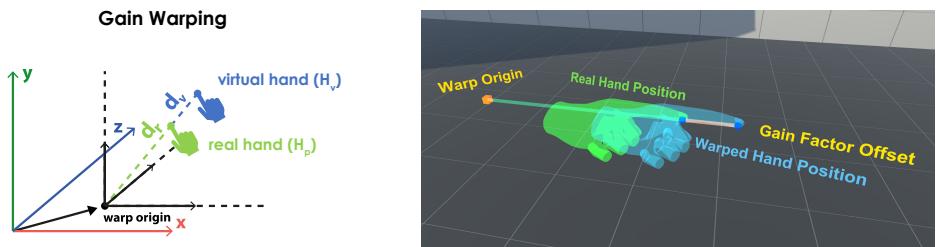


Figure 8.3: Illustration of gain warping. © 2019 IEEE.

Algorithm 2 Pseudocode of the Gain Warp Algorithm

Input: real hand position H_p , warp origin O , gain factor g **Output:** warped virtual hand position H_v

```

1: procedure GAINWARP( $H_p, O, g$ )
2:    $\vec{d}_r = H_p - O$                                  $\triangleright$  unwarped offset from origin
3:    $\vec{d}_v = g \cdot \vec{d}_r$                            $\triangleright$  warped offset from origin
4:    $H_v = O + \vec{d}_v$                              $\triangleright$  final warped position
5:   return  $H_v$ 
6: end procedure

```

interaction with different warps applied, and had to state the direction of the hand displacement. From the results, we derived how much HR could go unnoticed. The experiment was approved by the ethical review board of our faculty.

For our investigation, we employed the method of constant stimuli with a symmetric 1AFC task (see Table 2.4). In this design, participants are repeatedly exposed to different amounts of HR while reaching with their hand in mid-air to a virtual target in front of them. Throughout all conditions, the main task of the participants was to determine the *direction* of the hand offset and they were instructed to fully concentrate on that task. Moreover, they were informed about the investigated redirection technique, and knew how it worked. As a consequence of this priming, all our investigated scenarios are classified as conservative, representing the worst case for unnoticeable redirection. Moreover, to derive meaningful detection limits, we decided to not only investigate low-level perceptibility in a single very conservative scenario where participants only focused on detecting the offset. Instead, we additionally considered the noticeability of redirection in two further scenarios employing secondary tasks and visual distraction, as well as auditory and vibrotactile distractions, respectively.

Scenario 1: No Distraction The first and most conservative scenario tested in our experiment did not distract the participants from their main task at all. Participants did not hear anything, they had no second task, and no vibrotactile cues were present. Thus, the first scenario is well suited to derive the most conservative lower bounds for the DTs, but at the same time is also less realistic.

Scenario 2: Audio & Vibration Distraction The second scenario better represents realistic application scenarios. As distractions might influence the detectability of HR [Kohli, 2009], here, we distract users with a combination of two additional modalities likely used in VR applications: auditory and vibrotactile feedback. Specifically, spatial sound and four head-mounted vibration cells let participants experience a virtual bee orbiting their head during the redirection. We used vibrotactile actuation at the head as we expected this to yield strong distraction and believe it will be included in future generations of HMDs. Our

vibration cell placement was based on the results of Myles and Kalb [2010], who showed that the forehead, occipital, and temple regions are suited best for vibrotactile cues by being most sensitive to vibration.

Scenario 3: Visual & Dual-Task Distraction Besides distraction by special effects, distraction can also ensue due to increased cognitive load. Thus, our third test scenario included both a second task to solve in parallel to the interaction and a visual distraction forcing users to look away from the virtual hand at least once. For this task, we were inspired by a common application scenario for redirected touching: the simulation of cockpit procedures [Kohli, 2013a]. Here, users must regularly look away from their hand in order to read cockpit instruments. Thus, in our scenario, participants had to look at a number displayed on a virtual panel in front of them to read it out loud while being redirected. Similar to Scenario 2, this represents a more realistic use case than Scenario 1.

8.3.1 Participants

12 participants volunteered to take part in the experiment (6 female, 6 male, avg. 28 years, between 20 years and 61 years old). 4 participants wore glasses or contact lenses, but all participants confirmed that they have normal or corrected-to-normal vision, that they do not suffer from hearing impairments and that their sensation of vibrotactile feedback is not in any form negatively affected. 11 participants were right-handed and 1 participant was left-handed. The participants rated how regularly they play 3D video games on a scale from 1 (= never) to 7 (= regularly). Here, different frequencies were present, with answers between 1 and 7 ($M = 2.33$, $SD = 2.15$). Participants also stated how often they use VR technology on the same scale. We obtained answers between 1 and 5 ($M = 2.09$, $SD = 1.38$). A third question on the same rating scale asked how often participants perform precise handicrafts. Here, all different ratings were obtained ($M = 3.75$, $SD = 2.05$).

8.3.2 Apparatus

The experiment was conducted in our lab. We used an HTC Vive¹⁴ HMD with two HTC Vive Controllers⁵³ and a separate HTC Vive Tracker⁵⁴. One controller was used by the participants to record their 1AFC answers, the other one was used for initial finger calibration. The IVE was developed with the Unity3D 5.6 engine¹⁷ and represents a small, plain terrace-like scene as shown in Figure 8.7. It was intentionally kept simple to prevent uncontrolled distraction. A desktop computer executed the VR application and logged the 1AFC responses. The experimenter logged all additional answers in a separate spreadsheet.

Headphones were used for auditory feedback and four small vibration cells (Adafruit Vibrating Mini Motor Disc #1201; $10\text{mm} \times 2.7\text{mm}$) controlled by a WeMos



Figure 8.4: (a) Vibration cells at the temple and lower occipital regions close to the neck. (b) Finger splint and the HTC Vive Tracker used for hand tracking. (c) *Vis-DT* condition. Participants reach for the green target with the virtual hand while reading out the displayed number. Real hand added for illustration. All images ©2019 IEEE.

D1 mini microcontroller delivered tactile feedback. They were fixed with medical tape at the two temple regions and at the lower occipital region close to the neck [Myles and Kalb, 2010] as depicted in Figure 8.4 (a). The VR application controlled the vibration by communicating wirelessly with the microcontroller. In informal previous tests, we tested several vibration patterns and strengths to find suitable parameters for the experiment. We fixed the vibration strength to a comfortable but readily noticeable level. When called by the VR application, the microcontroller repeatedly activated a random set of the four cells for a random duration between 150ms and 500ms until the command to stop was received.

We used the HTC Vive Tracker for hand tracking, and during calibration attached the tracker on the back of the user's hand with a rubber band. We additionally used a finger splint that allowed participants to comfortably maintain a pointing hand position with the index finger pointing forward, as shown in Figure 8.4 (b). Besides being comfortable and reducing fatigue, this splint ensured that the real and virtual hand stayed spatially registered at all times as the rigid structure of the splint did not allow participants to move the index finger relative to the tracker on the back of the hand. Thus, the offset from tracker location to fingertip remained fixed. The touch-sensitive trackpad of one of the HTC Vive Controllers was used in an initial calibration step to calibrate this offset and to align the virtual hand model with the real hand. This spatial registration procedure represents a dynamic variant of the calibration used by Kohli [2010]. We rendered either a female or a male, and a right or a left virtual hand to account for the participant's gender and handedness. Both male and female hands had a fixed but realistic size and the male hand was slightly larger than the female hand model.

8.3.3 Procedure

Each participant was initially briefed with a prepared text explaining the concept of *Continuous Hand Redirection* and the course of the experiment. The experimenter attached the finger splint and the HTC Vive Tracker on the back of the participant's hand, and the four vibration cells at the participant's head.

Participants sat on a chair wearing headphones and the HMD. Upon entering the IVE, participants were introduced to all three redirection types in a short training session. To ensure that the participant understood and noticed the applied redirection, we demonstrated the largest redirection used in the tests during training (values determined by informal pre-testing; $\alpha = +/ - 14^\circ$ for horizontal and vertical redirection, $g = 0.75$ and $g = 1.25$ for gain redirection).

At the same time, the interaction to be performed by participants during each trial was practiced: seeing a small green sphere appearing at the start location 30cm beneath and 30cm in front of the head, the participant was supposed to touch the sphere with the virtual index finger. When touched, HR was applied with the warp origin set to the location of the start sphere. We chose the origin to be at this comfortable position in front of the user, as it is just outside the zero-warp zone defined by Cheng et al. [2017b], and we see this as a representative location for desktop setups. Upon activation of the warp, the sphere relocated to the target position. The distance to the target was likewise chosen to be representative for typical desktop-scale interaction distances. For the horizontal and vertical redirection, this target position was another 40cm away from the start position straight in front of the user. To account for gain factors $g < 1$ where participants had to reach further with their real hand, the target position, when gain warping was applied, was only 30cm from the start location to ensure reachability.

Upon relocation of the sphere, participants were to move their hand in order to touch the target sphere with the virtual finger. This required participants to compensate for the redirection warp. To ensure consistency and comparable hand movement speeds, each participant was asked to perform this movement within around 3s to 4s. Finally, when reaching the target with the virtual hand, a question appeared on the controller held in the second hand. This symmetric 1AFC question asked participants to state in which direction the virtual hand was displaced during the movement. Participants had to decide between the answers *right* or *left* for *horizontal* displacement and between *up* or *down* for *vertical* displacement. In the *gain* condition, participants had to state whether the virtual hand moved *faster* or *slower* than the real hand. Answers were logged by pressing the corresponding controller button. While equivalent to the experiment procedure, no data was recorded during training. In a final training session, participants could practice reading out numbers displayed in front of them to become acquainted with the dual task in Scenario 3.

Once training was completed, the actual experiment started. To test the three types of redirection (in the following abbreviated *Horiz* for horizontal, *Vert* for vertical, and *Gain* for gain warping) in the three introduced scenarios (abbreviated *None* for no distraction, *Audio-Vib* for audio & vibration distraction, and *Vis-DT* for visual & dual-task distraction), each participant performed nine runs in sequence. Following the method of constant stimuli, in each run, the interaction was performed repeatedly with different redirection parameters α or g applied, and after each interaction, a 1AFC question was answered.

In *Vis-DT* runs, participants additionally read out a random 4-digit positive integer while reaching, which appeared on a display in front when the hand progressed 20% along the way from start to target as shown in Figure 8.4 (c). This forced them to look away from the hand at least once during the redirection. We logged the correctness of the read numbers. In contrast, the distracting virtual bee in *Audio-Vib* did not require participants to take any specific action. They just had to try to focus on the main task of determining the displacement direction.

After completing a run, participants denoted their agreement with nine concluding statements on a Likert scale from 1 (= strongly disagree) to 7 (= strongly agree), displayed in the IVE. These statements assessed their subjective impressions of interacting under redirection in the experienced condition. After completion of all nine runs, the participants filled out a SUS presence questionnaire [Slater et al., 1994] and a post-study questionnaire. The duration of the experiment was ca. 90 minutes, including introduction, calibration, training, all experimental runs, final questionnaires, and debriefing.

8.3.4 Design

Our psychophysical 1AFC experiment is designed as a within-subjects study. We distinguish two independent variables: (1) HR direction (*Horiz*, *Vert*, and *Gain*), and (2) user distraction type (*None*, *Audio-Vib*, and *Vis-DT*). Using a full-factorial design, we investigated nine conditions.

We assessed 11 dependent variables: (1) the perceived offset of the virtual hand as a symmetric 1AFC answer for each trial, (2) – (10) the nine subjective measures assessed as ratings on the Likert scale after each run, and (11) the interaction time measured to reach the target.

The order of the conditions was counterbalanced across participants. We used a Latin square counterbalancing over the three distraction types and for each distraction type, we used an additional Latin square counterbalancing over the three HR directions. The resulting counterbalancing of size $n = 6$ was completed exactly twice with 12 participants.

In each of the three *Gain* conditions, the interaction was performed 22 times resulting in 22 samples (1AFC answers). Using a step size of 0.05, we tested all 11 values between $g = 0.75$ and $g = 1.25$ (inclusive) twice in a randomized order. Similarly, each condition applying *Horiz* or *Vert* redirection collected 30 samples testing all redirection angles between $\alpha = -14^\circ$ and $\alpha = 14^\circ$ (inclusive) in steps of 2° twice in a randomized order. Thus, each participant contributed $6 \cdot 30 + 3 \cdot 22 = 246$ samples. With 12 participants, we obtained $12 \cdot 246 = 2952$ samples for the psychometric analysis in total.

8.3.5 Results

First, we summarize our estimates for the HR DTs of all nine individual conditions. We also derive overall thresholds for the three HR directions from the pooled samples of all three tested scenarios. Secondly, we summarize the results of the subjective responses.

Detection Thresholds

To analyze the performance of our participants in discriminating the hand offset direction, we used the obtained 1AFC answers to fit a psychometric function modeling the discrimination performance over the applied redirection. Analogous to how Steinicke et al. [2010b] derived DTs for RDW, we used a psychometric sigmoid function as its shape is a good approximation to model human response. Plotting the overall probability of our participants answering “The virtual hand was offset to the *right/up/was faster*” against the applied amount of virtual hand offset yields an s-shaped distribution of our samples. To derive the DTs, we fitted the sigmoid function

$$f(x) = \frac{1}{1 + e^{-\frac{x-a}{b}}}$$

to our sample distribution by optimizing the parameters a and b . We computed the Point of Subjective Equality (**PSE**) indicating when the virtual movement is perceived as equal to the physical movement. This is where f intersects the 50% probability. Additionally, we computed the discrimination thresholds, i.e., where f intersects the probability halfway between the random guessing level and the correct answer. For the upper redirection limit

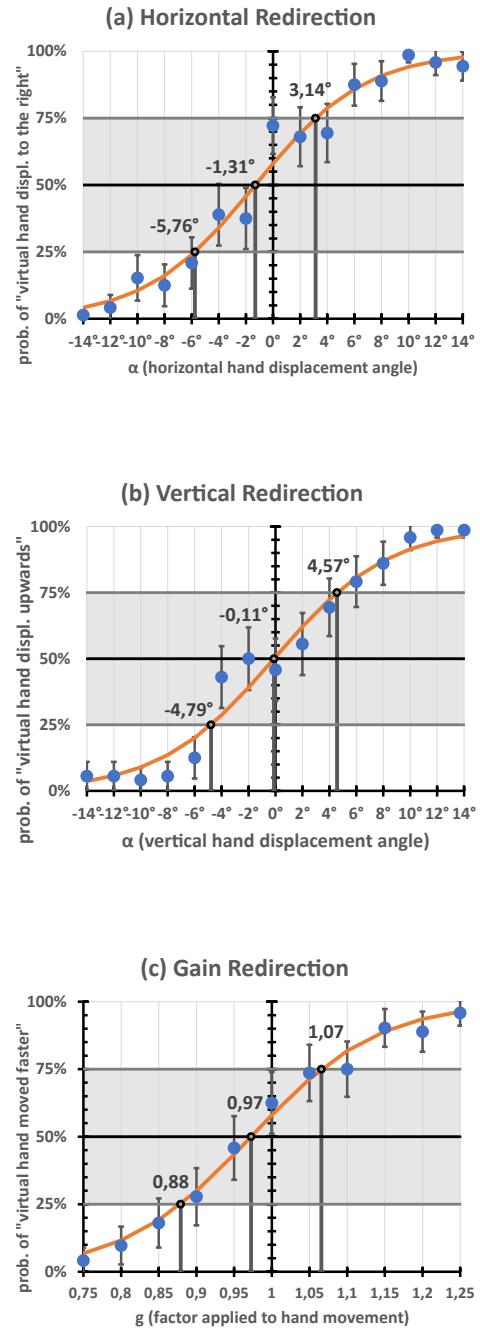


Figure 8.5: Pooled results of all three distraction scenarios for each HR direction. Charts show 95% confidence intervals, fitted psychometric functions f , PSEs, and derived DTs for (a) horizontal, (b) vertical, and (c) gain-based HR. All charts © 2019 IEEE.

		<i>None</i>	<i>Audio-Vib</i>	<i>Vis-DT</i>	<i>Mixed</i>
Horizontal Warping	Right	+3.81°	+2.26°	+2.94°	+3.14°
	PSE	-0.28°	-1.67°	-2.11°	-1.31°
	Left	-4.38°	-5.60°	-7.17°	-5.76°
	Range	8.19°	7.86°	10.11°	8.90°
Vertical Warping	Up	+4.48°	+4.55°	+4.65°	+4.57°
	PSE	+0.04°	-0.09°	-0.31°	-0.11°
	Down	-4.40°	-4.74°	-5.28°	-4.79°
	Range	8.88°	9.29°	9.93°	9.36°
Gain Warping	Faster	1.07	1.05	1.07	1.07
	PSE	0.98	0.96	0.98	0.97
	Slower	0.88	0.86	0.89	0.88
	Reach +	+13.38%	+15.71%	+11.92%	+13.75%
	Reach -	-6.50%	-5.11%	-6.63%	-6.18%

Table 8.1: Derived DTs for *Continuous Hand Redirection*. © 2019 IEEE.

(**Right** for *Horiz*, **Up** for *Vert*, and **Faster** for *Gain*), this is the redirection at which f intersects the 75% probability. Complementarily, for the lower redirection limit (**Left** for *Horiz*, **Down** for *Vert*, and **Slower** for *Gain*), this is where the 25% probability is intersected. The range in between these two amounts is the range of redirection that can go unnoticed, as users could not with certainty tell about the redirection. With the results of the *Gain* condition, we also computed the range within which the user's *real* hand can be unnoticedably redirected to reach further (**Reach +**) or less far (**Reach -**) than the virtual hand. The results are summarized in Table 8.1.

Continuing our analysis, we additionally used the samples of all three scenarios to derive even more robust estimates of the DTs. As the scenarios used different feedback channels for the distractions that can thus be regarded as orthogonal to each other, we pooled the samples of *None*, *Audio-Vib*, and *Vis-DT*. We thereby derived more realistic, but still conservative estimates while profiting from the tripled amount of samples per tested redirection. The resulting plots for this mixed scenario, the fitted function f , and the derived thresholds, are depicted in Figure 8.5 and summarized in Table 8.1 (column *Mixed*).

Distraction & Subjective Impressions

To further study the perceived degree of distraction in our different scenarios and the impact that our applied HR had on the quality of the VR experience, we conducted non-parametric Friedman tests to check for significant differences among conditions regarding each of the nine Likert scale ratings. For this analysis, we applied a significance level of $\alpha = .05$. Additionally, to check for significant effects of the two factors (HR direction and distraction type), we conducted Friedman tests comparing the results of the three HR directions and the three distraction scenarios, respectively. Significant differences indicated by the Friedman tests were investigated with pairwise Bonferroni-corrected Wilcoxon signed-rank tests in post-hoc analysis.

To verify our choice of distraction scenarios, we assessed the responses to

- **Distraction** ($M = 3.21, SD = 0.75$)⁶⁷:

"I felt distracted from the main task (determine hand offset direction)."

and found significant differences among the nine conditions ($\chi^2(8) = 59.455, p < .001$), among distraction types ($\chi^2(2) = 17.733, p < .001$) and among HR directions ($\chi^2(2) = 7.515, p = .023$). Post-hoc analysis (corrected significance level set at $p < .017$) confirmed the distraction of *Audio-Vib* ($M = 3.83, SD = 1.61$) ($Z = -2.941, p < .001, r = 0.60$) and *Vis-DT* ($M = 4.19, SD = 0.93$) ($Z = -3.063, p < .001, r = 0.62$) to be significantly higher than the distraction in our baseline scenario *None* ($M = 1.61, SD = 0.71$), as shown in Figure 8.6. Concerning HR directions, *Vert* ($M = 3.47, SD = 1.00$) was found to be significantly more distracting than *Gain* ($M = 2.97, SD = 0.73$) ($Z = -2.588, p = .008, r = 0.53$).

We did not find any meaningful significant differences in the participants' agreement regarding the following eight statements:

- **Disturbing Offset** ($M = 3.41, SD = 1.14$):

"The fact that the virtual hand representation was offset from the real hand position disturbed me."

- **10 Min.** ($M = 4.88, SD = 1.24$):

"I would not mind working under these conditions for a short amount of time (ca. 10 minutes)."

- **2 Hours** ($M = 3.08, SD = 1.14$):

"I would not mind working under these conditions for a longer amount of time (ca. 2 hours)."

- **Physical Exertion** ($M = 2.78, SD = 1.39$):

"The interaction was physically demanding."

- **Mental Exertion** ($M = 3.75, SD = 0.90$):

"The interaction was mentally demanding and I had to concentrate a lot."

- **Body Ownership** ($M = 4.78, SD = 1.04$):

"I had the feeling of interacting with my own hand in the virtual environment."

- **Hand Control** ($M = 4.88, SD = 1.20$):

"I had full control over the movements of the virtual hand at all times."

- **Sickness** ($M = 1.29, SD = 0.63$):

"When interacting, I felt uncomfortable (e.g. nausea, dizziness)."

We further conducted a Wilcoxon signed-rank test comparing the overall average willingness to work for a short period of time (**10 Min.**) ($M = 4.88, SD = 1.24$) against the willingness to work for a longer period of time (**2 Hours**) ($M = 3.08, SD = 1.14$) in a warped space and found the difference to be significant ($Z = -2.983, p = .001, r = 0.61$), as shown in Figure 8.6. Regarding interaction **Times**, we requested participants to complete the movement in 3s to 4s. The obtained timing

⁶⁷The single M and SD values reported here for each measure are of the *overall* ratings, i.e., participant-wise averages over all nine conditions.

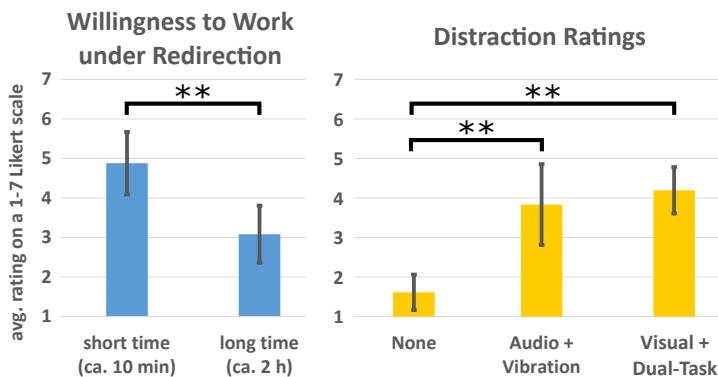


Figure 8.6: Left: Willingness to work for a short time (10 Min.) and a long time (2 Hours) with HR. Right: Perceived **Distraction** in our scenarios. Error bars depict 95% confidence intervals. Brackets indicate pairwise significant differences considering the normal p -value for the left chart and the Bonferroni-adjusted p^* -value for the right chart ($\{p, p^*\} < .05$ (*); $\{p, p^*\} < .01$ (**)). Both charts © 2019 IEEE.

measurements ($M = 3172ms$, $SD = 599ms$) verified comparable hand movement speeds as participants on average managed to perform the hand movement in the desired time. The dual task performance (i.e., correctly reading out the displayed number during interaction) was in each of the three *Vis-DT* conditions $> 95.5\%$. A Friedman test did not discover a significant difference across HR directions.

Post-study SUS presence counts ($M = 1.33$, $SD = 0.89$) ranged from 0 to 3. In a concluding questionnaire, we also asked participants on a scale from 1 (= not at all) to 7 (= I became very sick) for any experienced discomfort during the study. Results ($M = 1.75$, $SD = 1.21$) did not indicate any problems and are in line with participants' debriefing comments. Overall, in verbal debriefing, most participants were enthusiastic about the technique of HR.

8.4 Discussion of Continuous Hand Redirection

Our results provide insights into the order of magnitude at which HR based on continuous warping of the hand can go unnoticed. The following sections discuss these results and provide a thorough comparison to results of related research.

8.4.1 Detection Thresholds

The primary goal of our investigation was to capture the range of HR that can be applied in VR applications without the user being aware of it. The experiment showed that for both *Horiz* and *Vert* redirection, the range of unnoticeable warping is very similar. As a general reference, we propose the estimates derived from the pooled results of all three distraction types, i.e., the thresholds denoted in column *Mixed* in Table 8.1. Considering the *Horiz* redirection, we believe the

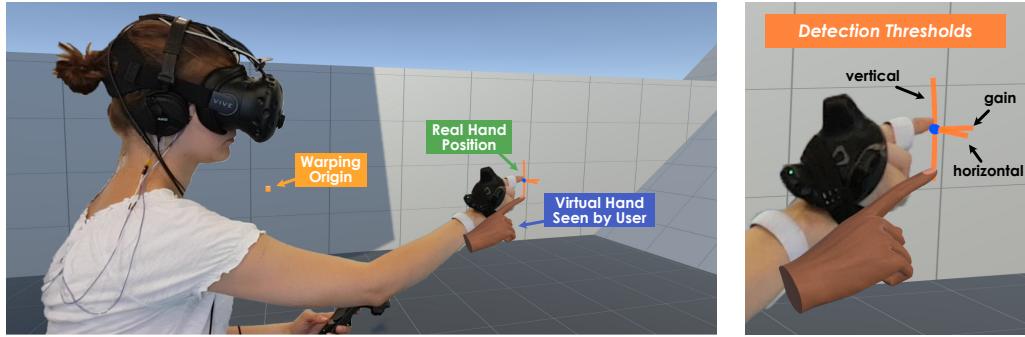


Figure 8.7: The ranges of unnoticeable HR found for worst-case redirection scenarios. According to our results, the virtual fingertip can, in worst-case scenarios, unnoticeably be displaced to any point along the orange lines. Both images ©2019 IEEE.

PSE bias (increasing to the left from -0.28° to -2.11° with increasing distraction) to be due to the predominance of right-handed participants. To derive recommendations, we thus take into account the derived angles relative to the PSE and conclude that in desktop-scale HR, the virtual hand can unnoticeably be displaced *horizontally* or *vertically* by $\approx 4.5^\circ$ to the *left/downwards* or to the *right/upwards* respectively, covering a range of $\approx 9^\circ$. Figure 8.7 visualizes this range in 3D.

Similar results were obtained for the *Gain* redirection, as here too, no noticeable increase of the redirection range was found with increasing user distraction. Users, however, seemed to detect an accelerated virtual hand better than a decelerated virtual hand. The possible downscaling range within limits ($g = 0.88 \rightarrow -12\%$) is almost double the possible upscaling range ($g = 1.07 \rightarrow +7\%$). From the DTs, we can derive how the *real* hand movement is affected when redirection is applied. This knowledge is of immediate value for haptic retargeting applications. Compensating for the discrepancy between real and virtual hand location, users grasp farther when trying to reach a virtual target with deceleration ($g < 1$) applied. Knowing from our results that a factor of $g = 0.88$ is still within limits implies that the user reaches $\frac{1}{0.88} \approx 1.1375$ times the distance to the virtual target. In turn, when accelerating the virtual hand ($g > 1$), the real hand only needs to reach $\frac{1}{1.07} \approx 0.9382$ times the virtual distance. We summarize our recommendations for HR DTs using the investigated *Gain* warp technique as follows: factors between $g = 0.88$ and $g = 1.07$ can unnoticeably be applied to the hand distance from the warp origin. This means that the user's real hand can be redirected unnoticeably to grasp up to 13.75% *further* or up to 6.18% *less far* than the virtual hand.

Based on previous results by Burns et al. [2006], we anticipated detection performance to decrease rapidly as the distraction increased. Our results (see Table 8.1), however, did not deliver strong evidence for that, although our distraction methods did work, as proven by the reported **Distraction** ratings (see Figure 8.6). While we could generate significantly higher distraction than in the *None* scenario, all our scenarios were likely still too conservative to substantially increase

the difficulty of detecting hand offsets. We believe the significant differences to the results of Burns et al. [2006] ($\approx 19.1^\circ$ when users were primed) to stem from differences in methodology, scenario, degree of user distraction, and/or type of offset investigated (gradually growing vs. fixed α).

However, our results support the JND values found by Lee et al. [2015] ($\approx 5.2\text{cm}$) which simultaneously rendered both the real and the displaced fingertip position as a sphere to derive requirements for finger tracking systems. Assuming the distance of 40cm from origin to target tested in our experiment, our estimation of 4.5° yields similar thresholds of $\approx 3.1\text{cm}$. Moreover, regarding the *Gain* condition, our results (DT of $g = 0.88$ for 30cm movements) were recently supported by the findings of Hartfill et al. [2021], who derived a DT of $g = 0.727$ for a similar *Gain* warping algorithm when testing movement distances of only 20cm .

Comparing our results with those of Abtahi and Follmer [2018] (horizontal remapping $\approx 49.5^\circ$, horizontal gain factor of ≈ 1.9 , vertical gain factor of ≈ 3.2) yields additional interesting insights. We believe the significantly higher DTs recorded in their experiment, compared to our results, to stem from the additional, continuous haptic feedback provided to the user's fingertip during the movement – a conclusion further supported by the results of Lee et al. [2015]. Feeling and seeing an edge underneath the fingerpad in Abtahi and Follmer's experiment probably served as a sort of haptic guidance, which might have increased the user's confidence in *not being redirected*. In this regard, we think that Abtahi and Follmer's study (with tactile signals during redirection) and ours (without tactile signals) complement each other. Considered together, they provide a more complete picture of the roles of visual and haptic feedback for HR.

Furthermore, a comparison of our results to those of Benda et al. [2020] (on constant hand offsets) and Esmaeili et al. [2020] (on scaled hand movements) yields interesting questions. In contrast to our conservative investigation of angular-based redirection, both found DTs to vary significantly between horizontal and vertical HR. We suspect this to be due to differences in terms of the HR technique employed as well as the tested scenario. Moreover, when applying constant hand offsets, Benda et al. [2020] found offsets that force users to grasp further to be noticed more easily than offsets that reduce physical reach distance – a finding that is in contrast to our results for gain-based HR. We suspect this difference to be due to differences in hand *speed* that occur when gain-based redirection is applied, in contrast to HR techniques that rely on fixed hand offsets and do not modify the speed of the visual hand. This phenomenon might further be related to the findings of van Beers et al. [2002], who found indication that "*in the depth direction, proprioception is weighted more heavily than vision*".

Lastly, compared to the *tolerance* ranges for HR ($\approx 40^\circ$) identified by Cheng et al. [2017b], our derived *imperceptibility* ranges are significantly smaller. We conclude that human hand-eye coordination is so good that even small discrepancies can be detected when the assessment of the hand movement only relies on visual feedback (i.e., no haptic guidance is present) and attention is paid.

8.4.2 Distraction & Subjective Impressions

The results for **Distraction** confirmed our scenario design as **Distraction** was significantly higher in *Audio-Vib* and *Vis-DT* compared to the baseline scenario *None*, even if *Audio-Vib* and *Vis-DT* were likely not distracting enough to show that DTs can be increased through user distraction. However, we still are confident that decreasing the user's attention on HR (e.g., through an engaging second task or additional stimuli), or not telling the user about HR being applied, can allow for a larger redirection to go unnoticed. This is supported by the comparison of our findings to related research [Matsuoka et al., 2002; Burns et al., 2006; Abtahi and Follmer, 2018; Esmaeili et al., 2020] and worthy of further investigation.

Regarding most remaining subjective scores, we found only minor differences between the conditions and did not find meaningful effects of the tested distractions or HR directions. The obtained SUS presence scores verified that the IVE was sufficiently immersive. While most participants stated their willingness to work under redirection for a short amount of time (**10 Min.**), they were significantly less willing to use the experienced redirection for a longer amount of time (**2 Hours**). However, the lack of a control condition without HR and the fact that many participants were rather inexperienced with VR makes it unclear how much this difference is due to a negative impact of the redirection. Future work should compare the results to the general willingness to be in VR. **Body Ownership** ratings were high throughout all conditions, as were **Hand Control** ratings. This indicates that, despite the fixation of the finger, the illusion was appropriate and participants could reasonably identify with the virtual hand.

Post-study results on discomfort were in line with the **Sickness** ratings assessed during the study and verified the absence of sickness issues. Overall, our subjective results emphasize that HR can work well in different distraction scenarios.

8.4.3 Limitations

We conducted the experiment with $N = 12$ participants, which, although at the lower end of the spectrum for psychophysical experiments, is a common participant set size in related research on RDW thresholds (cf. $N = 14$ reported by Steinicke et al. [2010b]). Having each participant contributing 246 samples to the 1AFC analysis in our experiment, and the fact that we did not experience large variations in detection performance between participants, as can be seen from the 95% confidence intervals in Figure 8.5, supports our experimental design.

Concerning the different distraction types, we did not test the derived thresholds for statistically significant differences. One way to do this would be to derive DTs for each individual participant in order to obtain a distribution of upper and lower thresholds and PSEs. By comparing these threshold distributions between distraction scenarios, statistically significant differences could be detected. However, leveraging the method of constant stimuli, this requires significantly more

samples per participant in order to derive individual thresholds, which would have rendered our experiment unfeasible, especially regarding time and user fatigue. A solution to this is to utilize a more efficient psychophysical method for threshold derivation, such as an adaptive staircase design, as we will employ in the evaluation presented in the following chapter.

Finally, our work is an important step towards the derivation of HR guidelines that apply even in worst-case application scenarios. While more complex 2D or 3D warps can be constructed by combining the investigated algorithms, a formal investigation of such combined HR was out of the scope of this chapter, yet represents a promising next step in this research direction.

8.5 Conclusion & Contribution to the Research Questions

The work in this chapter constitutes our first contribution to **RQ 3**. Our goal was to uncover what the human perceptual system is capable of detecting when VR systems employ intentional visual-proprioceptive conflicts to elicit perceptual illusions. Hence, we investigated how much HR can go unnoticed in common redirection scenarios as they would be encountered, for example, in applications that use haptic retargeting or redirected touching.

To achieve this goal, we investigated a popular HR approach, which continuously offsets the virtual from the physical hand of the user – a technique commonly employed to solve the *Colocation* challenge in proxy-based haptics [Azmandian et al., 2016b], and further capable of establishing proxy *Similarity* [Kohli, 2013a]. To systematically investigate how offsets are perceived, we defined three desktop-scale HR algorithms that, respectively, redirect the user’s hand horizontally, vertically, or along the depth direction. Motivated by the lack of (1) conservative lower-bound DTs for HR that correspond to scenarios in which the user is (2) visually perceiving only their displaced hand with a realistic appearance, while (3) reaching for a target in mid-air without the presence of a haptic signal at the hand, we conducted a psychophysical 1AFC experiment. In this experiment, we derived CDTs for horizontal, vertical, and gain-based hand warping leveraging one strictly conservative IVE and two conservative but more realistic scenarios with significantly higher user distraction. These scenarios employed auditory and tactile distraction, and combined visual and dual-task distraction, respectively. By pooling the results of all three scenarios, we derived general recommendations concerning unnoticeable HR.

The results of this chapter add to **RQ 3** by revealing that there is indeed a certain range of manipulation possible – even in worst-case scenarios – in which the VR system can control the user’s hand movement without the user noticing it. Moreover, we add to **RQ 3** by quantifying this range. Our psychophysical results show that for continuous HR, the virtual hand can be displaced *horizontally* or *vertically* by up to $\approx 4.5^\circ$ in either direction respectively, covering a range of $\approx 9^\circ$, without users being able to reliably detect redirection. Regarding the *gain*

technique, we found that factors between $g = 0.88$ and $g = 1.07$ can go unnoticed. Thus, the user's real hand can be redirected to grasp up to 13.75% further or up to 6.18% less far than the virtual hand.

These findings are of immediate value to applications that take advantage of haptic retargeting to establish proxy *Colocation* or employ redirected touching to achieve perceptual *Similarity*. In particular, our findings enable an **Improvement** of proxy-based haptics as the CDTs derived in this chapter inform about the amount of spatial proxy misalignment that can be solved in haptic retargeting scenarios without risking BIPs. Moreover, when warping the hand to simulate different object shapes with redirected touching [Kohli, 2013a], our thresholds can help narrowing down the set of shapes a proxy can represent without resulting in a semantic violation. Finally, our results might even be of value when combined with models that describe the haptic qualities perceived when manipulating the interaction with a proxy through pseudo-haptics. Considering, for example, the weight perception model by Samad et al. [2019], our CDTs might help to design believable haptic illusions that appear natural to the user.

Concerning the findings of this chapter, it is important, though, to keep in mind that the CDTs reported here are likely not ultimate, exact limits. Instead, they represent lower-bounds and might slightly vary as a function of the user and the IVE that is experienced. A more complete picture of the amount of HR possible in different scenarios can be achieved when considering our results in context and in conjunction with the results summarized in Tables 2.4 and 2.5. Our recommendations in this chapter are meant to serve as useful and well-founded reference points that underline the order of magnitude and the limited range of *unnoticeable* HR in worst-case scenarios. As such, they can serve as a baseline for developers to adjust the HR applied in their applications. Taking redirection angles and thresholds into account during the development of IVEs can, for example, inform developers as to which redirections might go unnoticed by users and which scene compositions result in HR beyond thresholds. Moreover, by keeping track of the angles and distances between the user's hand and props in vicinity, VR systems could use the reported CDTs to decide when best to start redirecting the hand. To ensure *unnoticeable* body warping, such systems would trigger HR early enough to stay within thresholds.

Following **RQ 3**, with the derivation of the CDTs, our investigation was also able to show the limitations implied by human perception in the context of HR. Most crucially in this regard, our results show that the minimum range of redirection that will go unnoticed in worst-case scenarios is rather narrow. This conclusion motivates research on novel HR techniques, which take into account perceptual phenomena that bear potential to allow for more HR to go unnoticed. We will follow up on this research path in the following chapter.

Chapter 9

Blink-Suppressed Hand Redirection

In the previous chapter, we studied the detectability of conventional Body Warping-based Hand Redirection (HR), which applies a continuous warping that gradually displaces the virtual from the real hand during interaction. The results of our psychophysical experiment revealed that even in worst-case scenarios, such approaches can redirect the user's hand within a certain range without users noticing the redirection. Yet, our results also revealed that this range can be narrow. Motivated by these findings, in this chapter, we turn towards an **Improvement** of proxy-based haptics by proposing a novel approach to HR. To advance the field and to add to **RQ 3**, we introduce a redirection technique that takes advantage of two perceptual phenomena that impede the perception of visual manipulations: blink suppression and change blindness. We present the concept of our new HR technique, and study its feasibility and detectability in a second psychophysical experiment. Moreover, we compare our proposed approach to the popular state-of-the-art HR algorithm by Cheng et al. [2017b], for which we additionally reveal CDTs.

A video⁶⁸ about the work presented in this chapter is available online. This chapter is based on the following publication. Images and parts of the text in this chapter, as well as the presented concepts, implementations, and results have been published previously therein:

Zenner, A., Regitz, K. P., and Krüger, A. (2021b). Blink-Suppressed Hand Redirection. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'21, pages 75–84. IEEE. © 2021 IEEE. Final published version available in the IEEE Xplore® Digital Library. DOI: [10.1109/VR50410.2021.00028](https://doi.org/10.1109/VR50410.2021.00028)



Video Link⁶⁸

⁶⁸Blink-Suppressed Hand Redirection Video. <https://bit.ly/3IkWLgs>

The publication is based on the following Bachelor's thesis, which was conducted under the supervision of the author of this PhD thesis:

Regitz, K. P. (2020). Leveraging Blink Suppression for Hand Redirection in Virtual Reality. Bachelor's thesis, Saarland University. Advisor: **Zenner, A.** © 2020 Kora Regitz.

9.1 Introduction

With HR being one of the main building blocks for techniques like haptic retargeting and redirected touching, past research has proposed a variety of implementations as outlined in Table 2.3 in chapter 2. Approaches proposed and studied in the past usually offset the virtual hand either by a constant amount (e.g., [Han et al., 2018; Benda et al., 2020]), or continuously increment offsets while the user moves their hand as studied in the previous chapter (e.g., [Kohli, 2013a; Azmandian et al., 2016b; Cheng et al., 2017b]). Moreover, as also discussed in chapter 2, the concept of HR is closely related to the technique of RDW [Razzaque, 2005]. While HR redirects the user's real hand during grasping, RDW redirects the user's real walking path while exploring IVEs through walking.

Past research on RDW examined how perceptual phenomena, such as change blindness, can be utilized to improve redirection. As introduced in subsubsection 2.2.2, change blindness describes "*the inability to detect changes to an object or scene*" [Simons and Levin, 1997], and has been leveraged in a variety of ways to enhance VR experiences. As we outlined in section 2.6, previous work hid changes to a virtual scene from the user's attention by applying manipulations when they were out of the user's sight [Suma et al., 2010, 2011a] or when changes occurred within the user's FOV but outside the visual attention area [Marwecki et al., 2019]. Moreover, and specifically important to this chapter, previous work on RDW also proposed to apply changes during blinks [Langbehn et al., 2018b; Nguyen and Kunz, 2018] or saccades [Sun et al., 2018].

Humans typically blink 10 to 20 times per minute [Doughty, 2002; Leigh and Zee, 2015; p. 241] and during eye blinks, i.e., when the user briefly closes their eyelids for 100ms – 200ms, the visual perception of users is largely suppressed, which leads to users temporarily being (almost) blind [Volkmann, 1986]. Langbehn et al. [2018b] as well as Nguyen and Kunz [2018] studied how this phenomenon of blink suppression and the resulting change blindness can be utilized to enhance RDW by rotating or translating the scene when the user's eyes are closed and found promising results. In the field of HR research, however, the potential of utilizing change blindness has not yet been systematically considered. Yet, from the perspective of VR research, blinks seem to be great opportunities to covertly manipulate the IVE. They are (1) reliably and easily tracked with off-the-shelf eye tracking solutions, (2) periodically occurring, (3) relatively long (compared to the average duration of saccades of only 50ms [Volkmann, 1986]), and (4) triggerable through external stimuli. Hence, blinks lend themselves to being used also for enhancing HR.

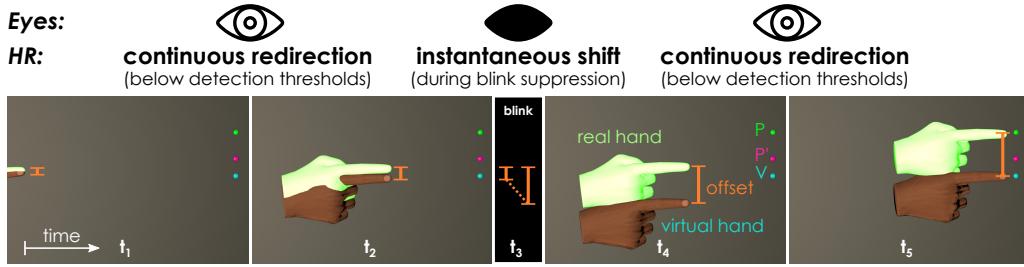


Figure 9.1: Concept of *BSHR*. Our proposed technique applies continuous HR below DTs when the user’s eyes are opened (from t_1 to t_2 , and from t_4 to t_5), and injects additional hand offset during blinks (t_3). © 2021 IEEE.

Motivated by this, in this chapter, we transfer the concept of leveraging change blindness to the domain of HR and propose *Blink-Suppressed Hand Redirection (BSHR)*, the first proof-of-concept body warping technique that makes use of blink suppression. *BSHR* is based on the body warping algorithm by Cheng et al. [2017b] and is designed to showcase the feasibility of leveraging blink-induced change blindness for HR. To verify that unnoticeable *BSHR* is possible with our approach we derive the CDTs of *BSHR* for redirection along three different spatial axes in a psychophysical experiment. Moreover, we contribute to **RQ 3** by deriving estimates for the CDTs of the original HR technique by Cheng et al. [2017b] and compare them to our novel approach that leverages blinks.

9.2 Blink-Suppressed Hand Redirection

The algorithm outlined in the following is the first to leverage blink-induced change blindness for HR. Our aim was to develop a proof-of-concept technique showcasing the feasibility of blink-suppressed HR and allowing us to study its perceptibility. For this, our approach supports two modes of redirection:

Mode 1: pure blink-suppressed redirection

Here, hand offsets are *only* introduced during blinks.

Mode 2: combined redirection

Here, hand offsets are introduced *before, during, and after* a blink.

9.2.1 Concept

We base our algorithm for *Blink-Suppressed Hand Redirection (BSHR)* on the body warping approach of Cheng et al. [2017b]. Figure 9.1 shows an illustration. *BSHR* redirects the user’s hand starting at an origin location O so that the virtual hand arrives at a virtual target V while the physical hand is redirected to reach a physical target P simultaneously. To achieve this, the introduction of hand offsets in the *BSHR* algorithm is governed by two central ideas:

1. While the user's eyes are *open* and the user reaches for the target, *BSHR* *continuously* increases the offset of the virtual to the real hand as in the original algorithm by Cheng et al. [2017b], but *only within ranges that go unnoticed (defined as a parameter)*.
2. When the user *closes* their eyes, i.e., during blink suppression, *BSHR* *instantaneously* changes the hand offset in a way that, after being compensated for, allows reaching the target by continuing to apply only unnoticeable warping as in step 1.

To implement this behavior, *BSHR* redirects the physical hand continuously towards a dummy location P' . This dummy location P' still lies within the range of unnoticeable continuous redirection, which ensures that while the user's eyes are open, only an unnoticeable amount of warping is applied. The additional offset required to reach the target P , but not achievable with unnoticeable continuous warping, is added during visual suppression, i.e., when the user blinks.

The motivation for combining continuous and instantaneous shifts stems from the domain of RDW, as here, researchers found the combination of continuous and discrete redirection techniques (e.g., adding discrete scene rotations during blinks) to successfully reduce the required space for reset-free walking and the number of resets [Nguyen and Kunz, 2018]. Our approach is further motivated by results on hand interactions by Han et al. [2018] and Benda et al. [2020], who evaluated fixed positional hand offsets for HR. While not adding such offsets during blink suppression, their studies revealed DTs of practical relevance [Benda et al., 2020] as summarized in Table 2.5, and found translational shift to outperform their interpolated reach technique [Han et al., 2018].

Reducing the Noticeability of Continuous Warping

Following the concept introduced above, *BSHR* consists of both continuous and instantaneous warping. To minimize the risk of continuous warping being detected while the user's eyes are open, *BSHR* only applies continuous redirection below DTs. As thresholds can vary with reaching distance and other interaction aspects [Abtahi and Follmer, 2018; Burns et al., 2006; Esmaeili et al., 2020; Gonzalez and Follmer, 2019], we do not hard-code thresholds into the *BSHR* algorithm but leave them free as an input parameter. For this chapter, we use the worst-case CDTs estimated in chapter 8 for desktop-scale continuous HR and use these thresholds for the algorithm's internal representation of the unnoticeability range. This *unnoticeability range* encompasses all real positions around V reachable with the real hand without exceeding any of the DTs for either maximum redirection angle (β_{max}), minimum gain (g_{min}), or maximum gain (g_{max}).

Our *BSHR* algorithm realizes both *Mode 1* and *Mode 2* of the *BSHR* concept. To apply pure blink-suppressed redirection (*Mode 1*), it suffices to set the threshold parameters to $\beta_{max} = 0$ and $g_{min} = g_{max} = 1$, as this effectively reduces the size of the unnoticeability range to 0, which, in turn results in *BSHR* not applying any

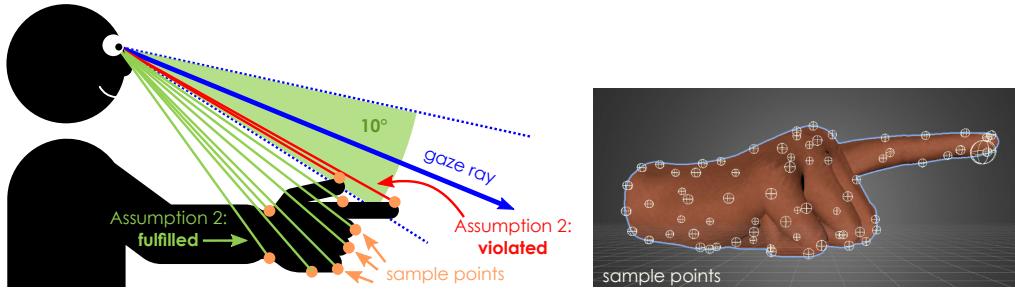


Figure 9.2: Check for Assumption 2. Left: It is checked for each sample point on the virtual hand model and each sample translated by \vec{b} , if its position is within the 10° focus area around the gaze ray. Right: Distribution of samples. Both images © 2021 IEEE.

continuous warping. In contrast, setting $\beta_{max} > 0$, $g_{min} < 1$, and $g_{max} > 1$ results in an unnoticeability range with a size greater than 0 as illustrated in Figure 9.3, and leads to the *BSHR* algorithm applying combined redirection (*Mode 2*).

Reducing the Noticeability of Instantaneous Warping

Inspired by Marwecki et al. [2019] and illustrated in Figure 9.2, we added a mechanism to reduce the noticeability of instantaneous hand shifts during blinks. This mechanism prevents the injection of hand offsets if the virtual hand rendering is likely to jump from outside to inside the user's visual focus area, or vice versa. For this, a subroutine checks where the virtual hand is rendered when the user closes their eyes, and then approximates where it would be rendered if the offset was changed during the blink. The hand offset is then only changed if neither of these two rendering locations intersect the user's visual focus.

9.2.2 Assumptions

The introduced measures and our general concept lead to three assumptions our proof-of-concept *BSHR* algorithm is based on:

Assumption 1: the user blinks ≥ 1 times while reaching for the target.

Assumption 2: the virtual hand does not intersect the user's visual focus area before and after the hand offset is changed during a blink (approximated here as everything $\leq 10^\circ$ from the gaze ray, representing the 5° foveal vision area [Marwecki et al., 2019] and additional 5° eye tracking tolerance).

Assumption 3: continuous angular warping up to β_{max} , and gain-based warping g with $g_{min} \leq g \leq g_{max}$, are below DTs and likely to go unnoticed⁶⁹.

⁶⁹We note that chapter 8 only considered the detectability of horizontal, vertical, and depth-based warping, but not their combination. For simplicity and to support arbitrary directions, *BSHR* assumes offsets also go unnoticed within the naïve combination of the DTs found in chapter 8.

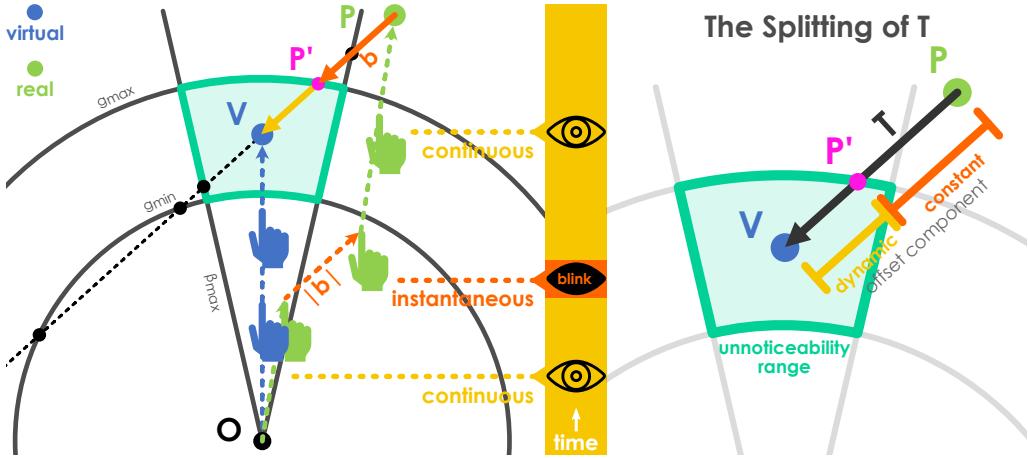


Figure 9.3: The orthogonal top-down view on the left shows V , P , P' , the unnoticeability range, as well as vectors and intersection points relevant to the computation of P' . We also sketch the trajectory of the real hand assuming perfect compensation for hand offsets. The right close-up illustrates how the total offset vector \vec{T} is split into a dynamic and a constant component. © 2021 IEEE.

Specifically, we assume: $\beta_{max} = 4.5^\circ$; $g_{min} = \frac{real}{virtual} = 0.94$; $g_{max} = 1.14$ based on the results of chapter 8.

For our evaluation of *BSHR* in this chapter, we ensured that assumptions 1 and 2 were met by only considering trials that fulfilled both conditions. Assumption 3 is indirectly built into the algorithm's computation of P' .

9.2.3 Algorithm

Algorithm 3 shows the pseudocode of *BSHR*. When starting a new redirection, the algorithm is initialized with the `INIT` procedure. The `UPDATE` function is called every frame to re-compute the virtual hand position H_v .

The *BSHR* algorithm offsets the virtual hand H_v from the physical hand H_p by applying two different translation components:

1. A *dynamic* offset component, i.e., a gradually growing translational offset, implemented through the continuously updated warp vector \vec{W} that represents continuous warping as proposed by Cheng et al. [2017b].
2. A *constant* offset component, i.e., a fixed translational offset that is toggled (switched on) during the first valid blink, implemented through the blink vector \vec{b} that represents instantaneous warping, similar to the approach by Benda et al. [2020].

Conceptually, each component is responsible for handling a different segment of the total offset vector $\vec{T} = V - P$ that is to be applied at the end of the redirection, as illustrated on the right in Figure 9.3.

Algorithm 3 Blink-Suppressed Hand Redirection

Input: Locations: warp origin O , real target P , virtual target V ;
Unnoticeability Range: DTs $0 \leq \beta_{max} \leq 180^\circ$, $0 < g_{min} \leq 1$, $1 \leq g_{max}$;
Frame-Wise: real hand position H_p , eye tracking data $eyes$.

Output: virtual hand position H_v

```

1: procedure INIT( $O, P, V, \beta_{max}, g_{min}, g_{max}$ ) ▷ executed once
2:    $P' = \text{ComputeDummyTarget}(O, P, V, \beta_{max}, g_{min}, g_{max})$  ▷ see Equation 9.8
3:    $\vec{b} = \vec{0}$  ▷ no constant warp before blink
4: end procedure
5:
6: procedure UPDATE( $H_p, eyes$ ) ▷ executed every frame
7:   if  $eyes_{closed}$  && Assumption2( $H_v, \vec{b}, eyes_{gaze}$ ) then ▷ check for valid blink
8:      $\vec{b} = P' - P$  ▷ apply constant warp
9:   end if
10:   $\alpha = \frac{|(H_p + \vec{b}) - O|}{|(H_p + \vec{b}) - O| + |(H_p + \vec{b}) - P'|}$  ▷ update dynamic warp ratio
11:   $\vec{W} = \alpha \cdot (V - P')$  ▷ update dynamic warp vector
12:   $H_v = (H_p + \vec{b}) + \vec{W}$  ▷ update virtual hand position
13: end procedure

```

The translation added through the dynamic component, i.e., through the frame-wise updated warp vector \vec{W} , is responsible for gradually introducing the offset that can be covered by unnoticeable continuous warping, i.e., the vector from the dummy target P' to V . To ensure that only unnoticeable warping is applied, the INIT procedure places P' at the closest location to P on \vec{T} that is still within the unnoticeability range. The remaining redirection offset, i.e., the vector from P to P' , is handled by the constant component, i.e., by an *instantaneous* translational hand offset applied during the first valid blink. For this, a blink vector \vec{b} is set to $\vec{0}$ (meaning no constant offset applied) before the first blink, and is changed to $\vec{b} = P' - P$ when the user validly blinks for the first time during the redirection (resulting in the constant offset component being applied). Following this scheme, the virtual hand position H_v is computed as

$$H_v = (H_p + \vec{b}) + \vec{W} \quad (9.1)$$

For realizing the *dynamic* component, the computation of \vec{W} is based on the approach of Cheng et al. [2017b], recomputing \vec{W} every frame as:

$$\vec{W} = \alpha \cdot (V - P') \quad (9.2)$$

where

$$\alpha = \frac{|(H_p + \vec{b}) - O|}{|(H_p + \vec{b}) - O| + |(H_p + \vec{b}) - P'|} \quad (9.3)$$

and

$$\vec{b} = \begin{cases} \vec{0}, & \text{before the 1}^{\text{st}} \text{ valid blink} \\ P' - P, & \text{else} \end{cases} \quad (9.4)$$

Comparing this algorithm to the original algorithm by Cheng et al. [2017b] makes apparent that, while the original algorithm computes hand offsets with respect to the location of the physical hand H_p , the *BSHR* algorithm computes hand offsets with respect to the (dynamic) reference location $(H_p + \vec{b})$. This reference position equals H_p before the user has validly blinked as, in this case, \vec{b} still equals $\vec{0}$. Once the user blinks validly for the first time during the reach, however, the blink vector \vec{b} is set to $\vec{b} = P' - P$, and is applied to the hand rendering in Equation 9.1. As a consequence the user will compensate for the now-applied constant translational offset \vec{b} by shifting their physical hand by $-\vec{b}$.

In order to account for this constant translation component introduced during the blink, the reference location of the dynamic component likewise shifts away from H_p by \vec{b} , effectively ignoring compensations that result from the fixed translational offset \vec{b} when computing \vec{W} . As a result, in order to reach V with the virtual hand, the user needs to compensate for both the constant translational offset $\vec{b} = P' - P$ applied during the blink, and for the full offset $V - P'$ added by the dynamic warping component. Such compensation results in the physical hand reaching the physical target P , as can be seen by solving for H_p while substituting \vec{b} and \vec{W} with the respective offsets, applied when H_v reaches V :

$$\begin{aligned}
 H_v &= (H_p + \vec{b}) + \vec{W} \\
 \Leftrightarrow H_p &= H_v - \vec{b} - \vec{W} \\
 &= V - (P' - P) - (V - P') \\
 &= V - P' + P - V + P' \\
 &= (V - V) + (P' - P') + P \\
 &= P
 \end{aligned} \tag{9.5}$$

■

Determining a Valid Blink

The occurrence of the first valid blink is determined using eye tracking data, specifically, by querying (1) if the eyes are closed, and (2) the eye gaze ray. To check for assumption 2, the 3D model of the virtual hand was populated with 68 invisible reference positions distributed over the surface of the hand as shown in Figure 9.2. To determine if the model was visible inside the user's visual focus, the angles of the gaze vector and the vector from gaze ray origin towards each reference position were checked against a threshold of 10° according to assumption 2. The same checks were also conducted for each reference position translated by $\vec{b} = P' - P$ to approximate the hand location after the blink.

Computation of P'

The computation of P' constitutes a central part of the *BSHR* algorithm as it splits the total offset vector \vec{T} into a part added through gradual warping with the

dynamic offset component, and a remaining part added as a *constant* offset during the blink. As per the concept of *BSHR*, P' is to be chosen so that the offset applied through gradual warping goes unnoticed even when the user's eyes are opened.

To achieve this, the dummy location P' is defined as the closest point to the physical target P that lies on \vec{T} , i.e., on the direct connection of P and V , but still is inside the unnoticeability range. As can be seen from Figure 9.3, potential optimal P' locations are:

1. P – the physical target itself, if within the unnoticeability range
2. $P'_{\beta_{max},\{right,left\}}$ – points where β_{max} is exceeded
3. $P'_{g_{max},\{1,2\}}$ – points where g_{max} is exceeded
4. $P'_{g_{min},\{1,2\}}$ – points where g_{min} is exceeded

To determine P' , the *BSHR* algorithm computes all of these candidate points using 3D line intersection methods. The computation of $P'_{g_{max},\{1,2\}}$, for example, is implemented by computing a 3D line-sphere intersection between the line r (covering the ray from V towards P) with:

$$r : V + t \cdot (P - V) \quad t \in \mathbb{R} \quad (9.6)$$

and a sphere at the origin O with radius $g_{max} \cdot |V - O|$, the surface of which geometrically represents the g_{max} threshold. The intersection of r with this sphere thus represents the position along r at which the g_{max} threshold is exceeded. In the same way, the points along r at which the g_{min} threshold is exceeded can be computed by intersecting r against a sphere at O with radius $g_{min} \cdot |V - O|$. As the ray origin V does not lie inside this sphere, here, either two, one, or no intersections exist. Finally, the angular threshold β_{max} is geometrically represented by two lines, which can be defined by rotating the vector from O to V towards the $\{right, left\}$ by angle β_{max} . To find $P'_{\beta_{max},\{right,left\}}$, the intersections of r with these lines are also computed. For computational robustness, we represented the β_{max} lines as two 3D planes and compute line-plane intersections.

Each of the computed intersections is represented by a value $t \in \mathbb{R}$ (see Equation 9.6) denoting the relative position of the intersection along the line ($t = 0$ representing V ; $t = 1$ representing P). To determine P' as the location that is closest to the target P but not exceeding any thresholds, the *BSHR* algorithm sorts all corresponding $t \geq 0$ in ascending order, resulting in a list where t_0 is the minimum, and t_6 the maximum value⁷⁰:

$$sort(\{t_P, t_{\beta_{max},right}, \dots, t_{g_{min},2}\}) = [t_0, t_1, \dots, t_6] \quad (9.7)$$

Accounting for a special case in which the g_{min} sphere is intersected twice, i.e., the threshold is first exceeded and then met again before any other threshold is

⁷⁰Note that less than seven points can be in this list, as for some geometric constellations, intersections might not be defined or yield $t < 0$.

exceeded, the algorithm finally returns:

$$P' = \begin{cases} V + t_2 \cdot (P - V), & \text{if } t_0 \text{ and } t_1 \text{ belong to } g_{\min} \\ V + t_0 \cdot (P - V), & \text{else} \end{cases} \quad (9.8)$$

9.3 Evaluation of *Blink-Suppressed Hand Redirection*

We conducted a psychophysical user experiment to study the detectability of *BSHR*. Our goal was to validate the feasibility of leveraging blink suppression for unnoticeable HR. To this end, we estimated the CDTs of four HR techniques:

1. *BSHR_{+0%}*: pure *BSHR* (*Mode 1*); no continuous warping.
2. *BSHR_{+50%}*: combined *BSHR* (*Mode 2*); incl. continuous warping up to 50% of DTs (cf. **H2a**, **H2b** below).
3. *BSHR_{+100%}*: combined *BSHR* (*Mode 2*); incl. continuous warping up to 100% of DTs (cf. **H2a**, **H2b** below).
4. *Cheng*: conventional HR as proposed by Cheng et al. [2017b]; unlimited continuous warping without instantaneous warping during blinks.

While the *BSHR_{+0%}* condition offsets the hand only instantaneously during blink suppression, *BSHR_{+50%}* and *BSHR_{+100%}* additionally redirect the hand when the eyes are open by applying continuous warps up to 50% and 100% of the DTs outlined in assumption 3, respectively. *Cheng* does not utilize any instantaneous shifts, but only applies unlimited continuous redirection [Cheng et al., 2017b].

In line with chapter 8 and previous research [Benda et al., 2020], we restrict our investigation to horizontal, vertical, and depth-related redirection. To prevent fatigue and to maintain an acceptable duration of the experiment, we estimated the CDTs of each participant and each of the four techniques for three directions:

- *right*: the virtual hand is offset towards the right (+X).
- *down*: the virtual hand is offset downwards (-Y).
- *towards*: the virtual hand is offset towards the user (-Z).

9.3.1 Hypotheses

Our study is designed to evaluate four central hypotheses. First, we expected unnoticeable pure *BSHR* (*Mode 1*, *BSHR_{+0%}*) to be feasible in practice, i.e., to achieve an unnoticeability range significantly greater than 0cm (**H1**).

$$\mathbf{H1: CDT}_{\text{all dir.}}(\text{BSHR}_{+0\%}) > 0\text{cm}$$

Moreover, as previous research did not yet reveal the DTs for the continuous warping technique by Cheng et al. [2017b], secondly, we expected *Cheng* to go unnoticed within the lower-bound thresholds found in chapter 8 (**H2**).

H2a: $CDT_{\{right, down\}}(Cheng) \geq \beta_{max} = 4.5^\circ$ (here: translation of 3.15cm)

H2b: $CDT_{\{towards\}}(Cheng) \geq g_{max} = 1.14$ (here: translation of 5.5cm)

Regarding the combined *BSHR* approach (*Mode 2*, $BSHR_{+50\%}$ and $BSHR_{+100\%}$), we expected the range of unnoticeable *BSHR* to increase as the amount of continuous warping increases (**H3**).

H3a: $CDT_{\text{all dir.}}(BSHR_{+0\%}) < CDT_{\text{all dir.}}(BSHR_{+50\%})$

H3b: $CDT_{\text{all dir.}}(BSHR_{+50\%}) < CDT_{\text{all dir.}}(BSHR_{+100\%})$

Finally, we expected the CDTs of $BSHR_{+100\%}$ to exceed those of *Cheng* (**H4**).

H4: $CDT_{\text{all dir.}}(Cheng) < CDT_{\text{all dir.}}(BSHR_{+100\%})$

9.3.2 Participants

17 volunteer participants were recruited from the local campus. The results of two participants had to be excluded from analysis as the staircase procedure for one of them did not converge, while the experiment took too long for the other participant. Hence, $N = 15$ participants (7 female, 8 male) completed the experiment. Participants were on average 25.5 years old ($SD = 3.5$ years, min. 20 years, max. 33 years). Assessed on a scale from 1 (= never) to 7 (= regularly), our participants covered a wide range of previous experience levels with 3D video games ($M = 3.8$, $SD = 2.65$, min. 1, max. 7), VR ($M = 2.6$, $SD = 1.73$, min. 1, max. 7) and manual crafting ($M = 4.4$, $SD = 1.41$, min. 2, max. 7).

9.3.3 Apparatus

Our study took place in a lab at our institution. The setup can be seen in Figure 9.4. We used a HTC Vive Pro Eye²⁰ HMD, a HTC Vive Pro Controller⁵³, and a HTC Vive Tracker⁵⁴ (v2018) tracked with SteamVR base stations 2.0. A notebook with an NVIDIA GTX 1070 graphics card was used to run the study software, which was implemented with the Unity 2019.3.0f6 engine¹⁷, the *VRQuestionnaireToolkit*⁴⁷ by Feick et al. [2020b], the *Unity Experiment Framework*⁴⁶ by Brookes et al. [2020], and the *SRanipal SDK*⁷¹ for eye tracking.

In addition, we developed the *Unity Staircase Procedure Toolkit*. The toolkit was used to implement the psychophysical staircase method used in this experiment (as well as the experiment in chapter 12) and is available online⁷² as open-source.

To track the dominant hand of participants, we applied the same tracking solution as in chapter 8 utilizing a tracker attached with a rubber band to the back of the hand. To maintain calibration and a static hand pose, participants wore a finger splint as shown in Figure 9.4 (b), which fixed the offset of the tracker on



GitHub Link⁷²

⁷¹*SRanipal SDK* Homepage. <https://bit.ly/3e4v1lF>

⁷²*Unity Staircase Procedure Toolkit* on GitHub. <https://github.com/AndreZenner/staircase-procedure>

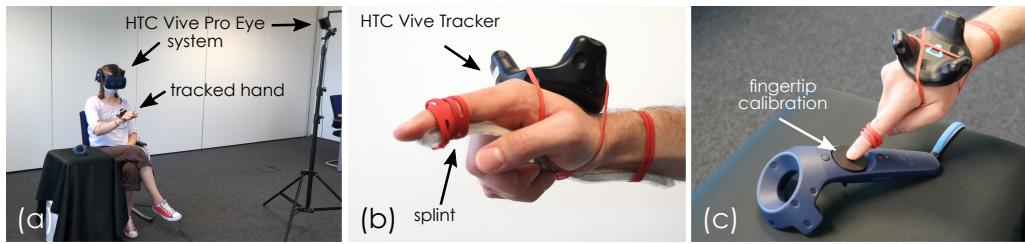


Figure 9.4: Experiment setup: (a) the real environment with the tracking system and calibration pedestal, (b) the hand tracking rig, and (c) the calibration of the offset from fingertip to tracker (as in chapter 8). All images © 2021 IEEE.

the back of the hand to the fingertip. This offset was (re-)calibrated before each experimental block. For this and as in chapter 8, participants repeatedly touched the touch-sensitive trackpad on the controller as shown in Figure 9.4 (c).

9.3.4 Procedure

Upon providing informed consent and demographic data, participants put on the VR equipment and sat on a chair inside the tracking volume as shown in Figure 9.4 (a). For hand calibration, the HTC Vive Controller was placed on a pedestal next to the participant and the offset from the fingertip to the HTC Vive Tracker was calibrated as shown in Figure 9.4 (c) and Figure 9.5 (a).

Participants started to practice the experimental trials in two test blocks ($BSHR_{+0\%}$ and *Cheng*; random directions), before the actual experiment and data recording began. To complete the experiment, each participant had to complete 12 blocks (a threshold estimation procedure for each of the four techniques in each of the three directions). To complete a block, a simple trial was repeated several times. In each trial, users were asked to reach towards a small virtual sphere initially rendered 125cm in front of them. When passing with their finger through a circular zone with radius 5cm centered 30cm below and 25cm straight in front of their head (highlighted as a red circle in Figure 9.5 (b)), HR and blink detection was activated and the virtual sphere relocated to the virtual target position V of the trial. V was located 40cm straight in front of their fingertip location when passing through the circular zone (in line with the experiment in chapter 8). While reaching for the target as shown in Figure 9.5 (c), the HR algorithm that corresponds to the current condition was applied, with the real target position P being offset from V along the axis corresponding to the block.

The amount of offset between real and virtual target (between P and V ; i.e., the stimulus) varied across trials following an interleaved staircase procedure, as described in the following section. To ensure conservativeness, participants were instructed to pay close attention to detecting any signs of redirection, and were told that both continuous and sudden hand offsets could occur at any time. Moreover, participants were asked to blink frequently throughout the experiment

(across all techniques) and informed that only trials in which they blinked at least once during the reach could be used for analysis. However, participants were intentionally not told about the fact that if a trial was determined to be invalid, the trial was completed as usual, with the offset being repeated immediately in the following trial.

To ensure that each trial used in the analysis met assumptions 1 and 2, a trial was considered invalid if the participant did not blink at all during the reach, or blinked but assumption 2 was not met. These requirements applied to all conditions to ensure comparability. To end a trial when the virtual finger reached V , participants responded to the 1AFC yes/no question “Did you notice a manipulation?” by touching one of the two virtual answering cubes shown in Figure 9.5 (d). After answering, the virtual target sphere relocated to its initial position and participants continued with the next trial in the block until the respective staircase procedure, and with it the block, terminated. In between blocks, participants were instructed to rest for at least 30s and to re-calibrate their fingertip tracking. After completion of all 12 blocks, the SUS presence questionnaire [Slater et al., 1994], the SSQ [Kennedy et al., 1993], and a concluding questionnaire were filled out in VR as can be seen in Figure 9.5 (e).

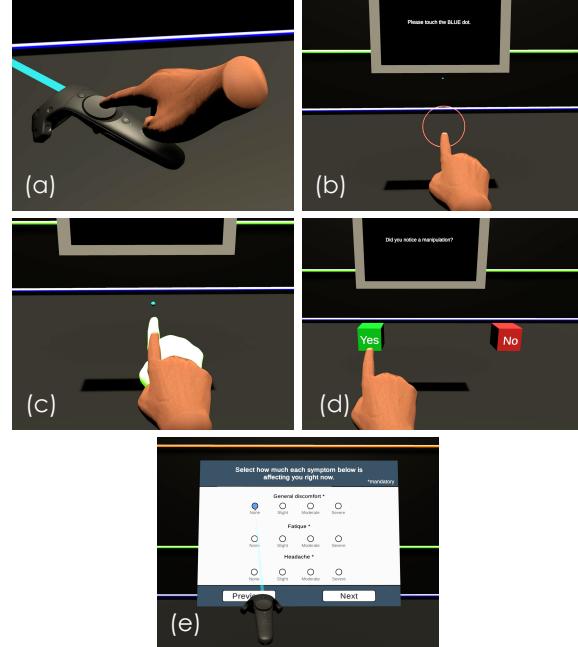


Figure 9.5: (a) Fingertip calibration, (b) circular start area, (c) virtual hand being offset *towards* the user (white real hand added for illustration only), (d) the 1AFC yes/no question, and (e) the SSQ. All images © 2021 IEEE.

9.3.5 Design

Our study has a within-subject design with two independent variables: redirection algorithm (i.e., the four HR techniques) and direction (i.e., the three offset directions). As a result, each participant completed 12 blocks and we used a 12×12 Latin square [Williams, 1949] to counterbalance blocks across participants. For each block, we employed a psychophysical DT estimation method [Kingdom and Prins, 2016b; pp. 29 ff.], specifically, a 1 up/1 down adaptive staircase method [Kingdom and Prins, 2016a; pp. 120 ff.], in order to approximate for

		H1		H3		H4		H2				
		Base	p'	$BSHR_{+0\%}$	p'	$BSHR_{+50\%}$	p'	$BSHR_{+100\%}$	p'	$Cheng$	p'	$Cheng_{exp}$
<i>right</i>	0cm	< .001		M = 2.65cm (SD = 1.26cm)	$\leq .038$	3.56cm (1.27cm)	$\leq .038$	4.34cm (1.47cm)	$\leq .027$	5.81cm (1.98cm)	$\leq .006$	3.148cm
<i>down</i>	0cm	< .001		M = 3.83cm (SD = 1.11cm)	$\leq .045$	4.94cm (1.47cm)	$= .372$	5.39cm (1.23cm)	$= .64$	5.63cm (2.13cm)	$\leq .006$	3.148cm
<i>towards</i>	0cm	< .001		M = 3.27cm (SD = 1.66cm)	$= .076$	4.26cm (1.73cm)	$= .076$	4.36cm (1.23cm)	$= .076$	4.63cm (1.88cm)	$= .519$	5.5cm

Table 9.1: CDTs for the four HR techniques along the three directions (p' indicates corrected p -value; sig. differences highlighted green). The top row indicates to which hypotheses the comparisons belong. A growth of the unnoticeability range can be observed with increasing continuous warping in $BSHR$. ©2021 IEEE.

each algorithm and direction the participant's CDT. The CDT, being the dependent variable in our study, represents the amount of redirection in *cm* that goes unnoticed when applying the respective algorithm in the respective direction.

We used an interleaved staircase implementation with an ascending (starting at the minimum redirection of *0cm*) and a descending (starting at the maximum redirection of *8cm*; determined during piloting) sequence, using a constant step size of *0.8cm*. If participants noticed the redirection in a trial (answering *yes*), the amount of redirection in the following trial of that sequence was decreased by the step size; otherwise (answering *no*) it was increased by the step size. A sequence terminated after five reversals, with the average of the last four reversals being taken as the sequence threshold estimate. The average of the ascending and descending sequence thresholds yielded the CDT.

9.3.6 Results

The SUS count ($M = 1.87$, $SD = 1.60$, min. 0, max. 5) and SUS mean ($M = 4.61$, $SD = 0.97$, min. 2.33, max. 6.33) scores verified our IVE to be generally immersive, while a relatively low SSQ total score ($M = 38.15$, $SD = 32.33$) did not indicate any cybersickness issues. In total, 6739 trials ($M = 449.3$, $SD = 121.6$ per participant) were completed, out of which 4581 trials ($M = 305.4$, $SD = 19.6$ p. p.) were valid and contributed to our analysis. The valid blinks of our participants lasted *115.9ms* on average ($SD = 34.5ms$). When blinking validly, participants closed their eyes when their hand had traveled 46.19% on average ($SD = 13.83\%$) along the way towards the target, and opened them again at 52.61% on average ($SD = 13.74\%$). To study our hypotheses, we analyzed the CDTs obtained for the 12 conditions, applying a significance level of $\alpha = .05$.

H1: Detectability of $BSHR_{+0\%}$ (Mode 1)

To investigate **H1**, we compared the obtained CDTs of $BSHR_{+0\%}$ for each direction against the baseline of *0cm*. After normality of the respective data was confirmed by a Shapiro-Wilk test, we performed three one-sample t-tests with Bonferroni

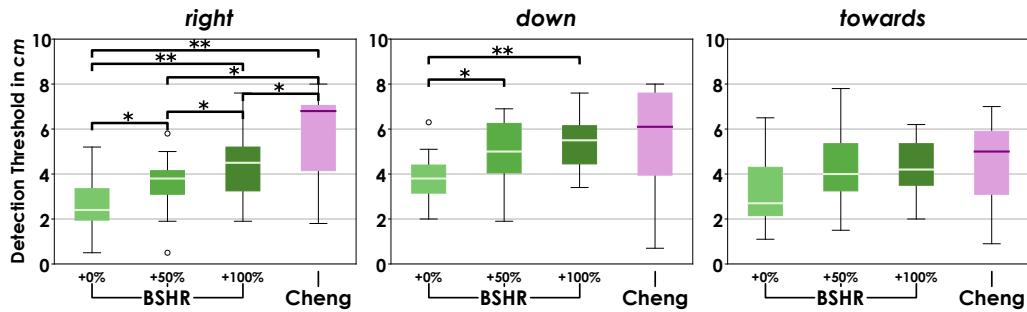


Figure 9.6: CDTs of the four HR techniques for each of the three tested directions. Brackets indicate sig. differences ($p' < .05$ (*); $p' < .01$ (**)). All charts © 2021 IEEE.

correction (corrected p values denoted as p'). The ranges of unnoticeable pure *BSHR* (i.e., *Mode 1*) implemented by $BSHR_{+0\%}$ with virtual hand offsets in direction *right* ($M = 2.65\text{cm}$, $SD = 1.26\text{cm}$), *down* ($M = 3.83\text{cm}$, $SD = 1.11\text{cm}$), and *towards* the user ($M = 3.27\text{cm}$, $SD = 1.66\text{cm}$) were all found to be significantly greater than 0cm (all $p' < .001$).

H2: Detectability of *Cheng*

To confirm **H2**, i.e., our assumption that continuous redirection with the *Cheng* algorithm [Cheng et al., 2017b] goes unnoticed within the worst-case DTs found in chapter 8, we compared the CDTs obtained for *Cheng* against these expected values (labeled *Cheng_{exp}* in Table 9.1). Since a Shapiro-Wilk test indicated a violation of the normality assumption ($p < .05$), we applied three one-sample Wilcoxon tests with Bonferroni correction. The results indicate the CDTs of *Cheng* to even significantly exceed the expected unnoticeability range of 3.15cm (i.e., 4.5° as found in chapter 8 in a distance of 40cm) for *right* ($M = 5.81\text{cm}$, $SD = 1.98\text{cm}$) and *down* ($M = 5.63\text{cm}$, $SD = 2.13\text{cm}$) (both $p' \leq .006$). The thresholds *towards* the user ($M = 4.63\text{cm}$, $SD = 1.88\text{cm}$) were not found to differ significantly from the expected threshold of 5.5cm (i.e., the real hand grasping 13.75% further than the virtual hand as found in chapter 8) ($p' = .519$).

H3 & H4: Detectability of Combined Redirection (*Mode 2*)

To analyze **H3** and **H4**, we compared the CDTs of the four HR techniques for each direction. Since normality could not be assumed in all cases according to Shapiro-Wilk tests, we performed a non-parametric Friedman test for each direction. To find pairwise differences, we applied post-hoc Wilcoxon signed-rank tests with Bonferroni-Holm correction. Figure 9.6 shows the obtained thresholds and indicates significant pairwise differences. Figure 9.7 shows a representative staircase plot and Table 9.1 summarizes the results.

For direction *right* ($\chi^2(3) = 28.034, p < .001$) and *down* ($\chi^2(3) = 14.331, p \leq .002$), the Friedman tests indicated that thresholds differed significantly across the HR

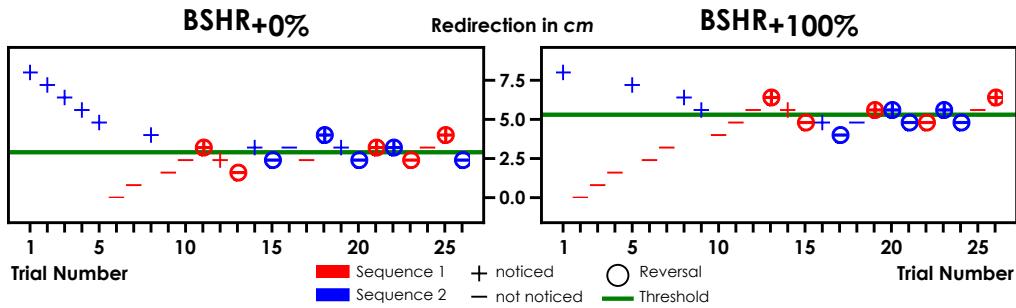


Figure 9.7: Staircase of participant #1 for HR to the *right*. An increase in DT is observed when continuous warping is added in $BSHR_{+100\%}$ compared to $BSHR_{+0\%}$. ©2021 IEEE.

techniques. Post-hoc results for *right* showed each individual technique to yield an unnoticeability range that is significantly different from any other technique's range (all $p' \leq .038$). For *down*, the pairwise test only showed the thresholds of $BSHR_{+0\%}$ to differ significantly from those of $BSHR_{+50\%}$ and $BSHR_{+100\%}$ (both $p' \leq .045$). For *towards*, the Friedman test did not indicate thresholds to differ significantly ($\chi^2(3) = 6.872, p = .076$).

Most importantly regarding **H3**, the hypothesized growth in CDTs from $BSHR_{+0\%}$ to $BSHR_{+50\%}$ was statistically confirmed for HR to the *right* (from $M = 2.65\text{cm}$ ($SD = 1.26\text{cm}$) to $M = 3.56\text{cm}$ ($SD = 1.27\text{cm}$); $Z = -2.355, p' \leq .038, r = .43$) and *down* (from $M = 3.83\text{cm}$ ($SD = 1.11\text{cm}$) to $M = 4.94\text{cm}$ ($SD = 1.47\text{cm}$); $Z = -2.613, p' \leq .045, r = .48$). The hypothesized increase of thresholds from $BSHR_{+50\%}$ to $BSHR_{+100\%}$ was significant only for *right* (from $M = 3.56\text{cm}$ ($SD = 1.27\text{cm}$) to $M = 4.34\text{cm}$ ($SD = 1.47\text{cm}$); $Z = -2.133, p' \leq .038, r = .39$).

Analyzing **H4**, the *right* threshold of $BSHR_{+100\%}$ ($M = 4.34\text{cm}, SD = 1.47\text{cm}$) was, in contrast to our expectations, significantly smaller than that of *Cheng* ($M = 5.81\text{cm}, SD = 1.98\text{cm}$) ($Z = -2.616, p' \leq .027, r = .48$). For *down* and *towards*, differences between $BSHR_{+100\%}$ and *Cheng* were non-significant.

9.4 Discussion of Blink-Suppressed Hand Redirection

Our results validate the feasibility of blink-suppressed HR and point to the value of combining it with continuous hand warping.

9.4.1 The Detectability of $BSHR_{+0\%}$ (Mode 1) and *Cheng*

Studying $BSHR_{+0\%}$, i.e., the sole application of instantaneous hand shifts during blinks, revealed CDTs significantly greater than 0cm for each tested direction as summarized in column $BSHR_{+0\%}$ in Table 9.1. This result confirms **H1** and proves the practical feasibility of pure $BSHR$ (Mode 1). The unnoticeability ranges found for $BSHR_{+0\%}$ correspond to 3.79° and 5.47° for *right* and *down* respectively, and to

a gain factor of 1.08 for *towards*, and are thus of the same order of magnitude as the CDTs found for continuous hand warping in chapter 8. Our results demonstrate for the first time that these ranges of unnoticeable HR can also be achieved by only leveraging periods of blink suppression, instead of applying manipulations while users observe the scene with opened eyes.

The opposite extreme, i.e., pure continuous hand warping with gradually increasing offsets as realized by the *Cheng* algorithm [Cheng et al., 2017b], likewise revealed thresholds of the same order of magnitude as the thresholds found in chapter 8 – further validating our previous results. Corresponding CDTs are summarized in column *Cheng* in Table 9.1. We found the CDTs of the *Cheng* algorithm along the horizontal and vertical axis to even exceed the lower-bound thresholds derived in chapter 8 to a statistically significant extent. Concerning redirection along the depth-axis, no significant deviation from the expected threshold (based on chapter 8) was found for the *Cheng* algorithm, with the CDT being only non-significantly below the expected 5.5cm. Our results thus also support **H2** and validate assumption 3: our expectation that within the worst-case thresholds found in chapter 8, the continuous warping of state-of-the-art algorithms like the *Cheng* algorithm also goes unnoticed. Consequently, the CDTs derived in the *Cheng* condition back our concept of combining continuous warping below the DTs derived in chapter 8 with instantaneous shifts during blink suppression.

9.4.2 The Detectability of Combined Redirection (Mode 2)

With **H3**, we investigated if combining (1) instantaneous hand shifts during blinks, with (2) continuous warping during phases where the user watches the scene with opened eyes, has any effect on the detectability of the redirection. As can be seen from Table 9.1, we found that the average range of unnoticeable redirection increased with increasing amounts of continuous warping added to the *BSHR* approach for every tested direction. Our statistical evidence, however, only partially supports **H3** as this rise in CDTs was found to be statistically significant only for *right* and partially for *down*. From our observations, users seemed to detect manipulations along the depth axis more easily. This seems backed by Benda et al. [2020] and might be related to how sensory integration is affected by direction [van Beers et al., 2002]. We speculate that this special role of depth-related perception might have contributed to the fact that an increase of CDTs did not become statistically striking for *towards*. Nonetheless, our data indicates a general tendency of thresholds to increase when combining both approaches (*Mode 2*), compared to pure *BSHR* (*Mode 1*).

The central idea of the proposed approach, i.e., allowing for as much continuous warping when the user looks at the scene as goes unnoticed, was implemented based on the thresholds derived previously in chapter 8, but can now be refined using the results of our study. Having derived CDT estimates that correspond specifically to the algorithm by Cheng et al. [2017b], the *BSHR* parameters can now be fine-tuned. Specifically, we recommend replacing the unnoticeability

range defined in assumption 3 with the now-known thresholds of *Cheng* summarized in Table 9.1. In line with this, we speculate that DTs can be further increased, maybe even up to a point where the combined technique allows for more redirection to go unnoticed than the continuous approach. This might occur if the entire unnoticeability range of the continuous approach is leveraged while users look at the scene, and an additional unnoticeable offset is introduced when the vision of users is suppressed. While we hypothesized this effect to occur already in this study, **H4** was not supported by our results as CDTs of $BSHR_{+100\%}$ did not exceed those of *Cheng* along any axis in our experiment. We suspect this to be a consequence of the non-optimal selection of the parameters β_{max} , g_{min} , and g_{max} , which were not directly based on the applied *Cheng* technique, but on the lower-bound estimates derived with a slightly different implementation of continuous HR in chapter 8. For example, the threshold of *Cheng* for offsets *towards* the user was lower than the configured range in $BSHR_{+100\%}$ of 5.5cm. Consequently, participants might have noticed the incremental warping in $BSHR_{+100\%}$ even before any offset was added during a blink. In the case of *right* and *down*, on the other hand, more continuous warping would have been possible within the unnoticeability ranges of *Cheng* than configured based on assumption 3. Hence, blink warps were applied “too early”, i.e., already for redirections smaller than necessary, and participants might have noticed these offsets added during blinks even for redirections below the *Cheng* thresholds. We are optimistic, however, that with more participant-specific parameters, combined redirection might achieve larger unnoticeability ranges than both individual approaches – a hypothesis to be investigated in future research. For this, a per-user calibration might help to determine the most optimal β_{max} , g_{min} , and g_{max} parameters for each individual.

9.4.3 Limitations

The *BSHR* algorithm proposed in this chapter is designed to allow for a first controlled investigation of the feasibility and the detectability of blink-suppressed HR. To reliably redirect the user’s real hand to reach P , however, assumptions 1 and 2 have to be fulfilled. These assumptions are to ensure consistency and comparability of our results, but limit the practical usability of *BSHR* as proposed in this chapter, as these assumptions might not always be met in in-the-wild scenarios. We thus aim to evolve the *BSHR* approach in future work, e.g., by integrating a fail-safe mechanism handling such situations to ensure that the user always reaches the intended target.

9.5 Conclusion & Contribution to the Research Questions

In this chapter, we took our second step towards an **Improvement** of proxy-based haptic feedback by advancing the field of HR. Following **RQ 3**, in this chapter, we aimed to take advantage of the knowledge about the worst-case detectability

of HR gained in chapter 8. Specifically, recognizing the lack of techniques that facilitate HR by exploiting particularities of human perception, such as visual suppression and resulting change blindness, our goal was to come forward with a novel HR algorithm that is closely tailored to human perception.

To this end, we add to **RQ 3** by presenting the first HR approach designed to take advantage of blink-induced change blindness. Motivated by recent advances leveraging blink suppression to enhance RDW [Langbehn et al., 2018b; Nguyen and Kunz, 2018], we developed the *Blink-Suppressed Hand Redirection (BSHR)* algorithm in order to study the feasibility of this novel take on HR, as well as the noticeability of the approach.

Our proposed technique grounds itself on the results of chapter 8, which showed that continuous HR with gradually growing offsets can, within the margin quantified in chapter 8, go unnoticed by users even when attentively observing the IVE. The fundamental concept proposed with *BSHR* is to split up HR into two separate phases: Firstly, while the user's eyes are opened, we propose to redirect the user's hand through continuous body warping (as proposed, e.g., by Azmandian et al. [2016a] and Cheng et al. [2017b], and as studied in chapter 8) *limited to a redirection amplitude below worst-case DTs*. By this, we introduce a certain amount of hand offset while ensuring that the manipulation remains unnoticed. Secondly, and in order to introduce the remaining hand offset necessary for redirection beyond the DTs of continuous HR, we propose to inject additional offsets as instantaneous, constant hand shifts when the user's vision is temporarily inhibited during a blink. Drawing from the effect of change blindness, *BSHR* thereby exploits an additional opportunity for covertly introducing manipulations, which has been left unused by previous techniques for HR—contributing to **RQ 3**.

To further uncover the limits of this approach, as well as to quantify the detectability of the state-of-the-art HR technique by Cheng et al. [2017b], we studied three different variants of *BSHR* alongside the algorithm by Cheng et al. in a psychophysical experiment. Here, we derived CDTs for (1) *BSHR* only instantaneously warping the hand during blinks ($BSHR_{+0\%}$), as well as for *BSHR* as proposed above and configured to continuously warp the hand up to (2) 50% ($BSHR_{+50\%}$) and (3) 100% ($BSHR_{+100\%}$) of the worst-case CDTs derived in chapter 8. Additionally, we also evaluated the technique by Cheng et al. [2017b] (*Cheng*), which does not leverage blinks. Our results verify the feasibility of blink-suppressed body warping as we found ranges of unnoticeable *BSHR* in the same order of magnitude as found in chapter 8 and for the conventional technique by Cheng et al. [2017b]. Moreover, we could show that for HR along the horizontal and vertical axis, more redirection can go unnoticed when combining continuous warping when the user's eyes are opened with instantaneous blink-suppressed shifts, compared to pure blink-suppressed warping.

These promising first results contribute to **RQ 3** and progress the field of HR. *BSHR*, e.g., could be employed for haptic retargeting to establish *Colocation* during proxy-based interactions taking advantage of regularly occurring spontaneous

blinks. Moreover, *BSHR* could be employed while exploring a proxy surface by manipulating the virtual hand on top of the discrepant virtual surface during blinks to establish *Similarity* with redirected touching. Finally, also pseudo-haptics might benefit from the concept of *BSHR* as blinks and resulting change blindness could be utilized, for example, to modify the offset between proxies and their virtual representation when rendering pseudo-weight.

In the context of **RQ 3**, our findings also highlight that the margins in which manipulations go unnoticed, even when blink-suppression is exploited, are rather narrow. This, in turn, highlights the importance of fine-tuning the parameters of *BSHR* to push the limits of unnoticeable HR. While we are optimistic that configurations of *BSHR* can be elaborated that covertly redirect the hands within ranges exceeding the CDTs derived here and in chapter 8, it remains to be shown that combining continuous HR and blink-suppressed shifts yields a direct advantage over conventional techniques. As such, our findings encourage continuing the exploration of approaches that take advantage of change blindness and visual suppression for HR. We are currently following up on our research presented here with our ongoing work studying the feasibility to unobtrusively trigger blinks on-demand. Moreover, we currently continue our work presented in this chapter with our ongoing work on novel HR techniques that aim to take advantage of saccadic suppression [Sun et al., 2018] in addition to blinks.

Chapter 10

Hand Redirection Toolkit

In the previous chapters, we introduced algorithms for rotational and gain-based body warping to study the detectability of HR, and proposed *Blink-Suppressed Hand Redirection (BSHR)*, a first body warping approach suitable to study the potential of blink-induced change blindness for HR. As a result, we also increased the need for an accessible platform that makes the various hand redirection techniques discussed among researchers accessible to developers. To contribute to the clarity and accessibility of research in this domain, in this chapter, we turn towards lowering the barriers of using and experimenting with hand redirection. To this end, we introduce a first open-source software framework that facilitates the creation of new techniques and offers reference implementations of several popular hand redirection algorithms proposed in the past. Additionally, it includes support of the algorithms introduced in chapters 8 and 9 of this thesis. We present our toolkit's architecture alongside an evaluation with expert users to conclude our work on **RQ 3** and our **Improvement** of proxy-based haptics through contributions to real-time virtual techniques.

A video⁷³ about the work presented in this chapter is available online. This chapter is based on the following publication. Images and parts of the text in this chapter, as well as the presented concepts, implementations, and results have been published previously therein:

Zenner, A., Kriegler, H. M., and Krüger, A. (2021a). HaRT – The Virtual Reality Hand Redirection Toolkit. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'21, pages 1–7. ACM. © 2021 André Zenner and co-authors. Final published version available in the ACM Digital Library.

DOI: [10.1145/3411763.3451814](https://doi.org/10.1145/3411763.3451814)



Video Link⁷³

⁷³Hand Redirection Toolkit Video. <https://bit.ly/3yX31ao>

The publication is based on the following Bachelor's thesis, which was conducted under the supervision of the author of this PhD thesis:

Kriegler, H. M. (2020). A Toolkit for Hand Redirection in Virtual Reality. Bachelor's thesis, Saarland University. Advisor: **Zenner, A.** © 2020 Hannah Kriegler.

10.1 Introduction

During the making of this thesis, the topic of hand redirection has received significant attention in the HCI and VR research communities as a tool to enhance ETHF systems [Abtahi et al., 2019; Gonzalez et al., 2020], to improve ergonomics in VR [Montano Murillo et al., 2017], and, most important for this thesis, to establish *Similarity* and *Colocation* in the context of proxy-based haptics [Nilsson et al., 2021a]. As our review of related work in chapter 2 illustrates, hand redirection allows the VR system to control the user's physical reaching motion, e.g., by displacing the virtual hand rendering as it is common in HR implementations. By this, hand redirection enables techniques such as haptic retargeting [Azmandian et al., 2016b] and redirected touching [Kohli, 2013a].

As the number of applications that can benefit from hand redirection increases, so does the number of algorithms for implementing it, as our review of body, world, and hybrid warping techniques in subsection 2.6.2 shows. Yet, up to now, no common framework or platform exists, which can serve as an entry point for researchers and practitioners interested in testing, using, improving, or evaluating hand redirection techniques. Knowledge and implementation details about hand redirection approaches are only available in research papers which leads to algorithms being continuously re-implemented, test scenarios being re-build from scratch, and the general development and evaluation of isolated solutions. The lack of reference implementations for the most common techniques impedes comparisons with new hand redirection approaches and hinders practitioners of non-research applications to apply hand redirection.

In order to address the consequential need for an easy-to-use framework, and inspired by the efforts previously taken in the domains of HCI and VR, in this chapter, we present the *Hand Redirection Toolkit (HaRT)*. Doing so, we follow an avenue that has proven to be of support for VR research in the past. The fact that VR researchers and developers mostly use a common set of hardware (i.e., VR HMDs, tracking systems, etc.) and software (i.e., 3D engines, SDKs, etc.) has led to the evolution of several open-source toolkits such as the *VRQuestionnaireToolkit*⁴⁷ by Feick et al. [2020b], the *Unity Experiment Framework*⁴⁶ by Brookes et al. [2020], our *Immersive Notification Framework* [Zenner et al., 2018b], and, most related to our work in this chapter, the *Redirected Walking Toolkit* by Azmandian et al. [2016a]. The latter serves as an inspiration for the *HaRT*, implementing methods for RDW in an easy-to-use and easy-to-extend way supported by additional functionality to simulate and analyze RDW scenarios.

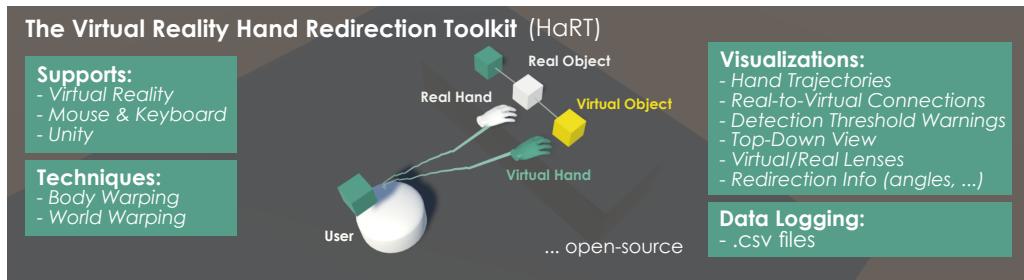


Figure 10.1: Summary of the central features of the *Hand Redirection Toolkit (HaRT)*.

The *HaRT* joins the ranks of the aforementioned frameworks as an open-source Unity package, which provides reference implementations of different hand redirection techniques that have been proposed and studied scientifically in the past. It comes with pre-built scenes and allows to easily add new hand redirection algorithms by exposing a useful class hierarchy and related data structures. Documented in an online wiki with step-by-step tutorials, the toolkit aims to fuel progress in this domain by lowering the barrier to hand redirection development. In addition to supporting common VR systems, it also comes with simulation functionality that enables testing and developing hand redirection scenarios without a VR system. Moreover, logging and visualization features provide support for evaluations and VR application development. This chapter introduces the structure and main features of the *HaRT*, and reports on a first qualitative user study conducted with expert VR developers.

10.2 The *Hand Redirection Toolkit*

The *Hand Redirection Toolkit (HaRT)* targets one of the most widely used 3D engines in the field of VR, i.e., the Unity engine¹⁷, to provide convenient access to common techniques for body and world warping. As a Unity package, it can easily be integrated into existing projects, offers pre-implemented algorithms and sample scenes, and exposes a modular class hierarchy with convenient data structures for easy extension. An online wiki in the open-source repository documents the framework, which encompasses real-time visualization and analysis features. Using the toolkit, it is no longer required to set up a VR system in order to work on implementations of hand redirection as movements of the real hand can be simulated using only mouse and keyboard. The *HaRT* was developed in the tradition of the *Redirected Walking Toolkit* by Azmandian et al. [2016a], which offers similar features for the domain of RDW, such as reference implementations of common gain-based RDW techniques and resetting controllers. It likewise features analysis features, takes into account DTs, and offers a software architecture that is easy to extend.

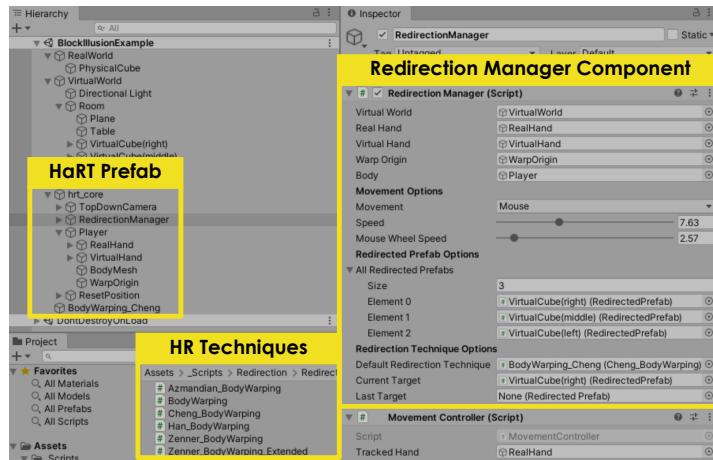


Figure 10.2: The Redirection Manager in Unity exposing settings and parameters.

10.2.1 Overview



GitHub Link⁷⁴

The *HaRT* is based on three central components. At its core, the *Hand Redirection Component* encompasses the implementation of the different hand redirection algorithms alongside required classes, data structures, and management logic. The *User Movement Component* integrates common VR SDKs and implements a script to use mouse and keyboard to move the real hand when no VR system is available. Finally, visualizations and logging functionality is implemented as part of an *Analysis Component*.

All relevant resources of the toolkit, such as the source code, Unity packages, and the wiki with a documentation and tutorials can be found online in an open-source repository.⁷⁴

10.2.2 Hand Redirection Component

As part of the *Hand Redirection Component*, the Redirection Manager script represents the heart of the hand redirection integration in a Unity scene. Using its Unity Inspector interface, shown in Figure 10.2, central parameters and settings can be applied (e.g., defining which objects in the scene belong to the real, and which to the virtual world). Moreover, the Redirection Manager starts and ends individual redirections. Each mapping of a virtual to a real object is represented by a Redirection Object script, which is added to the virtual object that is to be mapped. The script holds references to Virtual-To-Real Connection objects in the scene, each representing an individual mapping of a physical location to its corresponding virtual location. By this, Virtual-To-Real Connection objects define the individual R/V offsets present in

⁷⁴Hand Redirection Toolkit on GitHub. <https://github.com/AndreZenner/hand-redirection-toolkit>

the scene, which can be adjusted manually in the Unity Editor. Redirection techniques can be set for each mapping individually (via the corresponding `Redirection Object`) and a default technique can be set in the `Redirection Manager`. To trigger custom application code in response to the redirection, custom event callbacks can be registered that will be triggered, for example, when a redirection starts or ends.

The hand redirection techniques are implemented as individual scripts inheriting from the base classes `BodyWarping`, `WorldWarping`, or `3DInterpolation`. Each technique realizes their respective logic by overwriting an `Init` (called upon activation of a redirection), `ApplyRedirection` (called frame-wise during redirection), and `EndRedirection` (called upon termination) method. Within these methods, developers have access to all relevant data structures (e.g., real and virtual target locations). Due to this modular architecture, new hand redirection algorithms can be added easily by inheriting from one of the base classes or pre-implemented approaches. In order to add hand redirection support to a scene, users of the toolkit can simply drag a prepared prefab (`hrt_core`) into their scene, or start off modifying a sample scene. Due to its illustrative nature, we chose the 3-cubes illusion presented by Azmandian et al. [2016b] as a sample scene scenario shown in Figure 10.1 and Figure 10.3.

At this time, the *HART* encompasses nine pre-implemented techniques:

- *body warping*:
 - approach by Azmandian et al. [2016b]
 - approach by Cheng et al. [2017b]
 - approach by Han et al. [2018]
 - approach based on Kohli [2013a]
(leveraging inverse distance weighting [Shepard, 1968])
 - *Continuous Hand Redirection* as in chapter 8
 - *Blink-Suppressed Hand Redirection* as in chapter 9
- *world warping*
 - rotational approach based on Azmandian et al. [2016b]
 - translational approach based on Azmandian et al. [2016b]
 - combined approach based on Azmandian et al. [2016b]

Support is also indicated in the rightmost column of Table 2.3. In the future, we look forward to integrating additional algorithms.

10.2.3 User Movement Component

The *User Movement Component* is responsible for processing the user's motion input, i.e., the real head, hand, or controller movement. The toolkit is designed to

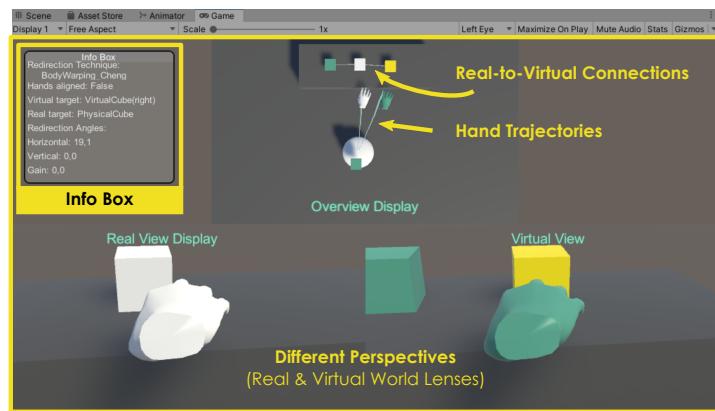


Figure 10.3: Overview of visualizations supported by the *HaRT*.

work with the widely-used VR platform SteamVR, supporting HTC Vive⁷⁵ and Oculus Rift systems.⁷⁶ Additionally, the toolkit supports tracking the real hands with Leap Motion.³⁵ For quick testing of a scene, toolkit functionality, or novel hand redirection concept, however, setting up a VR system seems tedious. For this reason, the toolkit includes a script that allows users to simulate real head and hand movements in a scene using mouse and keyboard. The movement method can be set via a dropdown menu in the *Redirection Manager*.

10.2.4 Analysis Component

The *Analysis Component* implements additional functionality that aims to lower the barrier to understanding, evaluating, and debugging hand redirection techniques and scenes. The toolkit creates log files in the `.csv`-format, which store for each frame a timestamped data set consisting of basic information about an applied redirection such as the positions and rotations of head and hands, selected target, and hand redirection technique. The resulting files can be used, for example, to analyze hand trajectories. Moreover, the *Analysis Component* provides five real-time visualizations to support developers, each of which can be individually switched on and off. Figure 10.3 depicts the *Scene Overview* visualization, which effectively implements the concept of a real-, and a virtual-world lens only displaying the content of the respective environment. These lens views are

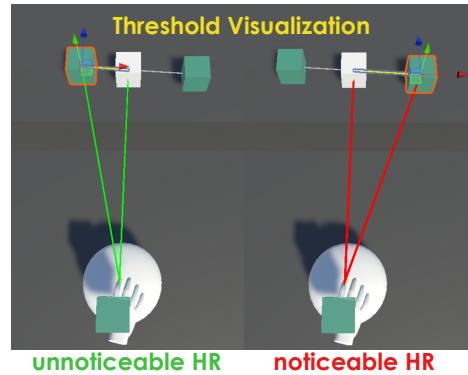


Figure 10.4: The threshold visualization.

⁷⁵HTC Vive Homepage. <https://bit.ly/3pxue0j>

⁷⁶Oculus Rift Homepage. <https://ocul.us/3HkKZ4K>

complemented by a top-down scene overview rendering both virtual and real objects simultaneously. *Trajectory Renderings* display the path traveled by the real and virtual hands in the 3D scene for visual comparison, and *Real-to-Virtual Object Connections* are displayed as line renderings during Edit and Play mode in Unity. This provides an overview of the R/V mappings configured in the scene. In order to create comfortable experiences, developers can also configure DTs in the Unity Inspector view of the BodyWarping scripts. As shown in Figure 10.4, *Threshold Warnings* then visually indicate if a redirection is going to be unnoticed by users, i.e., below thresholds, or not. Finally, an *Info Box* rendered in the Play view in Unity summarizes the current status of redirections at runtime.

10.3 Evaluation of the Hand Redirection Toolkit

To shed light on the usability of the toolkit and gather feedback from the target user group, we invited experts in Unity and VR development with an academic or industry background to participate in a qualitative remote user study. The study was approved by the Ethical Review Board of the Faculty of Mathematics and Computer Science at Saarland University. In this study, the experts were asked to use the toolkit and its wiki to solve a set of representative tasks using Unity. The goal of the evaluation was to collect whether the toolkit is supportive when (1) setting up a hand redirection scenario, (2) comparing two different hand redirection techniques, and (3) implementing a new hand redirection technique. Moreover, we asked participants (4) whether or not they regard the visualization features as useful.

10.3.1 Methodology

Our study was conducted remotely via Microsoft Teams⁷⁷ and was designed as an observation of different interactive tasks that participants completed using the toolkit, interleaved with assessments of the user experience, and concluded with a semi-structured interview. During the interview, participants were encouraged to think aloud and the experimenter observed the participant's interactions inside Unity through activated screen-sharing.

The study did not involve a VR system, instead participants simulated their head and hand movements using the mouse and keyboard functionality. Participants were sequentially introduced to four different tasks with varying complexity and independently solved these tasks by using only the information available in the toolkit and wiki. Interviewers observed this process, took the time for each task, took notes of problems, filled a binary task completion file for each task, and only intervened if participants could not solve tasks by themselves. After each task, participants filled out a short version of the User Experience

⁷⁷Microsoft Teams Homepage. <https://bit.ly/3FBFcaz>

Questionnaire (UEQ-S) [Schrepp et al., 2017; Schrepp, 2019]. Upon completion of all tasks, participants rated the usefulness of the visualization features provided by the toolkit by stating their agreement to the statement “*The <Visualization> was useful*” on a scale from 1 (= fully disagree) to 5 (= fully agree).

In the concluding semi-structured interview, participants were asked about the advantages and disadvantages of the toolkit, as well as their ideas for future improvements. The interviews were analyzed by means of a thematic analysis. Each session lasted approximately two hours and was recorded after the participant gave informed consent.

10.3.2 Participants

$N = 5$ experts (1 female, 4 male) with solid knowledge in Unity and experiences in VR development volunteered to take part in the study. Knowledge of hand redirection was not a requirement and only two experts had first-hand experience in implementing hand redirection, while all others only experienced hand redirection as a user before, or only heard of the concept but never experienced it. Participants were between 23 and 27 years old and every participant completed the interview.

10.3.3 Procedure

At the beginning of the study participants joined the video call and completed a consent form. The experimenter then introduced the concept of hand redirection and the way that body and world warping work from a technical perspective. After this, participants were granted access to the toolkit and wiki. Their first task was then to read an introductory section in the toolkit’s documentation, followed by creating a new Unity project and starting the screen-sharing.

Once the Unity project was created, participants completed four different tasks. The first task asked them to execute the *Getting Started* instructions in the wiki, i.e., downloading and importing the Unity package, followed by opening and testing a provided sample scene. Once completed, participants were to *Create a Body Warping Scenario*. They were instructed to set up a redirection scenario similar to the 3-cubes illusion by Azmandian et al. [2016b]. For this, a prepared scene and 3D models were provided. The task focused on adding a pre-implemented hand redirection technique to the scene and correctly configuring the *Redirection Manager*. In the third task, participants had to *Compare Redirection Techniques* by adding two different hand redirection algorithms to a single scene. Finally, *Adding a New Redirection Technique* was the goal of the last task. Here, the experimenter provided pseudo-code for a hand redirection algorithm that participants were to extend the toolkit with. To achieve this, participants added a new script implementing the pseudo-code using the toolkit’s class hierarchy and data structures. After each task, participants answered a

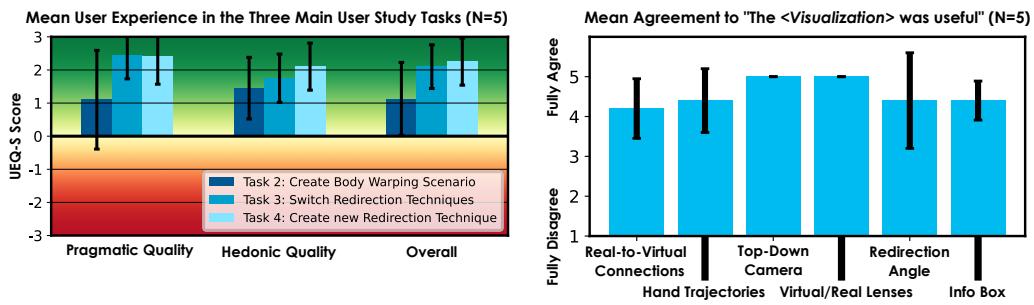


Figure 10.5: Left: UEQ-S results. Right: Agreement to “*The <Visualization> was useful*”.

UEQ-S through LimeSurvey.⁷⁸ At the end, the toolkit’s visualization features were introduced to the participant, who rated their usefulness. A semi-structured interview concluded the study.

10.3.4 Results

The observations revealed mostly positive user experiences, while also surfacing some usability issues of the evaluated version of the *HaRT*. Most of the errors made in the study concerned setting up the virtual scene while re-building the 3-cubes illusion. Here, for example, some participants wrongly placed Virtual-To-Real Connection objects in the scene, sometimes forgot to configure all settings inside a script, or assigned the hand redirection techniques only locally for individual warps instead of configuring them as the default redirection technique. Also, participants occasionally skipped a step described in the wiki. We assume some of these issues were rooted in the observed behavior of the participants, as four out of five only quickly read over the wiki instructions and approached the tasks mainly relying on their general Unity knowledge and a trial-and-error method to explore the toolkit. This, and a strong focus of participants on the example pictures in the wiki, resulted in them missing steps only described in the textual documentation. One of the participants, in contrast, first completely read all relevant sections in the documentation before working on the tasks in Unity. This participant did not make any mistakes throughout the study.

Concerning the toolkit’s usability, the results of the UEQ-S regarding pragmatic, hedonic, and overall quality are summarized on the left in Figure 10.5. The mean scale results show a positive evaluation of the toolkit’s usability across all tasks, with the most complex and error-prone second task being rated lowest but still in the neutral to positive range [Schrepp, 2019]. It is noteworthy that our study did not involve a counterbalanced order of tasks since the tasks built up on each other. Increased familiarity with the toolkit could thus be a reason for increased usability scores in later tasks. This, in turn, would provide a first indication that users can get familiar with the toolkit already after a relatively short time.

⁷⁸LimeSurvey Homepage. <https://bit.ly/3z2zfBg>

As can be concluded from the results depicted on the right in Figure 10.5, participants perceived all visualization features included in the *HaRT* as highly useful. The thematic analysis of the semi-structured interviews provided additional insights into the experts' opinion of the toolkit. When asked about the toolkit's advantages, all experts agreed that the toolkit and involved workflows save time. Moreover, four experts mentioned that they perceived it as effortlessly to switch between different hand redirection techniques inside the toolkit. Three experts also mentioned that they like the toolkit's structure, and the overview visualization for quick testing, perceived it as easy to add a new hand redirection technique, and found the toolkit to provide the basic functionality required to work with hand redirection. Themes regarding the toolkit's disadvantages and potential improvements were mostly mentioned only by individual experts. Most importantly, here, experts proposed making the wording of some toolkit components more intuitive, and to improve certain wiki instructions as, for example, the relationship between a *Redirection Object* and a *Virtual-To-Real Connection* object was perceived as confusing by one expert.

10.4 Discussion of the *Hand Redirection Toolkit*

Our first evaluation of the *HaRT* with expert users revealed that the general structure of the toolkit allows to achieve its main goals. As such, the *HaRT* presents a first ready-to-use tool in the toolbox of VR researchers and practitioners interested in working with hand redirection.

10.4.1 Benefits

Our qualitative study could reveal several benefits of the *HaRT*, as well as several points that can be further improved as the toolkit develops in future iterations. The experts' performance and feedback indicates that the toolkit is supportive when (1) setting up a new hand redirection scenario, (2) comparing two hand redirection techniques, and (3) implementing novel hand redirection algorithms. Moreover, (4) the visualization features offered by the toolkit were perceived as useful when working with the toolkit in a non-VR mode.

Our user behavior observations further highlighted the importance of intuitive wording inside the class hierarchy and documentation, which can inform future iterations of the wiki. Additionally, the experts underlined that video instructions and animations inside the documentation should ideally cover all steps of a tutorial. We also inferred from the expert interviews that, in order to provide a clean and helpful documentation, the wiki should separate basic explanations of hand redirection concepts and background information from step-by-step tutorials. This would make it easier for users of varying experience levels to find the information that is most relevant to them. However, despite the shortcomings of the wiki at the time of evaluation, our observations and collected usability

results indicate that users can get acquainted with the toolkit quickly. Thus, in summary, the toolkit was found to be time-saving and a helpful support when working with hand redirection in Unity.

10.4.2 Limitations

At this point, we also like to mention the limitations of the evaluation presented in this chapter. Most notably is the small sample size as we only recruited $N = 5$ experts, which does not allow us to draw statistically significant conclusions. Moreover, constrained by the COVID-19 pandemic situation in November 2020, we conducted our study in a remote setting and tested the toolkit's functionality in a non-VR mode only, which might have impacted the experts' impression of the toolkit. Lastly, to maintain a reasonable study duration, we reduced the scope of our evaluation to body warping. Testing further functionality (e.g., logging) and hand redirection techniques (e.g., world warping) is left to future studies.

10.5 Conclusion & Contribution to the Research Questions

Concluding Part IV, in this chapter, we took our third step towards an **Improvement** of proxy-based haptics by enhancing the accessibility of real-time virtual techniques. Specifically, we contribute to **RQ 3** by enabling broad and open access to the domain of hand redirection with our open-source Unity toolkit *HaRT*.

Up to now, integrating hand redirection into VR projects required careful research and re-implementation of existing algorithms, detailed information about which often only being available in academic research papers. Motivated by these circumstances that hinder a wider adoption of hand redirection, we presented the *Hand Redirection Toolkit (HaRT)*, a lightweight framework for integrating hand redirection into Unity projects. Our toolkit is designed to serve as a central contact point for the VR community interested in experimenting with, integrating, or improving hand redirection. It offers easy access to reference implementations of techniques proposed in past research (including the HR techniques studied in chapters 8 and 9), and exposes a useful class hierarchy and data structures to ease the implementation of novel algorithms. It supports common VR systems, offers a non-VR development mode, and a range of logging and visualization functionalities. Moreover, a detailed documentation including videos and tutorials provides a starting point for those new to the concept of hand redirection and experts alike. A first user study with $N = 5$ expert Unity developers indicated the toolkit to be well-structured, time-saving, fast to get acquainted with, and providing the basic functionality required to work with hand redirection.

By putting the *HaRT* into the hands of the research community, we enable everyone interested in experiencing the potential of hand redirection, as well as the limitations imposed by human perception, to "get hands-on" with the technique

in line with **RQ 3**. Furthermore, to keep the *HaRT* up-to-date as the field of hand redirection advances, we invite researchers and practitioners to contribute to the public repository of the *HaRT*. We look forward to contributions that might include extensions or improvements of the toolkit or its wiki, as well as new hand redirection techniques.

Chapter 11

Concluding Remarks

Our work in Part IV contributed an **Improvement** of proxy-based haptic feedback in an indirect way. Instead of directly proposing novel techniques for building proxies, as we did in Part III, in this part, we contributed to the flexibility of props by conducting fundamental research in the field of hand redirection. Through this, Part IV advanced a real-time virtual technique suitable to tackle the challenge of proxy *Colocation* by granting the VR system control over the user's real hand movement (e.g., when employing hand redirection to realize haptic retargeting or redirected touching).

Along this research path, we revealed the order of magnitude in which *Colocation* challenges in proxy-based haptics can be solved while leaving the user clueless of virtual manipulation techniques being applied. The results of chapters 8 and 9 motivate further research on how hand redirection techniques, especially those based on body warping (HR), can be improved to maximize the margins in which the VR system can retain control over the user's interaction with physical props.

We like to end Part IV with a conceptual analogy, specifically, an analogy between real-time virtual techniques that "*play with senses*" [Lécuyer, 2017] in VR and the field of cybersecurity: If one considers the interaction with physical proxies as a competition for control over the user's physical movement between the user and the VR system, real-time virtual techniques, like HR, can be regarded hacks of human perception applied by the VR system. In analogy to how hacks in the field of cybersecurity take advantage of software vulnerabilities, *perceptual hacks*, i.e., successful visual-haptic illusions, exploit the "vulnerabilities" of the human brain to achieve their goal of gaining control over the user's actions. Such *perceptual vulnerabilities* can be, for example, visual dominance (i.e., the particular way sensory integration works), or the phenomenon of change blindness (i.e., the susceptibility to miss changes that occur during recurring and inevitable

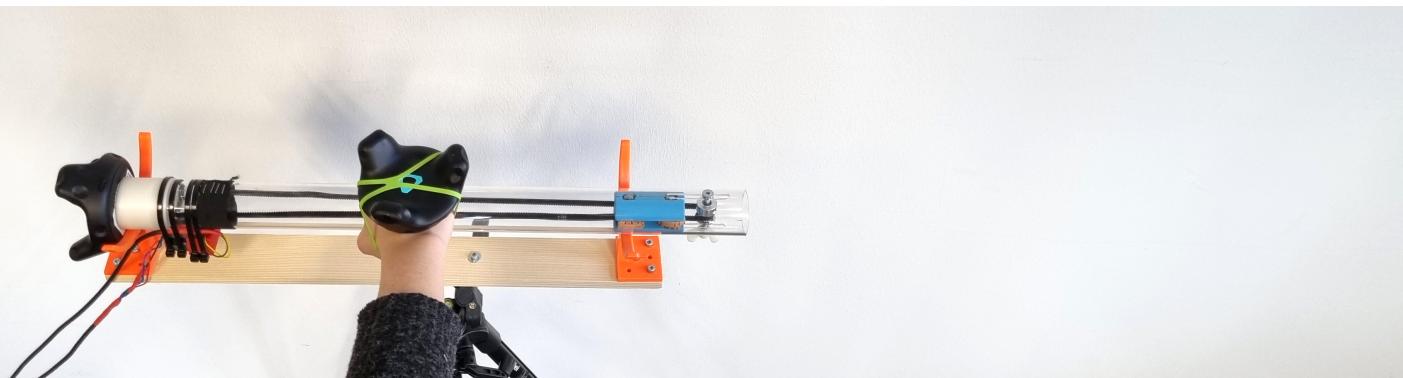
moments of visual suppression). Our work in chapter 8 highlighted how much unnoticeable HR can be achieved in the worst case when applying hacks based solely on visual dominance – the current “go-to attack vector” in the field of HR. Complementing this, chapter 9 informs about the potential of hacks that solely rely on blink-induced change blindness – an attack vector that is already known in related fields (such as RDW), but can be considered novel and a “zero-day attack” in the field of HR.

Following this line of thought, we put forward the idea of creating perceptual hacks that exploit not only one, but several perceptual vulnerabilities in orchestration. By this, we aim to increase the chances of an “unnoticed attack” that maximizes the control that the VR system gains over the user’s interaction, while leaving the user clueless and in the illusion of being themselves in full control. With our combination of continuous warping and blink-suppressed instantaneous shifts in chapter 9, we took a first step into this research direction. Our results with *BSHR* are promising as they show that collectively exploiting multiple perceptual vulnerabilities can indeed go unnoticed in totality. Moreover, they indicate an increase in the range of unnoticed manipulation with combined techniques compared to just exploiting one of the vulnerabilities (i.e., change blindness in our example). Hence, with our estimation of the CDTs for continuous HR in chapters 8 and 9, as well as our derivation of CDTs for blink-suppressed *BSHR* in chapter 9, we laid the foundation for future research and “novel attacks”. Our results can now be used for exploring if even more powerful HR techniques than those currently known can be created, which, with better tuning of the perception-related parameters, would allow for unprecedented control during proxy-based interaction.

This analogy, i.e., considering redirection techniques hacks of human perception, has also recently been put forward in the research community. Tseng et al. [2022], for example, started discussing the risks that come with the power of visual-haptic illusions and with the publication of psychophysical results. Using redirection approaches and knowledge about the boundaries of their detectability, malicious VR applications could use techniques like HR against users and abuse the increased control of the system over their real movements and actions. An attacker could, for example, covertly redirect users towards objects they would not normally touch, which could result in physical harm to the user or others in the environment, or to damage of surrounding property (an approach termed “*arm-movement puppetry attack*” [Tseng et al., 2022]). To mitigate such threats, research on techniques that can detect and prevent malicious manipulations is needed as the field of HR progresses [Tseng et al., 2022].

Part V

Enhancing Proxy-Based Haptics The *Combined* Approach



... *teamwork* makes the dream work

Research Question 4

Improvement

This part addresses **RQ 4:**

How can the physical approach of DPHF and the virtual approach of HR be combined to improve the flexibility of proxy-based haptic feedback?

Chapter 12

Combining Dynamic Passive Haptics and Hand Redirection

With our work on **RQ 2** in Part III, we examined how the flexibility of proxy-based haptic feedback for VR can be improved through real-time physical techniques. Specifically, we proposed the concept of Dynamic Passive Haptic Feedback (DPHF), i.e., dynamic props that leverage only minimal actuation to adapt their passive haptic qualities and by this are able to establish *Similarity* with various virtual objects. Conceptually orthogonal to this, in Part IV, we advanced the flexibility of proxy-based haptics by contributing to the field of real-time virtual techniques in the context of **RQ 3**. Here, we studied the margins in which Body Warping-based Hand Redirection (HR) can be employed without raising the user's awareness that their physical movements are being remote controlled by the VR system. With this knowledge, it is now possible to establish *Colocation* in proxy-based interaction through perceptually unnoticeable haptic retargeting.

With our contributions of Parts III and IV in place, in this concluding chapter, we take the next step towards an **Improvement** of proxy-based haptics and argue that both orthogonal concepts can be combined and that their combination yields further benefits. Specifically, we put forward that the combination of a real-time physical with a real-time virtual technique can better solve the challenges of *Similarity* and *Colocation* than the individual techniques alone can do. To demonstrate and validate this, we build on our work in Part III and Part IV, and propose a combination of DPHF and HR that we systematically compare to the individual techniques. We describe how we combine DPHF and HR to render haptic impressions in a proof-of-concept scenario, define metrics to measure the performance in solving the challenges of *Similarity* and *Colocation*, and perform two thought experiments. From these thought experiments, we derive two hypotheses which we eventually investigate in a controlled user experiment.



Video Link⁷⁹

A video⁷⁹ about the work presented in this chapter is available online. This chapter is based on the following two publications. Images and parts of the text in this chapter, as well as the presented concepts, implementations, and results have been published previously therein:

Zenner, A. and Krüger, A. (2020). Shifting & Warping: A Case for the Combined Use of Dynamic Passive Haptics and Haptic Retargeting in VR. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST'20 Adjunct, pages 1–3. ACM. © 2020 André Zenner and Antonio Krüger. Final published version available in the ACM Digital Library. DOI:⁸⁰ [10.1145/3379350.3416166](https://doi.org/10.1145/3379350.3416166)

Zenner, A., Ullmann, K., and Krüger, A. (2021c). Combining Dynamic Passive Haptics and Haptic Retargeting for Enhanced Haptic Feedback in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2627–2637. © 2021 IEEE. Final published version available in the IEEE Xplore® Digital Library.

DOI: [10.1109/TVCG.2021.3067777](https://doi.org/10.1109/TVCG.2021.3067777)

12.1 Introduction

In order for proxy-based haptic feedback to be successful, past research identified two central challenges that both need to be solved by interactive VR systems. As we have introduced in section 2.4 [Lohse et al., 2019; Strandholt et al., 2020; Nilsson et al., 2021a,b], these are the challenges of:

1. *Similarity*: proxies and virtual objects must be sufficiently similar in terms of their haptic properties to convey convincing perceptions.
2. *Colocation*: proxies and virtual objects must be spatially colocated to enable seamless interactions.

To tackle these challenges, several approaches have been proposed by previous research and in this thesis. Among those that received significant research attention in this thesis and the research community lately are the techniques of:

1. **Dynamic Passive Haptic Feedback (DPHF)** (proposed and studied in Part III): a hardware-based technique that extends props with actuation enabling a dynamic adaptation of their physical properties at runtime. The technique aligns a prop's passive haptic feedback to different virtual objects – targeting primarily the challenge of *Similarity*.
2. **Body Warping-based Hand Redirection (HR)** ([Kohli, 2013a; Azmandian et al., 2016b], studied in Part IV): a software-based technique manipulating the virtual hand rendering to control the real hand movement. By leveraging hand redirection for haptic retargeting [Azmandian et al., 2016b], users can seamlessly touch virtual objects that are dislocated from their proxies – targeting primarily the challenge of *Colocation*.

⁷⁹Combined DPHF and HR Video. <https://bit.ly/3Ju65zK>

⁸⁰As this publication cannot be accessed via its DOI, we provide a direct link to the publication in the ACM Digital Library here.

While individually successful, up to now, these two concepts have been considered only in isolation when it comes to rendering most haptic qualities. Looking at their respective strengths, however, both techniques promise to complement each other well – hinting at a potential yet unused in the domain of proxy-based haptics. This potential is further emphasized by the works of Ban et al. [2014] and Gonzalez et al. [2021b], who started to explore the benefits of combining shape-changing DPHF and HR. Moreover, the interest in research on combined techniques is fueled by successful combinations of HR and ETHF as reviewed in section 2.7. Here, for example, Abtahi et al. [2019] and Gonzalez et al. [2020] leveraged HR to compensate for *Colocation* issues inherent to ETHF. In their systems, the authors used HR to redirect the user’s hands to robotic elements providing just-in-time haptic feedback. The addition of HR thereby allowed for reliably reaching robotic proxy elements despite spatial inaccuracies that result, for example, from aerial drone hovering or time constraints.

Motivated by these previous findings and the lack of research on combining DPHF and HR, in this chapter, we combine both approaches in a proof-of-concept scenario focusing on the haptic quality of weight distribution. In other words, we transfer the idea of combining DPHF and HR for the first time to the domain of inertia-based haptics as studied in chapter 5 and apply a combined approach for rendering weight shift. By this, we aim to showcase and validate that combined DPHF and HR techniques are not only suitable to enhance shape *Similarity* [Ban et al., 2014], but likewise help solving the challenges of *Similarity* and *Colocation* in scenarios where proxies with varying weight distribution are employed, as we advocate in chapter 5.

12.2 Combined DPHF and HR

In order to study how dynamic proxy adaptation can act in concert with redirection of the user’s hand, we start by defining a proof-of-concept scenario, which will constitute the focus of our investigation. This proof-of-concept scenario is chosen so as to revolve around a haptic quality different than shape, distinguishing itself from the scenario investigated by Ban et al. [2014]. Following the definition of our scenario, we outline the corresponding haptic rendering strategies of DPHF, HR, and *Combined DPHF and HR*.

12.2.1 Proof-of-Concept Scenario

Our choice of the proof-of-concept scenario was governed by several considerations. In particular, we aimed to choose a scenario for our investigation, which:

1. concerns a basic and versatile haptic feedback property
2. only involves low-complexity actuation and principles
3. involves effects that can easily be observed and explained visually

Based on these criteria, we opted for the haptic property of *weight distribution*, which is central to chapter 5, as a scenario-defining property. The haptic dimension of weight distribution (1) nicely aligns with these requirements, and (2) has seen significant research interest [Lim et al., 2021] with several DPHF devices based on weight shift having been proposed recently [Krekhov et al., 2017; Shigeyama et al., 2019; Liu et al., 2019; Sagheb et al., 2019]. Hence, we defined our proof-of-concept scenario to be concerned with a basic challenge for proxy-based haptics, namely, *rendering the weight distribution inside a virtual stick*.

In particular, our scenario involves a user reaching out to grasp and lift up a virtual rod inside the IVE. Following the paradigm of proxy-based haptics, this rod is physically represented by a tubular proxy of the exact same diameter and length as the virtual stick. In our proof-of-concept interaction, users will see the virtual stick lying horizontally in a stand in front of them, grasp the virtual stick with their virtual hand at its geometric center, and lift it up vertically as shown in Figure 12.1.



Figure 12.1: A participant lifting a physical proxy in our proof-of-concept scenario. © 2021 IEEE.

While the users in our proof-of-concept scenario will lift and hold the virtual stick, our investigation will assess how they perceive the stick's weight distribution. To perceive the weight distribution of an object in our hand, our perceptual system makes use of sensors in our muscles, tendons, and skin, as introduced in our review of dynamic touch in subsubsection 2.2.3 [Turvey, 1996; Kingma et al., 2004]. In this perceptual process, two physical parameters are of importance:

1. the *lever*, defined by the distance between the grasp location (i.e., where the force lifting the object acts vertically upwards) and the object's Center of Mass (CoM) (i.e., where gravitation pulls the object downwards).
2. the *moment of inertia*, which can be viewed as the rotational resistance defined by the distance of the object's mass to the hand.

While our work in chapter 5 took advantage of the perception of inertia and weight-shifting proxies to convey virtual objects varying in weight, length, and thickness, in this scenario, we will utilize a weight-shifting proxy identical in construction to *Shifty* for investigating the *perception of balance* of a virtual stick.

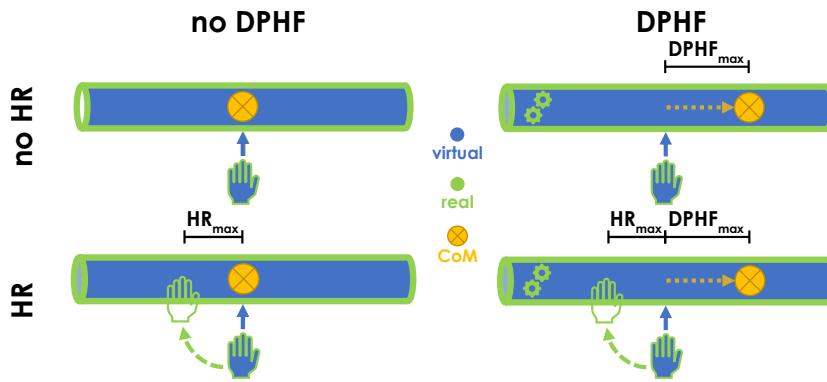


Figure 12.2: The four conditions we compare in this chapter. Our baseline *BL* (upper left) is a balanced passive proxy that does not employ any adaptation, nor hand redirection. The DPHF technique (upper right) relocates the physical CoM of the proxy while a 1-to-1 hand mapping is used. The HR technique (lower left) redirects the real hand to an offset grasp location while the physical CoM stays at the proxy’s center. The combined *DPHF+HR* technique (lower right) first shifts the physical CoM in one direction and additionally redirects the real hand in the opposite direction to increase the lever. ©2021 IEEE.

12.2.2 Haptic Rendering Techniques

To simulate varying states of balance in our proof-of-concept scenario, i.e., different weight distributions inside the virtual stick, we employ three different proxy-based haptic rendering strategies:

1. *DPHF*: leveraging *only* a dynamic proxy, which adapts its weight distribution to produce different states of balance
2. *HR*: leveraging *only* body warping, which redirects the real hand to an offset grasp location to make the user perceive different states of balance
3. *DPHF+HR*: leveraging a combination of the two approaches

In our thought experiments and user study we will compare these three rendering techniques to each other. In addition, our user study will involve a baseline condition (*BL*) as control.

Baseline (*BL*) The baseline condition *BL* represents conventional PHF and only renders a static, balanced proxy state. It does not use any HR, but only employs a 1-to-1 hand mapping, and does not involve any physical adaptation of the proxy. This condition is implemented with *Shift*.

Dynamic Passive Haptics (DPHF) The first actual rendering technique makes use only of a DPHF prop to convey different states of weight distribution inside the virtual stick. When the technique of DPHF is employed, the user’s virtual hand is registered to the real hand in a 1-to-1 manner, i.e., no HR is applied. As a

result, when reaching for the geometric center of the virtual stick, the user will also grasp the proxy at its geometric center.

Conceptually, for this rendering technique we assume a proxy that can dynamically locate its CoM along its main axis in the range $com \in [-DPHF_{max}, DPHF_{max}]$, with 0 being the geometric center of the prop. When the proxy is perfectly collocated with the virtual stick and its CoM location is shifted to a position com , the user experiences a lever of $|com|$ as the hand grasps at the center location 0. Thus to render a virtual CoM at location com_{target} , the proxy shifts its internal mass in the corresponding direction to move its physical CoM to com_{target} .

To implement this technique we use the weight-shifting proxy *Shifty* introduced in chapter 5. As in our previous experiments, *Shifty* uses its stepper motor to shift its mass along its tubular acrylic body. The proxy's implementation is identical in construction to that in chapter 5. The only difference is that in this experiment, the proxy is tracked with an HTC Vive Tracker⁵⁴ attached on its right end, which acts as a counterweight to the motor located on the left side as can be seen in Figure 12.1. Our implementation achieves $|DPHF_{max}| = 5cm$.

Hand Redirection (HR) Our second actual rendering technique, the HR technique, does not involve any proxy actuation, but instead relies on unnoticeable HR to convey weight shifts inside the virtual stick. As illustrated in Figure 12.2, the technique utilizes only a passive, balanced proxy (as used for *BL*), which in our experiment is implemented by *Shifty* keeping its CoM stationary at the center (i.e., in the state $com = 0$). When grasping the virtual stick, users will reach for its geometric center with their virtual hand. With HR being applied, the virtual hand is incrementally offset from the real hand (i.e., the mapping is no longer 1-to-1), and users will compensate for the offset by redirecting their real hand trajectory. As a result, this technique lets users grasp the physical stick at a location that is spatially offset from the virtual stick's center, introducing a lever.

Following the definition of DPHF introduced in chapter 4, the adaptation of the proxy's CoM in the DPHF technique takes place when the user is not in contact with the proxy so that it goes unnoticed by users. Hence, to achieve comparability between the DPHF and HR techniques, we also implement HR so that it goes unnoticed by the user. For this, the maximum range of redirection applied in the HR technique is limited to the worst-case CDTs derived in chapter 8 (i.e., the real and virtual hand drift apart horizontally at an angle of 4.5°).

Consequently, in order to convey a virtual CoM at com_{target} , the HR technique redirects the user's real hand to the location $grasp = -com_{target}$ on the physical proxy. By this, the technique ensures the grasp location to be $|com_{target}|$ away from the physical CoM (which remains at position 0) and in the opposite direction of the virtual shift. Constrained by the unnoticeability ranges, the grasp location on the proxy is limited to the range $grasp \in [-HR_{max}, HR_{max}]$.

We implement the HR technique in our experiment using the body warping approach by Cheng et al. [2017b], which, according to our findings in chapter 9,

can be assumed to go unnoticed within the thresholds derived in chapter 8. The technique by Cheng et al. [2017b] continuously increments the offset between the virtual and the real hand as the user reaches from an origin location (here: 60cm in front of the center of the virtual stick) to the virtual grasp location (here: the center of the virtual stick). We use the calibration outlined below in paragraph 12.3.4 and an HTC Vive Tracker⁵⁴ attached to the back of the user’s hand for tracking, as can be seen in Figure 12.8. Applying the thresholds found in chapter 8, our setup yields a maximum unnoticeable redirection range along the proxy’s body of $|HR_{max}| = \frac{\sin(4.5^\circ)}{\cos(4.5^\circ)} \cdot 60\text{cm} = 4.72\text{cm}$.

Combined DPHF and HR The third actual rendering technique combines DPHF, i.e., physical relocation of the CoM, and HR, i.e., redirection of the physical hand, as follows: In order to render CoM shifts up to $|DPHF_{max}|$, the combined technique applies a 1-to-1 hand mapping and relies solely on DPHF, i.e., weight shifts inside the proxy, to achieve the desired lever. Only when the target lever $|com_{target}|$ exceeds the shift capabilities of DPHF, the technique starts applying HR in order to increase the distance of the grasp location to the physical CoM. By this, the combined technique effectively increases the physical lever beyond the shift range of the DPHF device. Formally, this means that when

$$|com_{target}| > |DPHF_{max}| \quad (12.1)$$

the user’s real hand is redirected by

$$\min(|com_{target}| - |DPHF_{max}|, |HR_{max}|) \quad (12.2)$$

while in addition, the DPHF proxy performs a maximum weight shift. Equation 12.2 thereby represents the remaining lever distance not covered by DPHF alone, capped at the maximum unnoticeable HR range. It is noteworthy that the redirection of the real hand is always towards the direction opposite to the weight shift so as to increase the lever effect given by:

$$lever = |grasp - com| \quad (12.3)$$

12.3 Evaluation of Combined DPHF and HR

To validate that the combined proxy-based rendering technique *DPHF+HR* is advantageous compared to the individual physical and virtual strategies, we conducted a systematic evaluation. In this evaluation, we examined how well the different rendering techniques can solve the *Similarity* and *Colocation* challenges. To do this, we first define two metrics that quantify the “success” in tackling the challenges in our proof-of-concept scenario. Based on these metrics, we then first evaluate the different approaches from a theoretical point of view, before we turn towards validating our theoretical results in a user experiment.

12.3.1 Metrics

In order to compare the performance of the different rendering techniques in handling the challenges of *Similarity* and *Colocation* when rendering the weight distribution of a virtual stick, we define the following two metrics.

Similarity Metric The first metric captures what it takes for a VR system to solve the challenge of *Similarity*: the ability to provide, for each interactable virtual object, a haptic experience that is most similar to what the user would expect from the interaction in reality. Mapped to our scenario, this means that the system needs to provide convincing haptic sensations for as many different weight distributions inside the virtual stick as possible.

Thus, we regard a rendering technique that conveys a greater range of virtual CoM locations as superior to a technique that cannot provide as large a range of CoM locations since a larger rendering range implies that more virtual objects can be simulated. Consequently, the performance of the techniques in regard to *Similarity* is measured by:

$$\text{lever}(\text{Technique}) = \frac{\text{the maximum virtual lever}}{\text{that can be conveyed with Technique}} \quad (12.4)$$

It is given by the average location at which users perceive the virtual CoM relative to the virtual grasp location at the center of the virtual stick.

Colocation Metric Our second metric considers what it takes for a VR system to provide an optimal solution to the *Colocation* challenge. Such a solution is achieved when users can reach for any virtual object, no matter where it is located in the IVE, and receive a sensation of touch that matches the way their virtual hands touch the virtual object. Transferred to our scenario, the system should provide the sensation of lifting a balanced stick for as many different virtual stick locations as possible, even if the virtual stick is spatially offset from its proxy.

We consequently regard a technique as superior if it allows for larger spatial offsets of prop and virtual stick to go unnoticed. Based on this, we gauge the performance in solving the *Colocation* challenge through:

$$\text{offset}(\text{Technique}) = \frac{\text{the maximum offset between proxy and virtual stick}}{\text{that goes unnoticed with Technique}} \quad (12.5)$$

We restrict our investigation to offsets along the proxy's main axis, leveraging the fact that the proxy has a uniform and symmetric tubular shape. This way, users can only detect offsets by feeling an unexpected imbalance of the virtual stick when lifting it up.

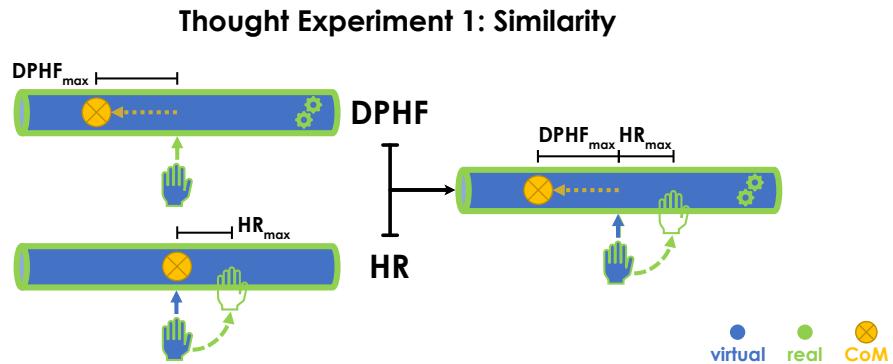


Figure 12.3: Conceptual sketches of how the DPHF, HR, and combined DPHF+HR strategies are applied in the first thought experiment arguing about *Similarity*.

12.3.2 Thought Experiments

Based on our introduced techniques and metrics, the expected performance of the different approaches in solving the *Similarity* and *Colocation* challenges can be thought-through. For this, we take the fundamental laws of physics into account and argue about the maximum achievable *lever* and *offset* from a theoretical perspective in two short thought experiments.

Thought Experiment 1: *Similarity*

In the first thought experiment, we derive the theoretically achievable *lever* when DPHF, HR, and the combined technique *DPHF+HR* are applied to convey weight shift inside a virtual stick. For this, and in line with our scenario, we assume a perfectly colocated prop and a user lifting the virtual stick by grasping it virtually at its geometric center. Figure 12.3 compares how the three different strategies can produce their respective maximum weight shift effect.

Using only a weight-shifting dynamic prop (DPHF), the maximum achievable lever distance is given by the maximum range that the device can shift its CoM, i.e., constrained by the physical specifications of the dynamic proxy. Hence, we conclude that

$$\text{lever(DPHF)} = |DPHF_{max}| \quad (12.6)$$

as the hand will grasp the proxy at its geometric center ($grasp = 0$) while the CoM of the prop is relocated maximally ($com = DPHF_{max}$).

Using HR, the user will physically lift a passive prop that also is identical in shape to the virtual stick and perfectly colocated with it. In contrast to DPHF, however, here, the proxy will not have adapted its weight distribution to convey a weight shift. Instead, it will remain balanced with its CoM fixed at the geometric center ($com = 0$). In this case, the maximum achievable shift effect is constraint by the maximum unnoticeable grasping offset towards the end of the stick that is supposed to feel more lightweight ($grasp = -HR_{max}$), i.e., by the detectability of

the HR approach. Thus, we conclude that

$$\text{lever}(\text{HR}) = |\text{HR}_{\max}| \quad (12.7)$$

Enlarging the range of weight shifts conveyable with DPHF or HR would require either an increase of the weight inside the prop, its shifting range, or redirection of the hand beyond perceptual DTs. Such measures, however, are not required when using both approaches in a combined way. As the laws of physics tell, the combined use of *DPHF+HR* should enable rendering of increased effect ranges without modifying the prop or redirecting beyond thresholds. Instead, when applying *DPHF+HR*, the weight of the prop can be shifted to a maximum extent ($\text{com} = \text{DPHF}_{\max}$), while the user's hand can be maximally redirected towards the opposite direction ($\text{grasp} = -\text{HR}_{\max}$), effectively yielding

$$\text{lever}(\text{DPHF+HR}) = |\text{DPHF}_{\max}| + |\text{HR}_{\max}| \quad (12.8)$$

Considering our implementation with *Shiftify*, this approach theoretically achieves a maximum lever distance of $\text{lever}(\text{DPHF+HR}) = 5\text{cm} + 4.72\text{cm} = 9.72\text{cm}$.

Thought Experiment 2: *Colocation*

In the second thought experiment, we derive the theoretically achievable *offset* that can covertly be compensated for by the techniques of DPHF, HR, and the combined technique *DPHF+HR*. For this, we again assume a user lifting a virtual stick at its geometric center. In contrast to the first thought experiment, however, this time the user expects the stick to feel balanced – even if the virtual stick and its proxy are not perfectly colocated. To achieve this, the system in this experiment employs the three haptic rendering techniques not to convey various states of imbalance (as in the first thought experiment), but to maintain the perception of a balanced stick. In other words, the system compensates for unwanted perceptions of imbalance that might arise as a consequence of spatial offsets⁸¹ between the virtual stick and its proxy. Figure 12.4 illustrates for each technique the maximum dislocation that can be compensated for while maintaining the perception of lifting a balanced virtual stick grasped at the center.

As Figure 12.4 shows, the weight-shifting prop can be dislocated only by up to

$$\text{offset}(\text{DPHF}) = |\text{DPHF}_{\max}| \quad (12.9)$$

when DPHF is used for such compensation, because the prop cannot align its physical CoM with the geometric center of the virtual stick for dislocations beyond $|\text{DPHF}_{\max}|$.

In turn, when employing only HR and a passive prop with fixed CoM at its center, the maximum displacement that can be compensated for is constrained

⁸¹For the sake of simplicity, in this scenario, we restrict our investigation to offsets along the stick's main axis and assume the virtual stick and its proxy to be equally oriented.

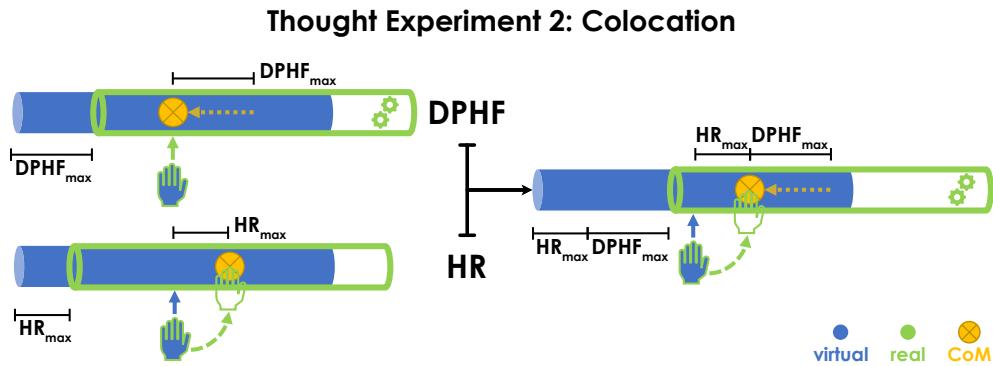


Figure 12.4: Conceptual sketches of how the DPHF, HR, and combined $DPHF+HR$ strategies are applied in the second thought experiment arguing about *Colocation*.

by $|HR_{max}|$. If displaced any further, the user's hand cannot unnoticed be redirected to the physical CoM, yielding a corresponding

$$offset(HR) = |HR_{max}| \quad (12.10)$$

As in the case of the first thought experiment, proxy modifications or redirection beyond thresholds would be required to increase the range of displacement that the individual techniques can compensate for, unless they are combined. To compensate for larger displacements, the laws of physics suggest that the combined technique could firstly shift the physical CoM by $|DPHF_{max}|$ towards the virtual CoM, and secondly bridge remaining spatial offset of up to $|HR_{max}|$ by applying HR. The maximum unnoticed displacement for $DPHF+HR$ is consequently expected to be

$$offset(DPHF+HR) = |DPHF_{max}| + |HR_{max}| \quad (12.11)$$

Conclusion

Our two thought experiments advocate the combined use of DPHF and HR. Theoretical considerations based on the basic laws of physics, the metrics outlined above, and the three introduced techniques show for our example scenario that the haptic effect ranges achievable with $DPHF+HR$ are greater than those achievable with DPHF and HR alone (in fact, they sum up). As a consequence, the first thought experiment suggests that $DPHF+HR$ allows individual dynamic props to represent even more virtual objects, showcasing the benefit of a combined use of DPHF and HR for solving the challenge of *Similarity*. Moreover, the second thought experiment leads us to expect that combined use of $DPHF+HR$ can also allow for larger prop displacements to go unnoticed than when applying only DPHF or HR alone for compensation. This demonstrates the potential of the combined techniques to solve the challenge of *Colocation* in proxy-based haptics.

12.3.3 Hypotheses

Based on the findings of our two thought experiments and the metrics for *Similarity* and *Colocation* introduced above, we formulate the following two central hypotheses, which we aim to validate in a user study:

H1: DPHF and HR can achieve greater perceived CoM shifts than the baseline *BL*, and the combined technique *DPHF+HR* can achieve greater perceived CoM shifts than both DPHF and HR.

Formally:

$$\text{lever}(\text{BL}) < \text{lever}(\text{DPHF}), \text{lever}(\text{HR}) < \text{lever}(\text{DPHF+HR})$$

H2: DPHF and HR can compensate for greater spatial offsets than the baseline *BL*, and the combined technique *DPHF+HR* can compensate for greater spatial offsets than both DPHF and HR.

Formally:

$$\text{offset}(\text{BL}) < \text{offset}(\text{DPHF}), \text{offset}(\text{HR}) < \text{offset}(\text{DPHF+HR})$$

Since $|\text{DPHF}_{\text{max}}| \approx |\text{HR}_{\text{max}}|$ in our implementation, we do not expect perceivable differences between DPHF and HR regarding both measures.

12.3.4 User Experiments

While the theoretical benefits of combining DPHF and HR become apparent from the considerations in our thought experiments, it remains unclear, if (1) the predicted benefits are achievable in a practical implementation, and if (2) the extension of the haptic rendering range would be perceivable with state-of-the-art implementations in use. Hence, we aim to verify the two hypotheses outlined above experimentally. By this, we intend to demonstrate the practical value of combining both techniques for the domain of weight shifting haptics.

Our study leverages established perceptual experiment designs, i.e., a psychophysical method, and validates two important aspects. In particular, our experiment uncovers whether a combination of a state-of-the art DPHF implementation (i.e., the weight-shifting proxy *Shifty* introduced in chapter 5) and unnoticeable HR (i.e., body warping based on Cheng et al. [2017b] within the CDTs derived in chapter 8) enhances the haptic rendering capabilities of proxy-based VR. Furthermore, it will reveal whether, if such augmentation is practically achievable, the effects are also perceivable for users.

To investigate **H1** and **H2**, our study is split into two separate experiments. In line with our thought experiments, both experiments employ the rendering techniques DPHF, HR, and *DPHF+HR*, as well as the control technique *BL*, as introduced in subsection 12.2.2. The *Similarity* experiment compares the techniques with regard to the maximum weight shift they can convey. Complementarily, the *Colocation* experiment applies a psychophysical method to capture how much the virtual stick can be offset from its proxy without users noticing it. The study

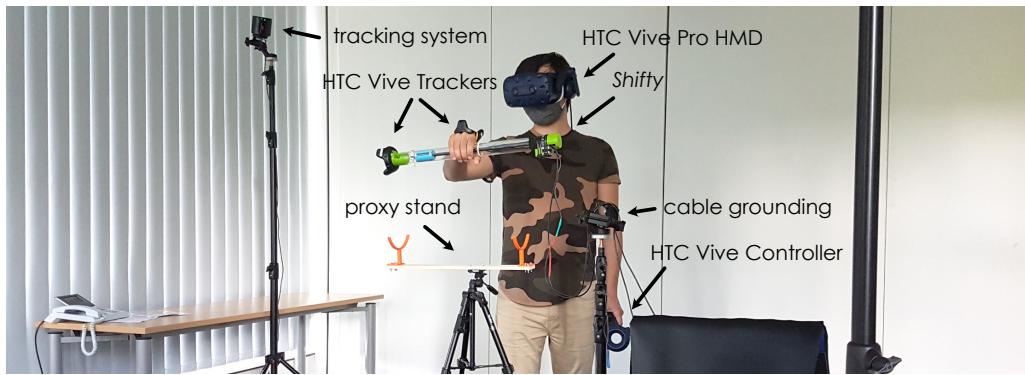


Figure 12.5: Experiment setup and involved components. The cable of the proxy has been grounded with a tripod to minimize gravitational pull. © 2021 IEEE.

was approved by the Ethical Review Board of the Faculty of Mathematics and Computer Science at Saarland University, and rigorous disinfection protocols were followed due to the COVID-19 pandemic situation.

Participants

$N = 24$ (6 female, 18 male) participants recruited from the local campus volunteered to participate in the study, each taking part in both experiments in a counterbalanced order. Participants were from 22 years to 36 years old ($M = 26$ years, $SD = 3.54$ years). All participants had normal or corrected-to-normal vision and were right-handed. We asked participants how often they play 3D video games, use VR systems, and work with their hands on a scale from 1 (= never) to 7 (= regularly). Our set of subjects covered a wide range of experience with 3D video games ($M = 4.17$, $SD = 2.30$) and VR systems ($M = 3.50$, $SD = 2.19$), both with responses ranging from 1 to 7. Responses regarding their experience in working by hand ($M = 4.12$, $SD = 1.94$) ranged from 2 to 7.

Apparatus

The study took place in a lab at our institution. A notebook with an NVIDIA GTX 1070 graphics card and an HTC Vive Pro⁵⁶ HMD were used to immerse participants visually and auditorily, using a tracking system with SteamVR base stations 2.0. The dominant hand of participants was tracked with an HTC Vive Tracker⁵⁴ (v2018) attached to the back of the hand with a rubber band. Participants used an HTC Vive Controller⁵³ in their non-dominant hand to answer questions in VR. The weight-shifting proxy used in both experiments was identical to *Shifty* as outlined in chapter 5. VR system and proxy communicated via WiFi and the proxy was tracked with an HTC Vive Tracker (v2018). A custom stand was placed in front of the participants as shown in Figure 12.5. Including attachments, the proxy weighed 615g and moved the mass from end to end in 2.8s.



Figure 12.6: View of the IVE showing the instruction display and the virtual stick. Participants align their hand with the hologram to start a trial. ©2021 IEEE.

The impact of the weight of the cables on the proxy's left end was carefully considered when designing the experiment by (1) minimizing it through grounding with a tripod (as can be seen in Figure 12.5), and (2) carefully choosing the directionality of the experiment. This means that any potential effect due to cables pulling down the left end of the proxy would only act against our hypotheses and only raise the bar for validating the expected effects. Moreover, to prevent participants from inferring information about the location of the weight from motor sounds, an obfuscation technique was employed. For this, the weight always moved to a random location first, before moving to its actual destination before each trial. The IVE was implemented using version 2019.3.7f1 of the Unity engine¹⁷, the *Unity Experiment Framework*⁴⁶ by Brookes et al. [2020], and the *VRQuestionnaireToolkit*⁴⁷ by Feick et al. [2020b].

Similarity Experiment

Following the first thought experiment, the *Similarity* experiment investigates **H1**, i.e., which of the rendering techniques can convey the strongest CoM shift inside the virtual stick.

Procedure After providing informed consent, each participant started with a calibration of the experiment setup. A point on the surface of the user's palm served as the reference location for the real hand in the scene. To calibrate this point, we established the accurate offset from the tracker on the back of the hand to this point on the palm, which is in contact with the prop when grasping it. For this, participants were asked to grasp the physical stick exactly as indicated by a green holographic hand displayed in VR. A haptic marker on the proxy surface felt underneath the middle finger indicated correct alignment during calibration. When correct alignment was achieved, calibration was completed and a transform representing the reference location of the real hand continued to follow the tracker with the calibrated offset for the remainder of the experiment.

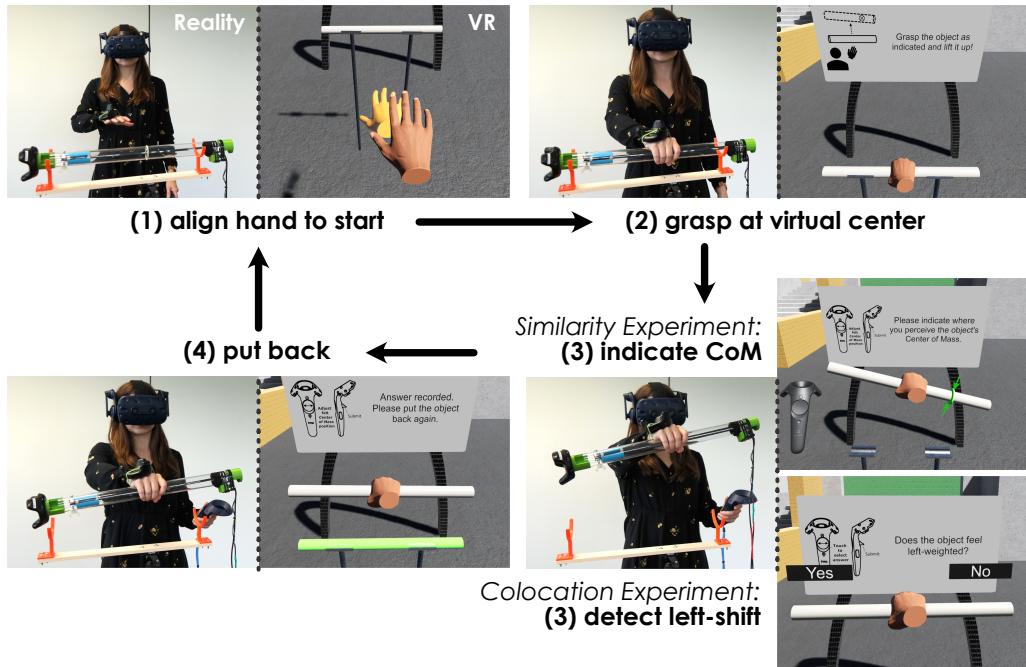


Figure 12.7: Step-by-step view of an experiment trial. The user starts the trial by aligning the calibrated virtual hand with a holographic hand rendered in VR. After 2s, any HR applied in the trial is activated as the user starts to grasp the virtual stick at its center. The user then lifts up the stick and provides her response. In the *Similarity* experiment the user indicates the location of the perceived CoM (green indicator) using the HTC Vive Controller in the second hand. In the *Colocation* experiment, the user answers a yes/no question asking if the stick feels left-weighted. After the answer is recorded, the trial ends by putting the object back down as indicated by the green hologram. The weight inside the proxy relocates between trials. ©2021 IEEE.

Upon completion of the calibration, participants practiced the experiment in four training trials, before the data collection started. When training was completed, each participant performed five trials for each of the four conditions in a randomized order. To complete a trial, participants first aligned their virtual hand for 2s with a holographic hand displayed in VR, shown in Figure 12.6. This holographic hand was located at the same height as the virtual stick and 60cm in front of its geometric center.

After 2s, a sound signaled that the participant could now reach for the virtual stick and grasp it at its center to lift it up vertically. Trials employing HR would warp the virtual hand during this phase as the user's hand approached the proxy. Virtual stick and physical proxy were perfectly colocated during all trials in the *Similarity* experiment. A carefully designed grasping animation of the fingers was played when participants approached the stick to enhance immersion.

After lifting up the stick, participants indicated where they perceived the CoM of the virtual stick to be located by marking a location on the stick with a virtual

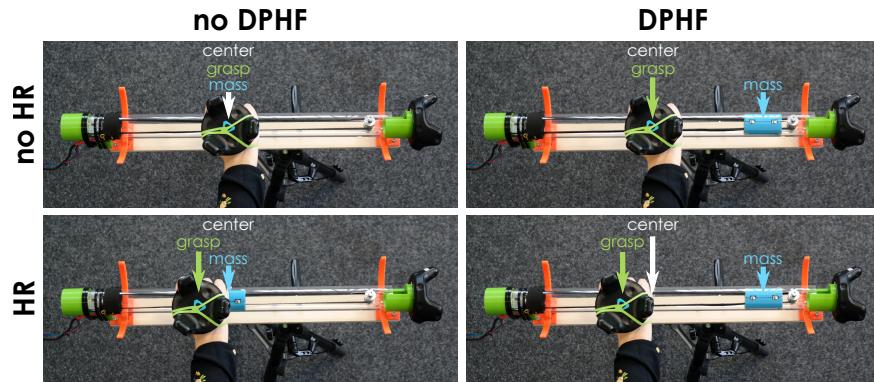


Figure 12.8: Real-world matrix view of our four conditions in the *Similarity* experiment. Each technique is shown rendering its respective maximum weight shift. © 2021 IEEE.

green indicator, as depicted in the center right VR view in Figure 12.7. To move the indicator on the stick, participants used the controller’s touchpad and they recorded their answer using the trigger button. When the perceived CoM was recorded, participants put back the stick by aligning it with a hologram displayed in VR and retracted their hand. This flow was then repeated for the next trial, and is illustrated in Figure 12.7.

Upon completion of all 20 *Similarity* trials, participants filled out a SUS presence questionnaire [Slater et al., 1994] and a SSQ [Kennedy et al., 1993]. Moreover, they were asked if they noticed their virtual hand moving differently than their real hand at any moment during the experiment. Finally, they completed additional demographic data. Upon leaving VR, participants were free to provide any further comments in a written form.

Design The *Similarity* experiment has a within-subjects design. The independent variable is the haptic rendering technique with its four tested implementations (*BL*, *DPHF*, *HR*, *DPHF+HR*). Each technique was configured to render its maximum achievable lever effect towards the right end of the stick as introduced in subsection 12.2.2 and sketched in Figure 12.2. Figure 12.8 depicts a detailed real-world view of the conditions, indicating the geometric center of the stick, the grasp location of the hand, and the location of the movable mass.

As a dependent variable, we assessed the *Similarity* metric, i.e., the perceived location of the virtual stick’s CoM in the continuous range $[-1, +1]$, where -1 represents the left end of the stick, 0 its center, and $+1$ the right end of the stick. Each participant completed four trials for training, and 20 trials with data recording (five times each of the four conditions in a random order). This results in $N \cdot 20 = 24 \cdot 20 = 480$ estimates of CoM locations (120 per condition) in total.

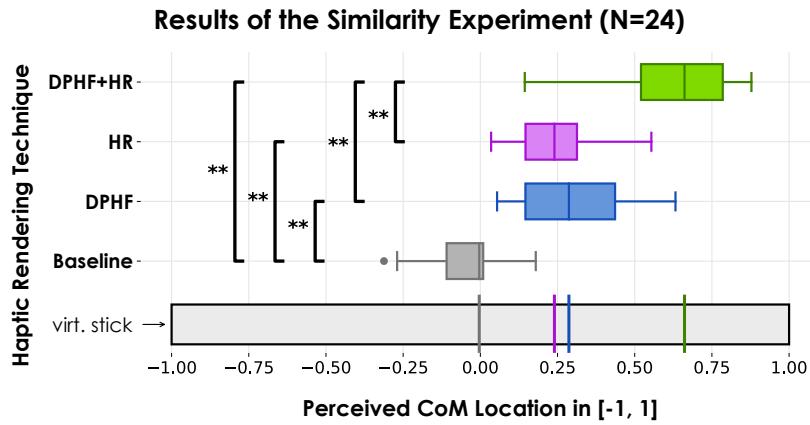


Figure 12.9: Perceived virtual CoM location along the stick in the *Similarity* experiment. Range from -1 (= left end) to $+1$ (= right end) with 0 indicating the geometric center of the stick. Brackets indicate statistically significant differences ($p' < .05$ (*); $p' < .01$ (**)). © 2021 IEEE.

Results It took participants on average $5.00s$ ($SD = .98s$) to reach for the virtual stick and lift it up, and $7.12s$ ($SD = 3.48s$) to indicate the perceived CoM location. 18 participants (75%) indicated that they did not notice the hand redirection in any trial. The main results of the *Similarity* experiment are summarized in Figure 12.9. To investigate **H1** and applying a significance level of $\alpha = .05$, we compared the perceived CoM locations of the four conditions using statistical tests. As, according to a Shapiro-Wilk test, normality could not be assumed for all conditions, we ran a non-parametric Friedman test indicating the average perceived CoM location to differ significantly across conditions ($\chi^2(3) = 66.05, p < .001$). To find pairwise differences, we performed post-hoc Wilcoxon signed-rank tests (p' indicating Bonferroni-corrected p values). Except for the comparison of DPHF and HR, the average perceived CoM locations of all remaining pairs were found to be significantly different (all $p' < .001$).

Most importantly concerning **H1**, we found the CoM shift perceived in the *BL* condition ($M = -0.054, SD = 0.122$) to be significantly smaller than the shifts perceived with DPHF ($M = 0.295, SD = 0.171$) ($Z = -4.286, p < .001, r = .62$) and HR ($M = 0.247, SD = 0.147$) ($Z = -4.286, p < .001, r = .62$). The CoM shift perceived when the combined technique of *DPHF+HR* was applied ($M = 0.611, SD = 0.217$) was found to be significantly larger than the shifts perceived when applying only DPHF ($Z = -4.286, p < .001, r = .62$) or HR ($Z = -4.286, p < .001, r = .62$). As expected, the shift ranges of DPHF and HR were not found to be significantly different ($p' = .193$).

Additionally, we computed the perceived virtual CoM shifts in centimeters by mapping the range of $[-1, 1]$ to the length of the virtual stick of $59.6cm$. We then compared the perceived levers in VR to the physical levers achieved in the rendering techniques using one-sample Wilcoxon tests and a Bonferroni-Holm correction. While the perceived shift in *BL* ($M = -1.60cm, SD = 3.64cm$) was

not found to be significantly different from 0cm ($p' = .101$), the tests revealed that users significantly overestimated weight shifts in VR (given by *lever*):

$$\begin{aligned} \text{lever}(HR) &= 7.37\text{cm} \stackrel{p'<.028}{>} 4.72\text{cm} = |\text{HR}_{\max}| \\ \text{lever}(DPHF) &= 8.78\text{cm} \stackrel{p'<.012}{>} 5.00\text{cm} = |\text{DPHF}_{\max}| \\ \text{lever}(DPHF+HR) &= 18.20\text{cm} \stackrel{p'<.004}{>} 9.72\text{cm} = |\text{HR}_{\max}| + |\text{DPHF}_{\max}| \end{aligned}$$

Finally, the results of the presence questionnaire for the SUS count ($M = 2.21$, $SD = 1.86$) (with answers ranging from 0 to 6), and SUS mean ($M = 4.72$, $SD = 1.22$) confirmed the IVE to be sufficiently immersive. Participants did not report any sickness issues, as is supported by the obtained SSQ total scores ($M = 26.96$, $SD = 19.01$).

Colocation Experiment

Based on the second thought experiment, the *Colocation* experiment studies **H2**, i.e., which of the techniques compensates best for unwanted CoM shifts when the virtual object is spatially offset from its proxy. To investigate this, we utilize a psychophysical method to determine for each technique the maximum displacement of proxy and virtual object that can go unnoticed. Following our *Colocation* metric, this measure is obtained by deriving for each technique the corresponding DT for real-to-virtual offset.

Procedure Just like in the *Similarity* experiment, participants started by completing the hand calibration procedure and a set of training trials before the data collection started. From the perspective of participants, the trials in the *Colocation* experiment were equivalent to those in the *Similarity* experiment, except for the question. In the *Colocation* experiment, participants were asked in each trial whether the virtual stick felt *left-weighted*, i.e., heavier on the left side, or not. Participants did not need to indicate the perceived CoM location as in the *Similarity* trials, but only answered the yes/no question using the controller as shown in the lower right of Figure 12.7.

A second difference from the *Similarity* experiment is that in the *Colocation* experiment, the proxy and the virtual stick were not always perfectly colocated, but the virtual stick's center was offset in each trial by a *stimulus* $\in [0, 20\text{cm}]$ towards the right from the physical stick's center. To control the stimulus across trials, we employed a staircase procedure as in chapter 9, configured as outlined in the following paragraph. While users could only see the virtual stick and were instructed to grasp it at its center, the haptic rendering techniques were used to compensate for unwanted physical lever effects caused by these dislocations.

Figure 12.10 depicts a detailed overview of how the techniques compensate for various amounts of real-to-virtual offsets. The illustration indicates the grasp

location of the virtual hand at the center of the displaced virtual stick with a blue arrow. The physical CoM location is shown in orange, and the real grasp location in yellow and observable from the hand in the pictures. After completing a staircase procedure for each of the four tested conditions, participants were asked to fill out the same questionnaires as in the *Similarity* experiment.

Design The *Colocation* experiment has a within-subjects design, with the independent variable being the technique employed to compensate for unwanted weight shifts caused by the dislocation of the virtual and physical stick (*BL*, *DPHF*, *HR*, *DPHF+HR*).

An estimation of the maximum unnoticeable offset of the virtual stick and its proxy along the stick's main axis serves as the dependent variable. This offset-related DT represents the *Colocation* metric (*offset*) and was determined by leveraging an adaptive method from the field of psychophysics, specifically, an interleaved staircase procedure [Kingdom and Prins, 2016c; p. 53] implemented with our *Unity Staircase Procedure Toolkit*⁷². Formally, we applied a 1 up/1 down method with a fixed step size and two interleaved sequences ($\Delta^+ = \Delta^- = 2\text{cm}$; min. stimulus of 0cm ; max. stimulus of 20cm) [Kingdom and Prins, 2016a; pp. 120 ff.].

The staircase procedure varied the offset of the virtual stick's center from the center of the proxy (i.e., the *stimulus*) across trials. It operated with an ascending sequence (plotted in red in Figure 12.11) starting at no offset, and a descending sequence (plotted in blue in Figure 12.11) starting at the maximum offset of 20cm . Each trial was randomly assigned to one of the sequences and when a stimulus was noticed in a trial (i.e., the participant answered that s/he noticed that the virtual stick felt imbalanced), the offset for the following trial was decreased by the fixed step size of 2cm . If, in a trial, the participant failed to notice a stimulus (i.e., the haptic technique was successful in compensating for any perceivable lever effect, and the virtual stick was perceived as balanced), the offset in the following trial of that sequence was increased by 2cm . If in one trial of a sequence the participant noticed an imbalance, while in the preceding trial s/he did not (or vice versa), the sequence logged a reversal. Each sequence was progressed until five reversals had occurred, out of which the last four reversals were averaged to estimate a threshold for the respective sequence. Finally, the average of the ascending and descending sequence thresholds was taken as a general offset DT estimate for each participant. The DT indicates how much dislocation can go unnoticed with the tested technique being used to compensate for undesired weight shifts.

Each participant completed an interleaved staircase procedure for each of the four tested conditions. A 4×4 Williams design Latin square [Williams, 1949] was used to counterbalance the order of conditions across participants.

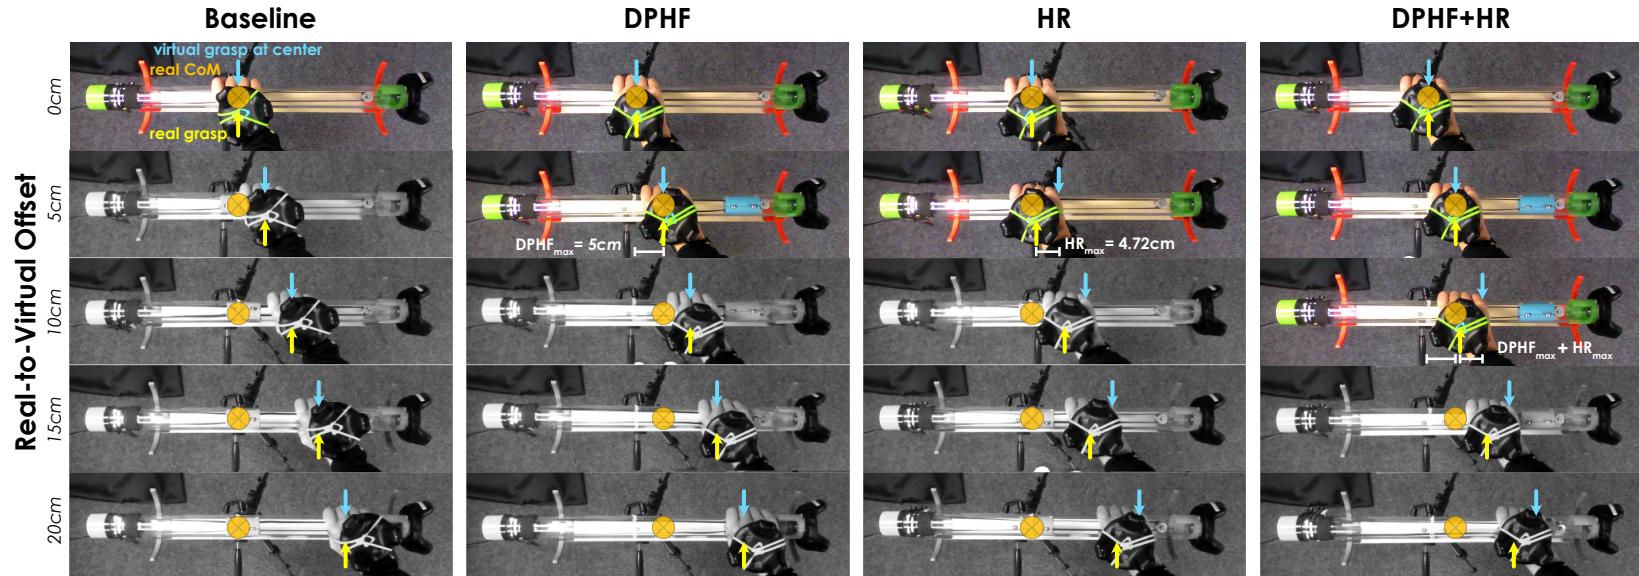


Figure 12.10: The four techniques to compensate for unwanted weight shifts as a result of the dislocation of virtual stick and proxy. Blue arrows indicate the center of the dislocated virtual stick (i.e., the virtual grasp location), yellow arrows show the grasp location of the real hand, and orange markers illustrate the proxy's CoM for five representative offsets. States without a perceivable *lever* (given by the distance from yellow to orange marker) are colored; imbalanced states are displayed in black and white. © 2021 IEEE.

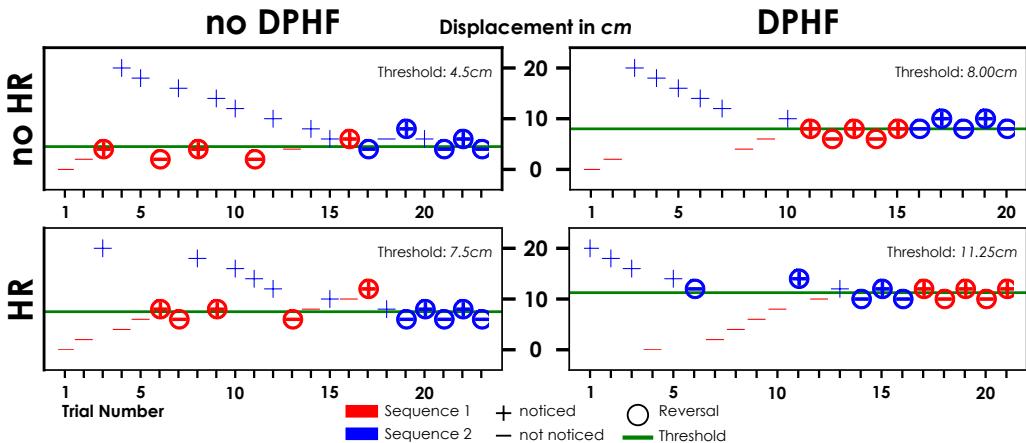


Figure 12.11: Interleaved staircase results of participant #16 in the *Colocation* experiment. The *stimulus* on the y-axis represents the offset of the virtual stick’s center from the center of the proxy in cm. The DTs that can be observed are in line with **H2**. © 2021 IEEE.

Results It took participants in the *Colocation* experiment on average 5.28s ($SD = 3.91$) to reach for the virtual stick and lift it up. The yes/no question was answered on average within 2.94s ($SD = 1.39$). 21 participants (87.5%) failed to notice the hand redirection during the experiment. Figure 12.12 summarizes the central results of the *Colocation* experiment. To arrive at these threshold estimates, which indicate the maximum offset of proxy and virtual stick that the techniques can compensate for in meters, a total of 2218 trials was conducted across all participants. On average, each participant completed 92.42 trials ($SD = 6.48$, $max. = 109$, $min. = 84$).

To investigate **H2**, we applied a significance level of $\alpha = .05$ and compared the DTs of the four techniques. A Shapiro-Wilk test indicated that normality could not be assumed for each condition and a Friedman test signaled the thresholds to vary significantly across techniques ($\chi^2(3) = 61.08, p < .001$). Post-hoc analysis with Wilcoxon signed-rank tests and Bonferroni correction revealed all pairwise comparisons of thresholds to be significantly different (all $p' < .001$), except for the comparison of DPHF and HR. Most meaningful regarding **H2**, the offset DT when no compensation for undesired weight shifts is applied in *BL* ($M = 3.95\text{cm}$, $SD = 2.19\text{cm}$) was found to be significantly smaller than when DPHF ($M = 7.13\text{cm}$, $SD = 2.95\text{cm}$) ($Z = -4.286, p < .001, r = .62$) or HR ($M = 6.74\text{cm}$, $SD = 2.55\text{cm}$) ($Z = -4.199, p < .001, r = .61$) were used for compensation. Moreover, the combined technique *DPHF+HR* was found to compensate for the most spatial offset, with DTs ($M = 10.13\text{cm}$, $SD = 2.99\text{cm}$) found to be significantly larger than when using only the individual techniques of DPHF ($Z = -4.076, p < .001, r = .59$) or HR ($Z = -4.286, p < .001, r = .62$). As expected, the thresholds of DPHF and HR were not found to differ significantly ($p' > .99$).

Interpreting these values, the *BL* threshold can be viewed as the users’ general “tolerance” for CoM offsets. Based on this consideration, we computed, for



Figure 12.12: Offset DT in meters as estimated by the staircase procedure in the *Colocation* experiment. It is the maximum offset of virtual stick and proxy that goes unnoticed. Brackets indicate statistically significant differences ($p' < .05$ (*); $p' < .01$ (**)). © 2021 IEEE.

each rendering technique, its actual contribution to the increased tolerance for real-to-virtual offsets by subtracting the *BL* threshold (participant-wise). We then compared the contribution of *DPHF+HR* ($M = 6.19\text{cm}$, $SD = 1.69\text{cm}$) to the sum of the contributions of the individual techniques *DPHF* and *HR* ($M = 5.98\text{cm}$, $SD = 1.83\text{cm}$). A Wilcoxon signed-rank test did not find the difference ($\Delta = 0.21\text{cm}$) to be statistically significant ($p = .853$).

Just as in the *Similarity* experiment, the SUS count ($M = 2.37$, $SD = 1.95$) (with responses from min. 0 to max. 6), SUS mean ($M = 4.79$, $SD = 1.27$), and SSQ total scores ($M = 34.75$, $SD = 25.18$) confirmed the IVE to be generally immersive and free of sickness issues.

12.4 Discussion of *Combined DPHF and HR*

Looking at both the theoretical and the empirical results of our thought and user experiments, we found that *Combined DPHF and HR* can indeed be superior to the individual techniques when it comes to solving the *Similarity* and *Colocation* challenges in the context of weight-shifting proxy-based haptics.

12.4.1 Enhancing *Similarity*

The conclusions of our *Similarity* experiment are in line with the hypotheses derived from our thought experiments. A comparison of the shifts conveyed with the three rendering techniques to the baseline perception verifies the general effectiveness of *DPHF*, *HR*, and *DPHF+HR* in rendering virtual mass distribution. All three conditions generated significantly stronger shift perceptions than *BL*. Moreover, the shifts achieved with the combined technique *DPHF+HR* were sig-

nificantly greater than those conveyed with DPHF and HR alone. The lever effect perceived with *DPHF+HR* ($M = 18.20\text{cm}$) even exceeded the sum of those perceived with DPHF and HR ($M = 16.16\text{cm}$) – although not significantly, according to a Wilcoxon signed-rank test ($p = .11$).

In general, our findings verify (1) that we could successfully combine both techniques in practice to improve the haptic rendering capabilities of a proxy-based VR system, and (2) that the benefits of combining DPHF and HR are perceivable with state-of-the-art implementations in use (i.e., a weight-shifting proxy like *Shifty* and body warping based on Cheng et al. [2017b]). Based on our empirical findings, we accept **H1** and provide a first proof of the capacity of combining an adaptive proxy and HR for conveying virtual weight distribution.

As an additional finding, our results also indicate that users tend to overestimate CoM shifts in VR. The perceived virtual shifts in the *Similarity* experiment were significantly greater than what we expected given the physical levers produced. The overestimation affected all rendering techniques, with similar ratios of *perceived virtual* to *rendered physical lever* (1.56 for HR, 1.76 for DPHF, 1.87 for *DPHF+HR*) – a finding worth exploring in future studies. From a practical perspective, this indicates that techniques that are only capable of producing small ranges of physical weight shifts can still convey larger effects in VR, which further increases their versatility. Moreover, the tendency to overestimate virtual shifts highlights the utility of our combined approach, as every extension of the rendering range results in an even larger extension of the range of virtual shifts conveyed. The way in which this overestimation is linked to the general performance of humans in estimating CoMs in the real world remains to be explored in future work.

Our findings transfer directly to applications that employ proxy-based haptics and render virtual mass distribution. In such applications, our combined technique allows a single proxy to convincingly represent a greater variety of virtual objects – e.g., a balanced virtual stick, a shovel, a sword, or a hammer – while maintaining haptic *Similarity*.

12.4.2 Enhancing Colocation

The *Colocation* results complement our findings regarding *Similarity*, and also align with the hypotheses derived in our thought experiments. In this chapter, we put forward how the haptic rendering of weight shifts can help solve *Colocation* issues in proxy-based VR. For this, we used the rendering techniques to cancel out undesired lever effects, which can be a result of spatial mismatches of virtual object and proxy. The offset DTs obtained in the *BL* condition represent the general tolerance for misaligned real and virtual CoMs, and indicate how much the proxy in our scenario can be dislocated before the misalignment is noticed. As highlighted by the significantly increased thresholds achieved with DPHF, HR, and *DPHF+HR*, compensation for such CoM offsets yields a larger

range of unnoticeable proxy dislocation and hence better solves the *Colocation* challenge. Additionally, we found *DPHF+HR* to achieve a range of unnoticeable dislocation that significantly exceeds the ranges obtainable with just DPHF or unnoticeable HR. The additional tolerance to proxy dislocations contributed by *DPHF+HR* (taking into account the general tolerance assessed in condition *BL*) closely matched the sum of the individual contributions of DPHF and HR, which coincides with the theoretical model outlined in the thought experiments and demonstrates the practical realizability. Based on this evidence we also accept **H2**. Moreover, our observation that 75% and 87.5% of participants in the *Similarity* and *Colocation* experiments, respectively, failed to notice the applied body warping further supports the DTs derived in chapters 8 and 9.

It is worth noting at this point again that the achieved improvements in terms of *Similarity* and *Colocation* came at no additional hardware costs compared to a DPHF solution. The only expense to improve the proxy-based haptic rendering is the added complexity introduced with hand redirection and the orchestration of both techniques during interaction. In practical implementations, our combined technique can be utilized to let a single dynamic prop represent a set of virtual objects at different locations inside the IVE. As long as they are within the unnoticeability range, users can remain unaware that they are all represented by the same proxy – solving the *Colocation* challenge for such cases.

12.4.3 Beyond the Proof-of-Concept Scenario

Our proof-of-concept scenario was chosen to propose the idea of combining DPHF and HR for enhancing kinesthetic proxy-based haptic feedback based on weight shift. The physical laws of mass distribution make this scenario perfectly suitable to study and showcase this approach. Moreover, the rendering of virtual mass-related properties can be considered to be of general importance when it comes to creating believable VR experiences, as can be seen from the increasing research interest in rendering virtual mass distribution [Fujinawa et al., 2017; Shigeyama et al., 2019; Yu and Bowman, 2020]. However, we believe DPHF and HR can also be combined when rendering other effects.

For example, DPHF proxies that can change their surface textures by relocating surface samples (like the *Haptic Revolver* by Whitmire et al. [2018], or the *Haptic Palette* by Degraen et al. [2020]) could benefit from simultaneous hand redirection to increase the number of samples that can be explored. In other scenarios, such as when leveraging air drag to render resistance effects in VR (as implemented by the *Drag:on* in chapter 6), gain-based HR could modify the user’s hand movement speed to intensify the perceived air resistance. Altering hand movement speeds while interacting with a controller such as the *ElastOscillation* by Tsai et al. [2020] could likewise yield novel effects of elasticity. Moreover, DPHF props rendering different degrees of stiffness [Murray et al., 2018] could potentially benefit from a combination with HR too, e.g., by redirecting the hand as the user explores the prop’s surface, or by applying gain factors when the proxy is squeezed.

Furthermore, our experiments only employed conservative HR within worst-case unnoticeability ranges (i.e., the CDTs derived in chapter 8). Yet, future implementations might take advantage of greater redirection amplitudes, which might go unnoticed in less conservative application scenarios, or lie within less restrictive tolerance thresholds [Cheng et al., 2017b].

The exact way in which DPHF and HR can best be combined in order to enhance proxy-based haptic rendering in VR depends on the physical laws underlying the targeted haptic effects. For the rendering of weight distribution, these laws are straightforward and primarily concern increasing or decreasing the lever distance or inertia. For other haptic dimensions, we would like to motivate future research to find effective approaches combining dynamic proxies and hand redirection.

12.4.4 Limitations

Our investigation showed the value of *DPHF+HR* in the context of our proof-of-concept scenario and in terms of the *Similarity* and *Colocation* metrics defined in subsection 12.3.1. The benefits reported in our experiment, however, were linked to certain assumptions. For example, we only investigated the interaction with a symmetric tubular object and we assumed the dislocation of the prop to only take place along the prop's main axis. Additionally, we assumed the proxy and the virtual object to be in the same orientation in order to not increase the complexity of our investigation.

Furthermore, a general limitation of our work presented in this chapter is that we only investigated a single haptic dimension, i.e., weight distribution, represented by a single implementation, i.e., the device and algorithms utilized. As a consequence, the metrics, principles, and algorithms applied here, while likely generalizable to haptic effects that have similar physical laws involved, probably do not work as-is for every imaginable haptic dimension or prop form factor. With our work, however, we contributed an important step towards a rigorous scientific analysis of the value of combining DPHF and HR and thus regard these limitations as a motivation for future research – especially in light of the promising findings uncovered in the investigated proof-of-concept scenario.

12.5 Conclusion & Contribution to the Research Questions

This chapter concludes the research in this thesis by embracing both of the previously investigated **Improvements** of proxy-based haptics at the same time. In particular, we contributed by investigating the combination of the physical technique of DPHF proposed in Part III, and the virtual technique of HR studied in Part IV. We found that for most haptic dimensions, both concepts have, up to now, not ever been studied in combination – with the exception of Ban et al.'s work [2014] on shape *Similarity*. Motivated by this, we continued to close this

fundamental gap. For this, we set out to answer **RQ 4** by showing how to improve the flexibility of proxy-based haptics with the first VR system that, both in theory and in practice, combines the two techniques for rendering kinesthetic effects, specifically, virtual weight distribution. Additionally, we consider for the first time both the challenges of *Similarity* and *Colocation* in the same context.

To do so, we first defined metrics for the two criteria of *Similarity* and *Colocation* with respect to the domain of weight-shifting proxies addressed in chapter 5. These metrics provide the basis for our comparison of the different approaches and consider their capacity of rendering mass distribution inside a virtual stick. We argue, based on these metrics and two thought experiments, that *Combined DPHF and HR* can provide feedback superior to that provided by the individual techniques.

Key to the success of the combined technique thereby is the complementary nature of physical weight-shifting and virtual HR. The compatibility of software-based visual-haptic illusions with hardware-based techniques enabled us to combine translations of the weight inside the proxy (DPHF in one direction) with redirections of the hand along the proxy (HR in the opposite direction) to control the perceived physical lever – a concept potentially transferable also to other domains such as shape, texture, or compliance rendering. As a result, our theoretical considerations in two thought experiments led us to hypothesize that when combining DPHF and HR to render weight shifts, the combined technique will provide significantly greater perceived shifts (**H1**). Moreover, we expected *DPHF+HR* to allow for significantly greater displacements between virtual stick and physical proxy to go unnoticed thanks to improved compensation of CoM mismatches (**H2**).

We add to **RQ 4** as we empirically show that the hypothesized effects can be achieved not only in theory but also in practice. For this, we utilized our proposed implementation based on the weight-shifting proxy *Shifty* introduced in chapter 5 and the body warping approach by Cheng et al. [2017b] applied within the thresholds derived in chapter 8. The statistical results of our perceptual and psychophysical experiments led us to accept both **H1** and **H2**, and we additionally obtained evidence that users significantly overestimate weight shifts in VR – an interesting aspect to be investigated in future studies. Furthermore, while we proposed concrete concepts and algorithms to improve the proxy-based rendering of weight distribution inside a virtual stick with *Combined DPHF and HR*, we further added to **RQ 4** by discussing starting points for future research. In this context, we outlined ideas on how the proxy-based rendering of various haptic qualities might benefit from a combination with hand redirection.

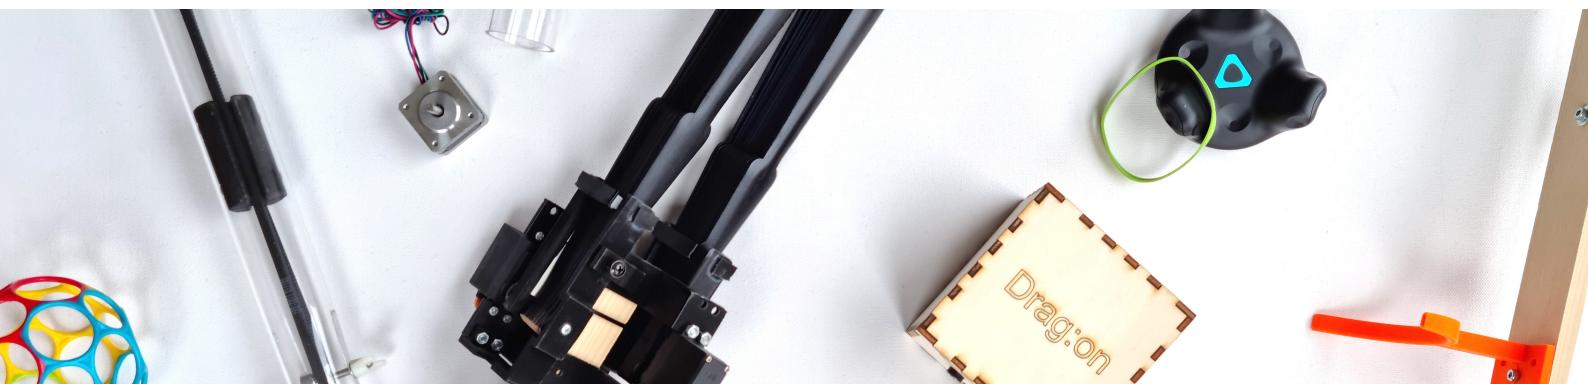
The promising findings of this chapter bear potential for lasting impact in the field of proxy-based haptics as combined real-and-virtual techniques might further pave the way towards the vision of an ultimate display. Combined techniques promise to get the most out of the particularities of human perception while taking advantage of the capabilities of physical approaches. In this context,

DPHF promises to add adaptive, multimodal physical feedback at comparably low costs to the equation.

While not all combinations of DPHF and hand redirection will yield benefits for proxy-based rendering, specific combinations can do, as our results validate. Yet, to reach the vision of an ultimate display, an important next step would be to further generalize the results achieved in this chapter. To benefit from combined techniques when rendering other haptic properties, future research needs to identify requirements and generally applicable algorithms. Potentially, these (or some of them) might be based on the principles introduced in this chapter for combining DPHF and HR to render weight shifts, namely: (1) employing both techniques in “opposite directions” to “maximize” effects as in our *Similarity* experiment, and (2) applying them in the “same direction” to “minimize” unwanted perceptions as in our *Colocation* experiment.

Part VI

General Conclusions



Chapter 13

Summary and Contributions

We conclude this thesis by summarizing our work and discussing our achievements in the context of the four research questions. For this, we first briefly recapitulate our efforts and then outline for each research question our main scientific contributions.

13.1 Summary

The goal of this thesis was to contribute to the field of VR haptics by pushing the boundaries towards Sutherland's 1965 vision of the ultimate display – a VR system capable of simulating arbitrary IVEs in a manner that is indistinguishable from reality and in which "*the computer can control the existence of matter*" [Sutherland, 1965]. As outlined in the introduction of this thesis, this vision of the ultimate display can be approached from two different angles when it comes to realizing haptic feedback: either by building on the idea of AHF or the concept of PHF.

With AHF being found to suffer from inherent drawbacks, such as a severely limited ability to provide multimodal haptic feedback, the need for costly and complex actuators, computationally expensive simulations, and severe workspace limitations, in this thesis, we concentrated on the approach of PHF. Yet, the naïve approach of using passive physical props to provide haptic feedback for arbitrary IVEs is struggling with its own severe limitations too. These, however, and in contrast to the challenges faced by AHF, promise to be solvable not primarily through advancements in engineering. Instead, solutions to the problems of proxy-based haptics seem achievable by mixing low-complexity active and passive concepts, by perception-inspired illusion techniques, as well as by the

combination of both approaches – making proxy-based haptics appear as the path of least resistance towards the realization of an ultimate display.

Motivated by this observation and our in-depth review of previous work on VR, human perception, and the different approaches to haptics in VR in chapter 2, we argue in Part I that PHF already achieves many of the qualities desirable for the ultimate display. In particular, proxy-based feedback can provide high-quality, multimodal, and realistic tactile and kinesthetic impressions of virtual objects inside an IVE. Yet, what still distances naïve PHF from Sutherland’s vision of the ultimate display is the *lack of control* the VR system has over the haptic impressions perceived by the user at runtime. It is with the research in this thesis that we aim to increase the flexibility of proxy-based haptics in order to grant the VR system *more control* over the haptic perceptions of users while maintaining the many advantages of haptic proxies. In the context of this endeavor, this thesis pursued two main goals: applying and improving proxy-based haptics.

13.1.1 Applying Proxy-Based Haptic Feedback

Firstly, we aimed to underline in Part II that proxy-based haptics bears great potential to lead towards the realization of an ultimate display. We did so by showcasing the **Applicability (RQ 1)** of haptic proxies in a novel application domain that has not ever been supported by proxy-based VR haptics before, namely, the domain of business process modeling.

Leveraging the large design freedom in the field of abstract data exploration, we developed the *Immersive Process Model Exploration* system in chapter 3. The system enables users to immersively explore business process models while tangibly transporting bits of information through a more-than-room-scale 3D IVE that represents the process graph. Motivated by a gamification component, users of the system experience internal dependencies of processes through multisensory feedback, including basic interactions with low-fidelity passive proxy objects. Our system thereby renders arbitrary process models haptically with only a constant number of four proxies, showcasing how design freedom can be leveraged to reuse props for larger virtual environments (e.g., by taking advantage of symmetric IVE and proxy layouts in combination with a resetting controller). In a user study, we compared our immersive interface with PHF to a version with vibrotactile AHF provided by commercial VR controllers, and to a traditional 2D interface on a tablet. Our results highlight a tradeoff between exploration efficiency and user interest with the tablet interface communicating processes the quickest and the proxy-based interface leading to greatest user interest. The findings of our work in chapter 3 thus lead us to conclude that proxy-based haptics is well-applicable in the domain of abstract data exploration.

13.1.2 Improving Proxy-Based Haptic Feedback

After we successfully applied conventional PHF in a novel application domain in Part II, we set course towards an **Improvement** of proxy-based haptic feedback. To contribute the next evolutionary steps towards the vision of the ultimate display, we followed two separate research paths in Parts III and IV, which we eventually joined in Part V. Our research was thereby guided by the two main obstacles for successful proxy-based haptic feedback introduced in chapter 2: the challenges of *Similarity* (focus of Part III) and *Colocation* (focus of Part IV). These challenges postulate that proxies need to deliver convincing haptic sensations during interaction (*Similarity*) while being perceived as spatially colocated with their virtual counterparts (*Colocation*). As such, the challenges serve as a compass on the way towards an ultimate display since a perfect solution of both challenges at the same time can be regarded the ultimate haptic feedback and would, theoretically, render the simulation indistinguishable from reality.

To push the scientific boundaries concerning the tackling of both challenges, our work on the **Improvement** of proxy-based haptics focused on two fundamentally different concepts. While our research in Part III concentrated on real-time physical (i.e., hardware-based) techniques, our work in Part IV was concerned with real-time virtual (i.e., software-based) approaches. To advance both research paths in a meaningful manner, we derived for each of the two concepts its own individual research question.

The Physical Approach Our work in Part III was guided by **RQ 2** and set out to fill the conceptual gap on the passive side of the Active-Passive Haptics continuum outlined in Figure 1.3. By this, **RQ 2** aimed to bring a new physical approach for system-*controlled* proxy-based haptic feedback to the table.

Our solution to **RQ 2** is the concept of *Dynamic Passive Haptic Feedback (DPHF)* introduced in chapter 4. We propose to combine haptic proxies with simple actuation mechanisms that add the capability of adaptation to the props. DPHF distinguishes itself from the concepts of AHF and ETHF in that it trades off actuation complexity and the degree of proxy usage differently. In particular, DPHF's tradeoff is much more in favor of simple mechanisms and the reliance on passive haptic impressions that still originate from the proxy's haptic qualities. DPHF props use their built-in actuators not to actively exert forces on the user, but instead only to vary their passive haptic qualities so as to align them with different virtual objects. By this, DPHF props are more flexible versions of PHF props and achieve *Similarity* with a greater variety of virtual counterparts. Just as PHF props, they provide high-quality multimodal haptic feedback while maintaining a low mechanical and computational complexity compared to ETHF. In contrast to PHF, however, DPHF grants VR systems control over the haptic perceptions of users at runtime, enabling them to command props to change their physical state during the simulation. These perceptions can further be controlled through visual-haptic illusions and the exploitation of visual dominance, as we

could demonstrate in the remainder of Part III and in Part V.

Following its introduction, in chapters 5 and 6, we then proved the practical realizability of the concept of DPHF in order to provide an answer to the remainder of **RQ 2**. Our goal in these chapters was to enhance the capabilities of haptic props to convey different kinesthetic qualities of virtual objects, such as their virtual weight, shape, material, and alike. To this end, we developed two novel haptic feedback concepts based on DPHF.

First, in chapter 5, we proposed to take advantage of changes in the proxy's *weight distribution* in order to convey the feel of different virtual objects. Specifically, we proposed to vary the proxy's inertia. To validate this approach, we developed the weight-shifting VR controller *Shifty* and evaluated it in two user experiments. In these, we could show that, thanks to dynamic touch [Turvey, 1996], adaptable proxy inertia in combination with appropriate visualizations can enhance the perception of virtual objects that differ in shape (e.g., length or thickness) or weight. Moreover, in a direct comparison with a PHF prop, we found the DPHF of *Shifty* to increase perceived realism and enjoyment significantly.

Following up on our achievements with *Shifty*, in chapter 6, we then introduced a second novel feedback mechanism based on DPHF designed to render kinesthetic impressions in VR. Here, we proposed to use shape-changing proxies that vary their *drag*, i.e., their air resistance, in conjunction with their inertia. By this, we aimed to exploit the controller movements that naturally occur during VR interactions for generating different haptic impressions of resistance. We implemented this concept and developed the shape-changing VR controller *Drag:on*. *Drag:on* utilizes two actuated foldable surface elements to increase and decrease its surface area either symmetrically or asymmetrically. The results of a user experiment showed that *Drag:on* can provide distinguishable levels of DPHF and increases haptic realism compared to the vibrotactile AHF of a commercial VR controller. In five different interactive scenarios tested in our user study, *Drag:on* could successfully convey different virtual mechanical resistances, virtual gas streams, and virtual objects differing in scale, material and fill state.

We concluded Part III with a discussion in chapter 7, which provides insights on how our two proposed DPHF concepts of *Shifty* and *Drag:on* theoretically complement each other. When considered in the bigger picture, the approaches of DPHF based on changes in weight shift and drag cover the entire object movement cycle. We conclude that, if appropriately combined, both techniques contributed in Part III might be capable of seamlessly rendering kinesthetic feedback during object holding, acceleration, and movement as outlined in Figure 7.1.

The Virtual Approach Conceptually orthogonal to Part III, our work in Part IV was concerned with **RQ 3**. Here, our aim was to advance proxy-based haptics by contributing to the research on Body Warping-based Hand Redirection (HR), a fundamental virtual technique, with which a VR system can *control* the user's real hand movement when interacting with haptic props. Our research on HR

was driven by three primary motivators: (1) the lack of worst-case DTs for HR techniques that employ gradually growing offsets and redirect the user's hand in mid-air while rendering a realistic hand visualization, (2) the lack of HR techniques that intentionally take advantage of change blindness during eye blinks, and (3) the lack of accessible software tools for HR.

Motivated by the first gap in research, in chapter 8, we investigated the detectability of desktop-scale HR in a psychophysical experiment. Here, we derived conservative lower-bound estimates for the DTs of *Continuous Hand Redirection*. In this way, we contributed to **RQ 3** by revealing the limitations that human perception implies for today's standard approach to HR. In the context of this study, we proposed algorithms for rotational and gain-based hand warping and tested the detectability of hand offsets in three scenarios that differed in the type and degree of user distraction. As a general result, we found unnoticeable angular HR to be possible within 4.5° horizontally and vertically, and gain factors between 0.88 and 1.07 to remain unnoticeable, even if users are primed and pay attention to detecting hand offsets. Our results thereby highlight that, even in worst-case scenarios, a certain margin exists within which the VR system can control the user's real hand trajectory without users even noticing it. This, in turn, can be used, for example, to solve *Colocation* issues through haptic retargeting [Azmandian et al., 2016b]. Yet, we also found this margin to be rather narrow. These findings form the basis of our research in chapter 9 and Part V, and are integrated into the software toolkit contributed in chapter 10.

Following up on our investigation of the state-of-the-art method to HR, we proposed *Blink-Suppressed Hand Redirection (BSHR)* in chapter 9. *BSHR* represents a novel approach to HR in that it combines the established strategy of continuously increasing hand offsets during reaching, with instantaneous hand shifts. The *BSHR* algorithm is an extension of the popular algorithm by Cheng et al. [2017b] and additionally utilizes eye tracking. That is, motivated by the second gap in research, we propose to apply instantaneous hand shifts during moments of visual suppression, in particular, during eye blinks, to take advantage of resulting change blindness. In order to study this novel approach and to further answer **RQ 3**, we conducted a psychophysical evaluation assessing the detectability of three variants of *BSHR* and the algorithm by Cheng et al. [2017b] as a baseline. While the first variant (pure *BSHR*) did not involve any continuous warping but only blink-suppressed shifts, the second and third variant combined continuous warping below the DTs derived in chapter 8 with instantaneous shifts during eye blinks. To implement our experiment, we developed the *Unity Staircase Procedure Toolkit*, which we released publicly in an open-source repository⁷² and reused in Part V. Our findings revealed pure *BSHR* to go unnoticed within margins of the same order of magnitude as those found for continuous warping in chapter 8. Yet, pure *BSHR* achieves this without manipulating the hand rendering in plain sight. Moreover, our experiment indicated that combining continuous warping below DTs and blink-suppressed shifts can further extend the range of unnoticeable redirection, compared to pure *BSHR*. Lastly, we quantified for the first time the

DTs of the common HR approach by Cheng et al. [2017b].

Our work on HR in Part IV was concluded by the introduction of the *Hand Redirection Toolkit (HaRT)* in chapter 10. Addressing the third motivator of this part of the thesis, we developed an open-source framework⁷⁴ for the Unity engine targeted at both novice and expert researchers and developers in the domain of VR. The toolkit was inspired by the RDW toolkit of Azmandian et al. [2016a] and provides support for implementing and evaluating hand redirection techniques. It is implemented as a unified platform that encompasses reference implementations of common hand redirection algorithms and a modular class hierarchy that allows adding novel algorithms. Moreover, it features simulation, logging, and visualization functionality, including threshold warnings, as well as an online documentation with explanations and videos. Besides popular algorithms proposed by related work, the toolkit also supports all HR algorithms proposed and investigated in this thesis. A qualitative expert study indicated the *HaRT* to be supportive and to provide a good UX.

The Combined Approach Having advanced research on real-time physical and virtual techniques with our work in Parts III and IV, we eventually entered a novel research domain by joining both approaches in Part V. To answer **RQ 4**, we investigated whether and how a combination of the physical approach of DPHF and the virtual technique of HR can better solve the challenges of *Similarity* and *Colocation* than the individual techniques alone can do.

For this purpose, we defined in chapter 12 the proof-of-concept scenario of rendering the weight distribution inside a virtual stick. This scenario was based on our research in chapters 5 and 8, so that we could draw from our previous results, in particular, by reusing our DPHF prop *Shifty* and the DTs derived for *Continuous Hand Redirection*. To objectively quantify the success of the individual techniques of DPHF, HR, and their combination, in tackling the challenges of *Similarity* and *Colocation*, we defined two respective metrics. These metrics were tailored to the scenario of investigation and the relative performance of the three different techniques was argued about in two thought experiments. From these thought experiments, we derived two central hypotheses, expecting the combined technique to outperform the individual approaches regarding both the maximum renderable weight shifts inside the virtual stick, as well as the maximum stick displacement that can be compensated for. Our expectations were ultimately confirmed by the results of two user experiments, one of which being a psychophysical study implemented with our *Unity Staircase Procedure Toolkit*. Our results highlight the great potential that combining real-time physical and virtual approaches bears for proxy-based haptic feedback. The scenario investigated in chapter 12 thereby points out that both the establishment of proxy *Similarity* and *Colocation* can profit from hybrid approaches – motivating further research along this path, at the end of which the ultimate display might be found.

13.2 Contributions

To highlight the advancements made throughout this thesis, we conclude with a summary of our main contributions. For this, we take a step back and view our work in the bigger picture. As introduced in chapter 1, our research was guided by four central research questions, each addressed in its dedicated part. The first question was concerned with studying the **Applicability** of proxy-based haptics, while the remaining three questions addressed the concept's **Improvement**.

To prepare our work, we started by conducting an extensive review of related research in chapter 2. This review, and especially the parts about proxy-based haptics covered in section 2.4 and the three sections following it, constitutes the first central contribution of this thesis. It represents a more comprehensive version of our review contributed in our recent publications [Nilsson et al., 2021a,b].

All other central contributions of this thesis originate from our work in Parts II, III, IV, and V, which dealt with **RQ 1**, **RQ 2**, **RQ 3**, and **RQ 4**, respectively. The major and minor contributions of these research endeavors were discussed in detail in the individual chapters, and we provided further discussion when closing Part III in chapter 7 and Part IV in chapter 11.

Research Question 1	Applicability
How can proxy-based haptic feedback support the domain of process model exploration?	
<i>Focus</i> Novel Application Domain	<i>Location</i> Part II (Chapter 3)
Theoretical Contributions	
Chapter 3	<ul style="list-style-type: none"> – Concept of a novel PHF-based VR system turning business process model exploration into a multisensory VR experience. – Comparison of the proposed VR interface to a conventional 2D interface revealing a tradeoff between process exploration efficiency and user interest. – Demonstration of the applicability of PHF and its value for the domain of process model exploration.
Technical Contributions	
Chapter 3	<ul style="list-style-type: none"> – Prototype of the <i>Immersive Process Model Exploration</i> system supporting arbitrarily large EPCs, incl. components for automatic IVE generation, logical process traversal, and haptic interactions (PHF and vibrotactile AHF).
Design Contributions	
Chapter 3	<ul style="list-style-type: none"> – Designs of IVEs, virtual objects, PHF proxies, and interactions suitable to communicate business process models in VR to non-expert users, and leveraging the great design freedom in the domain of abstract data visualization.

Major Contributions To complete previous discussions, four summary boxes in this section provide high-level overviews of the breadth of achievements of this thesis, detailing on our theoretical, technical, and design contributions.

The first box summarizes Part II, where we could demonstrate the applicability of proxy-based haptic feedback by successfully applying it for data exploration. Here, we could show the value of proxy-based haptics for a domain focused on learning and characterized by a large design freedom. Our results on **RQ 1** thus motivate the exploration of proxy-based haptics in further domains in the future.

Research Question 2	Improvement
<p>How can the gap in the haptics continuum be filled with a concept that enhances the flexibility of proxy-based haptic feedback and enables improved kinesthetic perceptions in VR with only minimal actuation?</p>	
<p><i>Focus</i> Physical Approach; Challenge of <i>Similarity</i></p>	<p><i>Location</i> Part III (Chapters 4, 5, 6)</p>
<p>Theoretical Contributions</p> <p>Chapter 4: – Novel, fundamental class of <i>Dynamic Passive Haptic Feedback (DPHF)</i> that complements the Active-Passive Haptics continuum for VR.</p> <p>Chapter 5: – DPHF-based approach to kinesthetic feedback that takes advantage of varying a proxy's inertia through weight shifting. – Insights on how adapting a proxy's inertia, in combination with different visualizations, can enhance the perception of virtual object shapes and weights. – Insights on how different audio-visual animations can help compensating for visual-haptic mismatch perceived during proxy transformation.</p> <p>Chapter 6: – DPHF-based approach to kinesthetic feedback that takes advantage of varying a proxy's drag through shape changing. – Insights on how adapting a proxy's drag and inertia, in combination with different visualizations, can enhance the perception of virtual resistances, gas streams, object scales, materials, and fill states.</p>	
<p>Technical Contributions</p> <p>Chapter 5: – Prototype of <i>Shift:y</i>, a weight-shifting DPHF-based VR controller, incl. an open-source repository, an experiment application for evaluating <i>Shift:y</i> with users, and a demonstration application encompassing three interactive scenarios.</p> <p>Chapter 6: – Prototype of <i>Drag:on</i>, a shape-changing DPHF-based VR controller, incl. an open-source repository, and an experiment application for evaluating <i>Drag:on</i> with users.</p>	
<p>Design Contributions</p> <p>Chapter 5 & Chapter 6 – Designs of physical mechanisms that provide ungrounded kinesthetic feedback while being realizable with low-cost materials, low-power actuators, as well as low-complexity mechanics and algorithms. – Designs of virtual objects, interactions, and animations that, when combined with the DPHF of <i>Shift:y</i> and <i>Drag:on</i>, elicit convincing visual-haptic illusions.</p>	

The second box reviews our answer to **RQ 2** in Part III, where we developed a novel class of haptic feedback for VR, mixing the ideas of PHF and AHF with a strong emphasis on proxy-based impressions. We argue that with *Dynamic Passive Haptic Feedback (DPHF)* we finally filled the theoretical gap between PHF and ETHF in the Active-Passive Haptics continuum.

Research Question 3	Improvement
<p>What limitations and potentials does human perception imply for the technique of Body Warping-based Hand Redirection (HR)?</p>	
<p>Focus Virtual Approach; Challenge of Colocation</p>	<p>Location Part IV (Chapters 8, 9, 10)</p>
<p>Theoretical Contributions</p> <p>Chapter 8: – Algorithms for rotational and gain-based <i>Continuous Hand Redirection</i>. – Quantification of how much HR in mid-air goes unnoticed, even in worst-case scenarios, when applying rot. or gain-based <i>Continuous Hand Redirection</i>.</p> <p>Chapter 9: – Algorithm for <i>Blink-Suppressed Hand Redirection (BSHR)</i> based on eye tracking, the first HR algorithm to take advantage of blink-induced change blindness. – Quantification of how much HR in mid-air goes unnoticed, even in worst-case scenarios, when applying (1) pure <i>BSHR</i>, (2) <i>BSHR</i> combined with subliminal continuous HR, and (3) the approach by Cheng et al. [2017b].</p> <p>Chapter 10: – Architecture of the <i>Hand Redirection Toolkit (HaRT)</i>, an open-source software framework for hand redirection, incl. a suitable class hierarchy and concepts for simulation, logging, and visualization features. – Insights on the opinions of expert Unity users about the <i>HaRT</i>.</p>	
<p>Technical Contributions</p> <p>Chapter 8: – Implementation of a psychophysical VR experiment to derive the perceptual DTs of <i>Continuous Hand Redirection</i> based on the method of constant stimuli.</p> <p>Chapter 9: – Implementation of a psychophysical VR experiment to derive the perceptual DTs of <i>BSHR</i> and the approach by Cheng et al. [2017b] based on the adaptive staircase method. – Implementation of the <i>Unity Staircase Procedure Toolkit</i>, an open-source software framework for psychophysical DT experiments based on the adaptive staircase method, incl. Unity support and a live plotting feature.</p> <p>Chapter 10: – Implementation of the <i>HaRT</i>, incl. six body warping and three world warping techniques, as well as an online documentation.</p>	
<p>Design Contributions</p> <p>Chapter 8: – Insights on the limits of human perception regarding the ability to detect visual-proprioceptive mismatches, which can help designing proxy objects, IVEs, and interactions.</p> <p>Chapter 9: – Insights on the potential of leveraging eye tracking and change blindness for covertly injecting visual-proprioceptive mismatches, which can help designing proxy objects, IVEs, and interactions.</p> <p>Chapter 10: – Designs of workflows and visualizations as part of the <i>HaRT</i>, which support the development of HR scenarios and are perceived as useful by expert users.</p>	

Moreover, in Part III, we proposed two novel approaches to kinesthetic VR haptics based on DPHF with *Shift:Y* and *Drag:On*. With both DPHF-based proxies, we demonstrated the great potential of visual-haptic illusions, which are based on adaptable weight distributions and shapes varying a proxy's inertia and drag as controlled by the VR system.

In the third box then, we compile the achievements of Part IV with regard to **RQ 3**. Here, we advanced the research on HR, which can improve a VR system's control over the conveyed proxy-based feedback by manipulating the real hand movement of users. We could answer some of the many questions that were left unanswered by previous research. Specifically, we could deliver insights on how, and to what extent, HR can be applied without users noticing that they are being manipulated. Our contributions were achieved through the proposal of novel algorithms and software tools, as well as through psychophysical investigations. By this, we did not only contribute to the domains of VR, haptics, and HCI, but also to the field of psychophysics.

Research Question 4		Improvement
<p>How can the physical approach of DPHF and the virtual approach of HR be combined to improve the flexibility of proxy-based haptic feedback?</p>		
<i>Focus</i>		<i>Location</i>
Combined Approach; Challenges of <i>Similarity</i> & <i>Colocation</i>		Part V (Chapter 12)
Theoretical Contributions		
Chapter 12	<ul style="list-style-type: none"> – Proposal of combining the physical technique of DPHF and the virtual technique of HR for achieving enhanced <i>Similarity</i> and <i>Colocation</i>. – Definition of perception-based metrics for <i>Similarity</i> and <i>Colocation</i> in the context of haptically rendering the weight distribution inside a virtual object. – Empirical validation that, and insights on how, the combination of a weight-shifting DPHF proxy (Part III) and unnoticeable HR (Part IV) can yield results for <i>Similarity</i> and <i>Colocation</i> that are superior to the individual techniques. 	
Technical Contributions		
Chapter 12	<ul style="list-style-type: none"> – Implementation of a perceptual VR experiment partly based on the adaptive staircase method, which combines <i>Shift:Y</i> (Part III) with unnoticeable HR (Part IV) to empirically investigate the performance of (1) DPHF, (2) HR, (3) combined <i>DPHF+HR</i>, and (4) naïve PHF as a baseline, in solving the challenges of <i>Similarity</i> and <i>Colocation</i> when rendering the weight distribution inside a virtual stick. 	
Design Contributions		
Chapter 12	<ul style="list-style-type: none"> – Designs of immersive interactions that not only involve a physical or a virtual real-time technique, but combine both for greater system control over the proxy-based haptic feedback, enhanced <i>Similarity</i>, and improved <i>Colocation</i>. 	

Finally, our overview is closed with the fourth box summing up our major contributions in the context of Part V and **RQ 4**. Having advanced proxy-based physical rendering and virtual manipulation techniques in Parts III and IV, we take proxy-based haptics even one step further by combining real-time physical and virtual approaches in Part V. We answer **RQ 4** by demonstrating that combining weight-shifting DPHF and HR below perceptual DTs is practically feasible and, most importantly, desirable. This conclusion is based on the results of our final investigation, which could show that with a combined technique, a VR system can better solve the challenges of *Similarity* and *Colocation* than with the individual techniques alone – further approaching the feedback quality expected from an ultimate display.

Chapter 14

Future Work and Closing Remarks

The summary of our contributions in the previous chapter showed that we could advance the flexibility of proxy-based haptics by answering **RQ 1** to **RQ 4** throughout Parts II to V. Yet, further research questions that were out of the scope of this thesis remain to be answered. Based on the results of our work, some of them can now be explored, while others emerged from our contributions. In this chapter, we close this thesis by summarizing starting points for future research.

14.1 Future Work

We already proposed ideas for future work in the individual chapters of this thesis. To complement these, we provide a catalogue of further research topics in the following, which might follow up on our work in Parts II to V.

14.1.1 Advancing Systems and Applications

The first class of topics is concerned with continuing our work on the **Applicability** of proxy-based haptics in Part II.

Multiuser Settings Future research efforts should consider multiuser applications. Here, research is necessary on how multiple immersed users can interact with (shared) haptic proxies while being in the same physical space (as investigated by Cheng et al. [2017a]) or in remote settings (as investigated by Auda et al. [2021]), and while DPHF and HR are involved. Additionally, it seems important to study how non-immersed bystanders can be involved in proxy-based VR experiences through, for example, semi-immersive solutions like AR or projections, as

we proposed in earlier work [Zenner et al., 2018a, 2019a]. By this, proxy-based VR systems targeting teamwork-focused domains could foster collaboration between immersed and non-immersed users.

Improving the *Immersive Process Model Exploration* System Continuing our work in Part II, our system for *Immersive Process Model Exploration* could further be extended. For example, future iterations of the system could add support for seated exploration, which would serve as a more accessible alternative to the room-scale experience proposed in chapter 3. By this, *Immersive Process Model Exploration* could be used also in desktop environments. This, however, would likely lead to the development of new prop-based interactions suitable for spatially constraint settings like office workspaces. Future work might also investigate props and visualizations that are more closely tailored to the process domain. Users could, for example, carry tangible versions of the documents or objects involved in the real processes through the IVE (instead of abstracted information packets), or meet virtual avatars of involved persons to facilitate memorization and recall. In addition, animations could simulate the relative duration of individual process steps to highlight time bottlenecks. By integrating features known from editor applications, the presented system could eventually evolve to a fully-featured VR tool for immersive process modeling supported by proxy-based haptics. Moreover, it would be interesting to study and compare the short- and long-term learning effects of efficiency-focused 2D and experience-focused VR exploration interfaces.

Transfer to Further Domains Apart from business process model exploration, future work should investigate how proxy-based haptics, DPHF, and HR can benefit further application areas. Domains of interest include, for example, (online) shopping (as briefly considered with our virtual pan store demonstration in chapter 5 and in our previous work on furniture retail [Zenner et al., 2020b]), rehabilitation, training, entertainment, and further domains, which profit from enhancements in feedback flexibility, haptic realism, and presence, such as those listed in Table 2.1. While we found proxy-based haptics to constitute a valuable alternative to traditional 2D interfaces for process model exploration, it remains to be explored if, and to what extent, these findings transfer to other domains.

14.1.2 Advancing the Physical Approach

The second topic area connects to our work on the **Improvement** of proxy-based haptics in Part III. While we could lay the foundations for DPHF in this thesis, additional research is required to further develop this novel class of VR haptics.

Improving the Approaches of *Shifty* and *Drag:on* We could show that the feedback of *Shifty* and *Drag:on* can give rise to convincing visual-haptic illusions,

conveying, for example, different virtual object shapes and weights. Yet, we think that even more illusions based on dynamic touch [Turvey, 1996] can be achieved. Future investigations should thus explore which virtual object properties can be convincingly simulated with DPHF proxies that change their weight distribution and air resistance. Moreover, with regard to DPHF-based controllers like *Drag:on*, we recommend exploring also the exploitation of aerodynamic lift. While we used *Drag:on*'s surface elements only for slowing down the user's hand, movement trajectories could be deflected laterally when variable surfaces on the controller are used as airfoils to produce lift. Interesting synergies might come to light when combining such physical hand redirection with virtual hand redirection as studied in Part IV. In addition to all that, the derivation of guidelines on how to create successful visual-haptic illusions based on DPHF should be a focus of future work. Finally, the haptic effects achievable through changes in inertia and drag seem to be generally understudied in the HCI domain. Hence, we propose to investigate the effect of adaptable inertia and air resistance on 3D UI interaction (like pointing onto a 3D menu) and 2D mouse-based interaction (like desktop gaming or classical UI navigation).

Combining the Approaches of *Shifty* and *Drag:on* While *Drag:on* already combines inertial and drag-based adjustments, the focus of our investigations in chapters 5 and 6 lay on the individual concepts of weight-shifting and drag-based DPHF, respectively. Yet, as outlined in chapter 7, we are optimistic that an advanced combination of both concepts could yield additional benefits. Future research should explore if, by combining the ideas of *Shifty* and *Drag:on*, the rendering of kinesthetic haptics throughout the whole object movement cycle can be improved as hypothesized in Figure 7.1. We invite researchers to build upon our open-source contributions when looking into this topic, as well as when exploring further improvements of the concepts of *Shifty* and *Drag:on*.

Exploring Further DPHF Dimensions Our exploration of DPHF focused on kinesthetic feedback. Other haptic cues, such as those relating to tactile perception of texture, temperature, compliance, etc. were out of the scope of this thesis (although subject to work we conducted in parallel [Degraen et al., 2019, 2020]). To fill this gap, we thus propose exploring how further haptic qualities, including tactile properties of virtual objects, can be rendered through DPHF. As the construction of dynamic props calls for multidisciplinary efforts [Lim et al., 2021], we suggest developing novel proxies by combining HCI and computer science knowledge with expertise from mechanical engineering and psychophysics. Advanced proxy designs that take advantage of perception-based computational fabrication [Fujinawa et al., 2017], methods for creating dynamically reconfigurable devices [Yang et al., 2022], and recent advances in soft robotics, smart materials, and modular robotics [Alexander et al., 2018] promise exciting new possibilities for adaptation. Beyond that, future research should explore how everyday objects can provide haptic feedback in VR, following the ideas we

discussed recently in our workshop on *Everyday Proxy Objects for Virtual Reality* at the *ACM CHI Conference on Human Factors in Computing Systems* [Daiber et al., 2021] and in our work on the integration of smart devices in VR [Makhsadov et al., 2022]. Our group's work on leveraging indoor climbing walls as proxies for virtual walls and mountain faces is one example for this [Kosmalla et al., 2017, 2020, 2022]. Yet, apart from specialized sports equipment, also widely available objects like kitchen utensils [Feick et al., 2022], tools, or items found in office spaces might lend themselves to being used as VR props.

Improving DPHF Rendering Algorithms In the experiments of this thesis, we took a straightforward approach to DPHF rendering by mapping the available proxy transformation (i.e., the shifting range of *Shifty*'s internal mass and the percentage of *Drag:on*'s fan opening) to the different conditions in the most naïve way. By this, for example, we mapped the heaviest virtual object to the greatest weight shift and the maximum surface area, respectively. In order to lift DPHF rendering to the next stage, more elaborate algorithms should be considered. We think that DPHF rendering should compute for each virtual object the user aims to interact with the most suitable proxy configuration. Such an optimization-based rendering could be driven either by physically-based heuristics, or be based on empirically collected data that reveals how different proxy states are perceived (similar to the approach by Shigeyama et al. [2019]). For this, physical and perceptual models of the DPHF proxies would be required, which, in the optimal case, would be multisensory and encompassing also the impact of visual feedback and multisensory integration [Ernst and Banks, 2002]. Once such rendering strategies have been defined, the next step should be to establish standardized methods for creating the required models, so that these rendering techniques can be easily applied to new proxies. Such models would then also reveal how much physical adaptation is required for convincingly simulating specific virtual objects when taking advantage of DPHF and visual-haptic illusions. This, in turn, could lead to optimized proxy designs with respect to manufacturing complexity, cost, and ergonomics, for example, reducing proxy weight or size while maintaining feedback fidelity.

Combining DPHF with Other Physical Rendering Approaches By now, most DPHF devices provide multimodal haptic feedback (i.e., impressions of shape, size, material, weight, etc.), but are only able to adapt a single haptic modality. Thanks to visual-haptic illusions, this might be sufficient for a basic simulation of various virtual objects (as we showed with *Shifty* and *Drag:on*). Nonetheless, DPHF props capable of multi-property adjustment seem desirable to improve feedback quality. Hence, we recommend looking into the combination of different DPHF approaches. Moreover, DPHF can be combined with other physical rendering techniques, such as vibrotactile feedback for enhanced *Similarity*, or ETHF for complete *Colocation* (for example, expanding on the work of Mercado et al. [2021a]). Furthermore, with *Shifty* and *Drag:on* delivering kinesthetic im-

pressions, a combination with tactile techniques promises to be a fruitful path towards the haptic feedback of next-generation VR controllers. Finally, as haptic feedback for VR becomes more powerful and realistic, future work should also consider potentially emerging safety issues. For example, we encourage research on how physical and psychological harm caused by realistic haptic feedback can be prevented, following the view of Clavelin et al. [2022].

Pursuing the Vision of the Ultimate Proxy In line with Sutherland’s vision of the ultimate display, the goal of DPHF research would be an *ultimate proxy*. Such a proxy would be *generic* while able to provide *realistic* feedback [Mündler et al., 2022], or in other words, it would be capable of adjusting all of its haptic dimensions and would provide optimal haptic feedback (at least for a large set of applications) while being controlled by the VR system. Projecting research on DPHF and its neighboring concept of ETHF far into the future, the ultimate proxy might represent the point where both concepts eventually converge. This becomes apparent when considering McNeely, who describes that “[t]he general solution to robotic graphics is to provide cellular robots that dynamically configure themselves into the desired shape and size, lock together and simulate the desired object. We call this a ‘roboxel,’ standing for ‘robotic volume element.’” [McNeely, 1993]. As such a roboxel-based solution to VR haptics would realize the ultimate proxy, we encourage research at the intersection of ETHF and DPHF.

14.1.3 Advancing the Virtual Approach

The third block of research ideas consists of topics that follow up on the **Improvement** of proxy-based haptics through our HR research conducted in Part IV.

Deepening the Knowledge About How HR Is Perceived We explored the detectability of HR in worst-case scenarios to derive the lower bounds of what is possible when covertly redirecting the user’s hand in desktop-scale settings. In many scenarios, however, it seems likely that more HR could go unnoticed than predicted by these lower bounds [Geslain et al., 2021]. Thus, future research should aim for developing a more complete perceptual model of HR. Such a model would, in the optimal case, tell about the boundaries of detectability and tolerance based on perceptual data. To collect such data, future experiments should study how the perception of HR varies with hand movement angle and speed (a topic that we explored recently in [Feick et al., 2021]), distraction, adaptation, fidelity of body visualization [Ogawa et al., 2021; Dewez et al., 2021], user activity (e.g., considering HR while lying down, sitting, standing, walking in place, during real walking or RDW), and DT estimation method [Grechkin et al., 2016]. Moreover, models to predict redirected hand trajectories and arrival times should be developed as proposed by Gonzalez et al. [2019, 2020]. Going even one step further, systems to predict the occurrence of a semantic violation

based on physiological signs or brain-computer interface signals might emerge. This also raises the question how to best maintain the sense of body ownership during HR [Dewez et al., 2021] and how much cognitive resources (e.g., in terms of cognitive load) are occupied when users are redirected below and beyond DTs – a topic also discussed in the field of RDW [Nilsson et al., 2018]. Finally, we see a need for research on potential negative side effects of HR. In this context, it is of concern whether HR can lead to problems (e.g., reduced task performance) when interacting in the real world after prolonged use of, and adaptation to, HR; how re-calibration to the real world could be supported; and whether spatial memory is affected by long-term redirection.

Improving HR Based on Change Blindness With *BSHR*, we proposed the first HR algorithm that utilizes eye tracking to leverage blink-induced change blindness. To continue this work, we propose two immediate next steps. Firstly, we suggest extending the *BSHR* algorithm with a fail-safe mechanism. This mechanism should ensure that users reach their physical destination under all circumstances, i.e., even if no blinks occur while reaching. To realize this, a fail-safe mechanism would sacrifice unnoticeability in favor of reachability by boosting continuous hand warping beyond DTs if users do not blink in time. In addition, enhanced *BSHR* algorithms should take advantage of all blinks that occur during a redirection, instead of only using the first blink. Secondly, we propose to conduct a follow-up study to investigate if combined *BSHR* with fine-tuned parameters can outperform the traditional approach of continuous warping regarding unnoticeable redirection. Such a study could now use the DTs of the algorithm by Cheng et al. [2017b] derived in chapter 9 as more optimal parameters for *BSHR*. Alternatively, follow-up experiments could also consider a per-user calibration of DT parameters. In this context, suitable per-user calibration methods that automatically select optimal *BSHR* parameters with minimal setup effort seem to be an interesting avenue for future work. Finally, we encourage continuing research on HR techniques that intentionally take advantage of change blindness. A guiding vision is that of future HR techniques that can flexibly exploit all kinds of change blindness opportunities, be it occluded views, objects leaving the FOV or the areas of focus and attention [Marwecki et al., 2019], blinks, saccades [Sun et al., 2018], or other forms of distraction.

Improving the *HaRT* To complement the preliminary results of our expert study on the *HaRT*, future work might conduct a follow-up evaluation with more users, for example, by reaching out to the GitHub audience. This study should be conducted with access to VR hardware and also cover world warping techniques. In the long run, we propose to keep the *HaRT* up to date as new techniques and findings emerge, in order to channel scientific advances on HR to the open-source community. Moreover, future work might extend our framework by adding support for hybrid hand redirection techniques, as well as multi-finger redirection. Apart from that, the *HaRT* could also be extended with implementations of

pseudo-haptic techniques [Dominjon et al., 2005; Lécuyer, 2009; Samad et al., 2019], per-user calibration methods, reach target prediction, and benchmarking functionality. A simulation of hand movements (in line with the idea of the *Simulated Walker* in the RDW toolkit by Azmandian et al. [2016a]) could eventually even enable simulation studies similar to the work by Nguyen and Kunz [2018], and thereby reduce the need for user experiments. Finally, our toolkit could be ported to other 3D engines.

Developing and Supporting Novel Hand Redirection Techniques Continuing our work in this thesis, we encourage exploring novel hand redirection techniques that leverage characteristics of human perception yet unused in this domain. Beyond blink suppression, for example, we currently work on algorithms that make use of saccadic suppression for HR – an approach that, just like the blink-based approach, could be transferred also to world and hybrid warping. Leveraging saccades would open many more opportunities for injecting hand offsets compared to blinks, each of which, however, would be much shorter and more difficult to track. To reduce the dependence on spontaneous blinks and saccades, future research might further study how VR systems can control the blink and gaze behavior of users. To this end, we recently investigated software- and hardware-based methods for triggering blinks on demand in VR, and collected first promising results. The idea is related to *Subtle Gaze Direction*, which has been used recently by Sun et al. [2018] to increase the user’s saccade frequency in VR. Furthermore, models to predict the occurrence of blinks and saccades would allow for improved redirection planning, potentially in combination with attention models similar to those employed by Marwecki et al. [2019]. Knowing about the user’s focus of attention in real-time would open up exciting opportunities for novel hand redirection techniques. These could inject changes only when the user’s attention lies elsewhere, or control the amount of warping based on what the user currently attends to. Such strategies might work well especially in combination with system-controlled distractors [Cools and Simeone, 2019]. On the other hand, the question of how introduced discrepancies can be reverted again has been left understudied. Here, moments of change blindness could be exploited to re-establish alignment of real and virtual hands [Hartfill et al., 2021]. In addition, we think that conducting research on bi-manual HR will become more important as VR interactions increase in complexity and realism, continuing, for example, research by Gonzalez and Follmer [2019]. Exploring how hand redirection approaches can be employed for redirecting other body parts, such as the user’s feet, might constitute another interesting path of research. Scenarios that involve the operation of pedals or precise foot movements (like VR climbing [Kosmalla et al., 2020]) might profit from techniques for foot redirection. Finally, we encourage future work to consider security aspects [Tseng et al., 2022] and to transfer the recent advances in the domain of hand redirection to practical use cases and real applications that go beyond research prototypes, integrating redirection in productive VR systems.

14.1.4 Advancing the Combined Approach

Lastly, future work concerned with the combination of multiple techniques makes up the fourth block of potential research topics, following up on our **Improvement** of proxy-based haptics in Part V.

Generalizing the Combination of DPHF and HR We showed that when rendering weight distribution, combining the physical approach of DPHF and the virtual technique of HR can better solve the challenges of *Similarity* and *Colocation* than the individual techniques alone can do. Future work should now focus on generalizing these findings. We encourage investigating how dynamic props and haptic retargeting can form a symbiosis also when rendering other haptic properties, such as texture, stiffness, shape, resistance, or elasticity. The goal would be the derivation of rules and a set of requirements that VR systems must meet to profit from combined techniques. As a next step in this line of research, we recently proposed a VR system, in which a small set of proxies can represent larger sets of virtual objects [Düwel, 2021]. To realize this, our system employs HR and DPHF props. The system contributes to a generalization as it employs computational optimization to propose suitable prop placements and mappings. For this optimization, perceptual aspects are taken into account, such as the amplitude of required HR and the degree of mismatch between haptic characteristics. Future research might expand on this approach.

Combining Further Techniques Finally, combining DPHF and HR is not the only way systems can take advantage of the breadth of physical and virtual techniques introduced by researchers. While many techniques are compatible, the combination of most of them is still severely under-explored. Hence, we propose to systematically investigate how DPHF and HR can be combined with other physical (e.g., ETHF, EMS, or AHF) and virtual strategies (e.g., pseudo-haptics, change blindness remapping, or RDW). Investigations could reveal the prerequisites, benefits, and drawbacks of various combinations of techniques. Resulting technical and theoretical findings, as well as empirical data could then be collected and documented in a central repository. Such a repository could serve as a guide for future investigations and as a reference for VR practitioners.

14.2 Closing Remarks

This thesis advanced proxy-based haptic feedback for VR. The goal of this thesis was twofold as we first showcased the applicability of haptic proxies, and secondly advanced the field through contributions to physical, virtual, and combined techniques that make proxy-based haptics more flexible. By this, we considered the two central challenges faced by proxy-based VR systems, namely, the challenges of *Similarity* and *Colocation*.

With our research, we could provide answers to the four central research questions this thesis was committed to, and made several theoretical, technical, and design-related contributions. The contributed concepts, algorithms, prototypes, and systems, as well as the empirical results of our experiments, bear potential to lay the foundation for techniques, which eventually might become an integral part of the VR systems of the future. By this, we could extend the knowledge of the scientific community and push the fields of HCI, VR, haptics, and psychophysics closer towards Sutherland's vision of the ultimate display.

It is to the submission of this thesis, that such advancements seem more timely than ever before, especially in light of the recent efforts undertaken by major players in the industry to push VR and the development of a *metaverse*⁸². Our contributions strongly motivate continuing our research endeavors as the concept of DPHF, perception-inspired HR techniques, and in particular their combination, might constitute basic building blocks for bringing immersive haptics to the ubiquitous virtual spaces of the future. As a result, we might eventually realize Sutherland's 1965 vision of the ultimate display further down this research path – with (dynamic) proxies and perceptual illusions serving as links between the worlds that bring the qualities of our physical reality to the virtual cosmos.

⁸²Metaverse on Wikipedia. <https://w.wiki/4Jbt>

List of Figures

1.1	A Typical VR System	4
1.2	Screenshot of <i>The Climb</i> by Crytek and a Virtual Kitchen Store	5
1.3	The Active-Passive Haptics Continuum	7
1.4	Thesis Outline	17
2.1	A 1930's Stereoscope and a Modern Cardboard VR HMD	20
2.2	The HMD by Sutherland [1968] and a Modern HMD	21
2.3	The Virtuality Continuum	23
2.4	Fish Tank VR, a CAVE, and HMD-based VR	27
2.5	The Perception-Action Loop	31
2.6	A Perception Model Adapted from Proske and Gandevia [2012] . .	37
2.7	The 2D Visual-Haptic Reality-Virtuality Continuum	45
2.8	Examples of World-Grounded Active Kinesthetic Feedback Solutions	47
2.9	Examples of Body-Grounded Active Kinesthetic Feedback Solutions	48
2.10	Examples of Ungrounded Active Kinesthetic Feedback Solutions .	49
2.11	Examples of Contact-Based Active Tactile Feedback Solutions . .	50
2.12	Examples of Contactless Active Tactile Feedback Solutions . . .	52
2.13	Examples of IVEs Augmented with Passive Haptic Feedback . . .	54
2.14	Previous Works Investigating Proxy Inertia, Shape, and Material .	56
2.15	Examples of Application Areas for Passive Haptic Feedback . . .	58
2.16	Examples of Alternative Approaches to Passive Haptic Feedback .	59
2.17	Orthogonality of the Challenges of Similarity and Colocation . .	63
2.18	Overview of Physical Approaches to Similarity and Colocation . .	64
2.19	Overview of Offline Physical Strategies	65
2.20	Examples of Reconfigurable Proxies	66
2.21	Overview of Real-Time Physical Strategies	67
2.22	Examples of Mass Distribution-Changing DPHF Proxies	69
2.23	Examples of Stiffness-Changing DPHF Proxies	70
2.24	Examples of Shape-Changing DPHF Proxies	71

2.25	Examples of Texture-Changing DPHF Proxies	71
2.26	Examples of World-Grounded ETHF	73
2.27	Examples of Body-Grounded ETHF	74
2.28	Examples of Ungrounded ETHF	75
2.29	Examples of Human Actuation-Based ETHF	75
2.30	Overview of Virtual Approaches to Similarity and Colocation	76
2.31	Overview of Offline Virtual Strategies	77
2.32	Example of Substitutional Reality	78
2.33	Overview of Real-Time Virtual Strategies	78
2.34	Examples of Pseudo-Haptic Feedback	79
2.35	Concept of Body Warping-Based Hand Redirection	82
2.36	Examples of Haptic Retargeting	89
2.37	Example of Redirected Touching	91
2.38	Hand Redirection, Haptic Retargeting, and Redirected Touching	93
2.39	Order of Haptic Retargeting and Redirected Touching	93
2.40	Example of Change Blindness Haptic Remapping	94
2.41	Examples of Redirected Walking	96
2.42	Examples of Redirected Walking for Enhanced Colocation	97
2.43	Examples of ETHF Combined With HR	99
3.1	<i>Immersive Process Model Exploration</i> – An EPC in 2D	111
3.2	<i>Immersive Process Model Exploration</i> – Concept	113
3.3	<i>Immersive Process Model Exploration</i> – First Person VR View	114
3.4	<i>Immersive Process Model Exploration</i> – 2D and 3D EPC Elements	114
3.5	<i>Immersive Process Model Exploration</i> – Interactive Platform Elements	115
3.6	<i>Immersive Process Model Exploration</i> – View of the Immersed User	116
3.7	<i>Immersive Process Model Exploration</i> – Illustration of Unlocking	116
3.8	<i>Immersive Process Model Exploration</i> – A 3D xor Platform	117
3.9	<i>Immersive Process Model Exploration</i> – Spatial Proxy Layout	118
3.10	<i>Immersive Process Model Exploration</i> – Registration of the Proxies	119
3.11	<i>Immersive Process Model Exploration</i> – 2D EPC of the Study Process	121
3.12	<i>Immersive Process Model Exploration</i> – Study Conditions	123
3.13	<i>Immersive Process Model Exploration</i> – Study Results (Part 1 of 2)	125
3.14	<i>Immersive Process Model Exploration</i> – Study Results (Part 2 of 2)	126

4.1	<i>Dynamic Passive Haptic Feedback – Concept</i>	138
4.2	<i>Dynamic Passive Haptic Feedback – Continuum</i>	138
4.3	<i>Dynamic Passive Haptic Feedback – Classification</i>	139
5.1	<i>Shifty – Concept</i>	143
5.2	<i>Shifty – Image of a User</i>	144
5.3	<i>Shifty – Components</i>	145
5.4	<i>Shifty – Prototype</i>	146
5.5	<i>Shifty – Virtual Objects in Experiment 1</i>	148
5.6	<i>Shifty – Experiment Setup</i>	149
5.7	<i>Shifty – Results of Experiment 1</i>	150
5.8	<i>Shifty – Interaction in Experiment 2</i>	152
5.9	<i>Shifty – Pick-Up Visualizations in Experiment 2</i>	154
5.10	<i>Shifty – Location in the Haptics Continuum</i>	160
5.11	<i>Shifty – Virtual Shopping Simulation</i>	161
5.12	<i>Shifty – Virtual Maintenance and Shelf Stocking Simulations</i>	162
5.13	<i>Shifty – Illusions Based on Weight-Shifting DPHF</i>	164
6.1	<i>Drag:on – Concept</i>	167
6.2	<i>Drag:on – States Investigated in the Study</i>	168
6.3	<i>Drag:on – Investigated Types of Interactions</i>	169
6.4	<i>Drag:on – 3D Rendering of the Prototype</i>	170
6.5	<i>Drag:on – Components</i>	171
6.6	<i>Drag:on – Software Architecture</i>	171
6.7	<i>Drag:on – Scale Scenario</i>	173
6.8	<i>Drag:on – Material, Ratchet, and Wagon Scenarios</i>	174
6.9	<i>Drag:on – Flow Scenario</i>	175
6.10	<i>Drag:on – Results of Part 1 of the Study</i>	177
6.11	<i>Drag:on – Ratchet Results of Part 2 of the Study</i>	178
6.12	<i>Drag:on – Realism Results of Part 2 of the Study</i>	178
6.13	<i>Drag:on – Wagon Results of Part 2 of the Study</i>	179
7.1	<i>Part III – Kinesthetic Feedback Across the Object Movement Cycle</i> .	186
8.1	<i>Continuous Hand Redirection – Horizontal Warping</i>	193
8.2	<i>Continuous Hand Redirection – Vertical Warping</i>	194
8.3	<i>Continuous Hand Redirection – Gain Warping</i>	194

8.4	<i>Continuous Hand Redirection</i> – Hardware Setup	197
8.5	<i>Continuous Hand Redirection</i> – Study Results (Part 1 of 2)	200
8.6	<i>Continuous Hand Redirection</i> – Study Results (Part 2 of 2)	203
8.7	<i>Continuous Hand Redirection</i> – Threshold Visualization	204
9.1	<i>Blink-Suppressed Hand Redirection</i> – Concept	211
9.2	<i>Blink-Suppressed Hand Redirection</i> – Assumption 2	213
9.3	<i>Blink-Suppressed Hand Redirection</i> – P' Computation	214
9.4	<i>Blink-Suppressed Hand Redirection</i> – Experiment Setup	220
9.5	<i>Blink-Suppressed Hand Redirection</i> – Experiment in VR	221
9.6	<i>Blink-Suppressed Hand Redirection</i> – Detection Thresholds	223
9.7	<i>Blink-Suppressed Hand Redirection</i> – Staircase Plot	224
10.1	<i>Hand Redirection Toolkit</i> – Features	231
10.2	<i>Hand Redirection Toolkit</i> – Settings	232
10.3	<i>Hand Redirection Toolkit</i> – Visualizations	234
10.4	<i>Hand Redirection Toolkit</i> – Threshold Visualization	234
10.5	<i>Hand Redirection Toolkit</i> – Study Results	237
12.1	<i>Combined DPHF and HR</i> – Investigated Interaction	248
12.2	<i>Combined DPHF and HR</i> – Sketch of Conditions	249
12.3	<i>Combined DPHF and HR</i> – Thought Experiment 1	253
12.4	<i>Combined DPHF and HR</i> – Thought Experiment 2	255
12.5	<i>Combined DPHF and HR</i> – Experiment Setup	257
12.6	<i>Combined DPHF and HR</i> – Virtual Scene	258
12.7	<i>Combined DPHF and HR</i> – Experiment Walkthrough	259
12.8	<i>Combined DPHF and HR</i> – Conditions in the <i>Similarity</i> Experiment	260
12.9	<i>Combined DPHF and HR</i> – Results of the <i>Similarity</i> Experiment	261
12.10	<i>Combined DPHF and HR</i> – Conditions in the <i>Colocation</i> Experiment	264
12.11	<i>Combined DPHF and HR</i> – Staircase Plot	265
12.12	<i>Combined DPHF and HR</i> – Results of the <i>Colocation</i> Experiment	266

List of Tables

2.1	Areas that Successfully Employ VR	29
2.2	Comparison of Active and Passive Haptic Feedback	61
2.3	Overview of Body Warping-Based Hand Redirection Research	84
2.4	Overview of Hand Redirection Thresholds Research (Part 1 of 2)	86
2.5	Overview of Hand Redirection Thresholds Research (Part 2 of 2)	87
2.6	Overview of Employed Assessment Methods	101
3.1	<i>Immersive Process Model Exploration</i> – Study Results	127
4.1	<i>Dynamic Passive Haptic Feedback</i> – Characteristics	140
5.1	<i>Shifty</i> – Conditions in Experiment 2	153
5.2	<i>Shifty</i> – Results of Experiment 2	155
6.1	<i>Drag:on</i> – Technical Data	170
6.2	<i>Drag:on</i> – States Investigated in the Study	172
8.1	<i>Continuous Hand Redirection</i> – Detection Thresholds	201
9.1	<i>Blink-Suppressed Hand Redirection</i> – Detection Thresholds	222

List of Algorithms

1	<i>Continuous Hand Redirection</i> – Rotational Warp Algorithm	193
2	<i>Continuous Hand Redirection</i> – Gain Warp Algorithm	195
3	<i>Blink-Suppressed Hand Redirection</i> – Algorithm	215

Abbreviations

1-DoF	One Degree of Freedom
1D	One-Dimensional
2D	Two-Dimensional
3-DoF	Three Degrees of Freedom
3D	Three-Dimensional
6-DoF	Six Degrees of Freedom
AFC	Alternative Forced-Choice
AHF	Active Haptic Feedback
API	Application Programming Interface
AR	Augmented Reality
AV	Augmented Virtuality
BIP	Break in Presence
BPMN	Business Process Model and Notation
BSHR	Blink-Suppressed Hand Redirection
C/D	Control/Display
CAVE	Cave Automatic Virtual Environment
CDT	Conservative Detection Threshold
CoM	Center of Mass
DFKI	German Research Center for Artificial Intelligence
DoFs	Degrees of Freedom
DPHF	Dynamic Passive Haptic Feedback
DT	Detection Threshold
EMS	Electrical Muscle Stimulation
EPC	Event-Driven Process Chain
ERM	Eccentric Rotating Mass
ETHF	Encountered-Type Haptic Feedback
FOV	Field of View
HaRT	Hand Redirection Toolkit

HCI	Human-Computer Interaction
HMD	Head-Mounted Display
HR	Body Warping-based Hand Redirection
IRL	Innovative Retail Laboratory
IVE	Immersive Virtual Environment
JND	Just-Noticeable Difference
LRA	Linear Resonant Actuator
MLE	Maximum-Likelihood Estimation
MR	Mixed Reality
NASA TLX	NASA Task Load Index
PHF	Passive Haptic Feedback
PSE	Point of Subjective Equality
R/V	Real-to-Virtual
RDW	Redirected Walking
RQ	Research Question
SDK	Software Development Kit
SR	Substitutional Reality
SSQ	Simulator Sickness Questionnaire
SUS	Slater-Usoh-Steed Presence Questionnaire
TEC	Thermoelectric Cooler
UEQ	User Experience Questionnaire
UEQ-S	User Experience Questionnaire (Short Version)
UI	User Interface
UX	User Experience
UXF	Unity Experiment Framework
VR	Virtual Reality
VT	Virtual Target

Bibliography

Abdullah, M., Kim, M., Hassan, W., Kuroda, Y., and Jeon, S. (2017). HapticDrone: An Encountered-Type Kinesthetic Haptic Interface with Controllable Force Feedback: Initial Example for 1D Haptic Feedback. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST'17 Adjunct, pages 115–117. ACM.

Abtahi, P. and Follmer, S. (2018). Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'18, pages 1–13. ACM.

Abtahi, P., Landry, B., Yang, J., Pavone, M., Follmer, S., and Landay, J. A. (2019). Beyond The Force: Using Quadcopters to Appropriate Objects and the Environment for Haptics in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'19, pages 1–13. ACM.

Achberger, A., Arulrajah, P., Vidakovic, K., and Sedlmair, M. (2022). STROE: An Ungrounded String-Based Weight Simulation Device. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'22, pages 112–120. IEEE.

Achberger, A., Heyen, F., Vidakovic, K., and Sedlmair, M. (2021). PropellerHand: A Hand-Mounted, Propeller-Based Force Feedback Device. In *Proceedings of the International Symposium on Visual Information Communication and Interaction*, VINCI'21, pages 1–8. ACM.

Achibet, M., Girard, A., Talvas, A., Marchal, M., and Lécuyer, A. (2015). Elastic-Arm: Human-Scale Passive Haptic Feedback for Augmenting Interaction and Perception in Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'15, pages 63–68. IEEE.

Achibet, M., Le Gouis, B., Marchal, M., Léziart, P.-A., Argelaguet, F., Girard, A., Lécuyer, A., and Kajimoto, H. (2017). FlexiFingers: Multi-Finger Interaction in VR Combining Passive Haptics and Pseudo-Haptics. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, 3DUI'17, pages 103–106. IEEE.

Achibet, M., Marchal, M., Argelaguet, F., and Lécuyer, A. (2014). The Virtual Mitten: A Novel Interaction Paradigm for Visuo-Haptic Manipulation of Objects Using Grip Force. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, 3DUI'14, pages 59–66. IEEE.

Alexander, J., Roudaut, A., Steimle, J., Hornbæk, K., Bruns Alonso, M., Follmer, S., and Merritt, T. (2018). Grand Challenges in Shape-Changing Interface Research. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'18*, pages 1–14. ACM.

Araujo, B., Jota, R., Perumal, V., Yao, J. X., Singh, K., and Wigdor, D. (2016). Snake Charmer: Physically Enabling Virtual Objects. In *Proceedings of the ACM Conference on Tangible, Embedded and Embodied Interaction, TEI'16*, pages 218–226. ACM.

Argelaguet, F., Hoyet, L., Trico, M., and Lécuyer, A. (2016). The Role of Interaction in Virtual Embodiment: Effects of the Virtual Hand Representation. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'16*, pages 3–10. IEEE.

Argelaguet, F., Jáuregui, D. A. G., Marchal, M., and Lécuyer, A. (2013). Elastic Images: Perceiving Local Elasticity of Images through a Novel Pseudo-Haptic Deformation Effect. *ACM Transactions on Applied Perception*, 10(3):17:1–17:14.

Ariza, O. (2020). *Wearable Haptic Technology for 3D Selection and Guidance*. PhD thesis, Universität Hamburg.

Ariza, O., Bruder, G., Katzakis, N., and Steinicke, F. (2018). Analysis of Proximity-Based Multimodal Feedback for 3D Selection in Immersive Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'18*, pages 327–334. IEEE.

Ariza, O., Freiwald, J., Laage, N., Feist, M., Salloum, M., Bruder, G., and Steinicke, F. (2016). Inducing Body-Transfer Illusions in VR by Providing Brief Phases of Visual-Tactile Stimulation. In *Proceedings of the ACM Symposium on Spatial User Interaction, SUI'16*, pages 61–68. ACM.

Ariza, O., Lubos, P., and Steinicke, F. (2015). HapRing: A Wearable Haptic Device for 3D Interaction. In Pielot, M., Diefenbach, S., and Henze, N., editors, *Proceedings of the Mensch und Computer Conference, MuC'15*, pages 421–424. De Gruyter.

Arora, J., Saini, A., Mehra, N., Jain, V., Shrey, S., and Parnami, A. (2019). Virtual-Bricks: Exploring a Scalable, Modular Toolkit for Enabling Physical Manipulation in VR. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'19*, pages 1–12. ACM.

Auda, J., Busse, L., Pfeuffer, K., Gruenfeld, U., Rivu, R., Alt, F., and Schneegass, S. (2021). I'm in Control! Transferring Object Ownership Between Remote Users with Haptic Props in Virtual Reality. In *Proceedings of the ACM Symposium on Spatial User Interaction, SUI'21*, pages 1–10. ACM.

Azmandian, M., Grechkin, T., Bolas, M., and Suma, E. (2016a). The Redirected Walking Toolkit: A Unified Development Platform for Exploring Large Virtual Environments. In *Proceedings of the IEEE Workshop on Everyday Virtual Reality*, WEVR'16, pages 9–14. IEEE.

Azmandian, M., Hancock, M., Benko, H., Ofek, E., and Wilson, A. D. (2016b). Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'16, pages 1968–1979. ACM.

Badshah, A., Gupta, S., Morris, D., Patel, S., and Tan, D. (2012). GyroTab: A Handheld Device That Provides Reactive Torque Feedback. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'12, pages 3153–3156. ACM.

Ban, Y., Kajinami, T., Narumi, T., Tanikawa, T., and Hirose, M. (2012). Modifying an Identified Angle of Edged Shapes Using Pseudo-haptic Effects. In Isokoski, P. and Springare, J., editors, *Proceedings of the EuroHaptics Conference*, EuroHaptics'12, pages 25–36. Springer.

Ban, Y., Narumi, T., Tanikawa, T., and Hirose, M. (2014). Displaying Shapes with Various Types of Surfaces Using Visuo-Haptic Interaction. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST'14, pages 191–196. ACM.

Barfield, W. and Danas, E. (1996). Comments on the Use of Olfactory Displays for Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 5(1):109–121.

Benda, B., Esmaeili, S., and Ragan, E. D. (2020). Determining Detection Thresholds for Fixed Positional Offsets for Virtual Hand Remapping in Virtual Reality. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*, ISMAR'20, pages 269–278. IEEE.

Benko, H., Holz, C., Sinclair, M., and Ofek, E. (2016). NormalTouch and Texture-Touch: High-Fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'16, pages 717–728. ACM.

Bergström, J., Mottelson, A., and Knibbe, J. (2019). Resized Grasping in VR: Estimating Thresholds for Object Discrimination. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'19, pages 1175–1183. ACM.

Bolte, B. and Lappe, M. (2015). Subliminal Reorientation and Repositioning in Immersive Virtual Environments using Saccadic Suppression. *IEEE Transactions on Visualization and Computer Graphics*, 21(4):545–552.

Borst, C. W. and Volz, R. A. (2003). Preliminary Report on a Haptic Feedback Technique for Basic Interactions with a Virtual Control Panel. In *Proceedings of the EuroHaptics Conference*, EuroHaptics'03, pages 1–13. EuroHaptics Society.

Bouzbib, E., Bailly, G., Haliyo, S., and Frey, P. (2020). CoVR: A Large-Scale Force-Feedback Robotic Interface for Non-Deterministic Scenarios in VR. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'20, pages 209–222. ACM.

Bouzit, M., Popescu, G., Burdea, G., and Boian, R. (2002). The Rutgers Master II-ND Force Feedback Glove. In *Proceedings of the IEEE Haptics Symposium*, HAPTICS'02, pages 145–152. IEEE.

Bowman, D. A., Koller, D., and Hodges, L. F. (1997). Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'97, pages 45–52. IEEE.

Braun, V. and Clarke, V. (2012). Thematic Analysis. In *APA Handbook of Research Methods in Psychology, Vol 2: Research Designs: Quantitative, Qualitative, Neuropsychological, and Biological*, APA Handbooks in Psychology, chapter 4, pages 57–71. American Psychological Association.

Bristow, D., Haynes, J.-D., Sylvester, R., Frith, C. D., and Rees, G. (2005). Blinking Suppresses the Neural Response to Unchanging Retinal Stimulation. *Current Biology*, 15(14):1296–1300.

Brookes, J., Warburton, M., Alghadier, M., Mon-Williams, M., and Mushtaq, F. (2020). Studying Human Behavior with Virtual Reality: The Unity Experiment Framework. *Behavior Research Methods*, 52(2):455—463.

Brooks, F. P., Ouh-Young, M., Batter, J. J., and Jerome Kilpatrick, P. (1990). Project GROPE – Haptic Displays for Scientific Visualization. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH'90, pages 177–185. ACM.

Burdea, G. C. (1999a). Invited Review: The Synergy Between Virtual Reality and Robotics. *IEEE Transactions on Robotics and Automation*, 15(3):400–410.

Burdea, G. C. (1999b). Keynote Address: Haptic Feedback for Virtual Reality. In *Proceedings of the International Workshop on Virtual Prototyping*, pages 87–96.

Burns, E. and Brooks, F. P. (2006). Perceptual Sensitivity to Visual/Kinesthetic Discrepancy in Hand Speed, and Why We Might Care. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST'06, pages 3–8. ACM.

Burns, E., Razzaque, S., Panter, A. T., Whitton, M. C., McCallus, M. R., and Brooks, F. P. (2006). The Hand Is More Easily Fooled than the Eye: Users Are More Sensitive to Visual Interpenetration than to Visual-Proprioceptive Discrepancy. *Presence: Teleoperators and Virtual Environments*, 15(1):1–15.

Butz, A. and Krüger, A. (2017). *Mensch-Maschine-Interaktion*. De Gruyter Oldenbourg.

Böhme, C., Hartmann, J., Kern, H., Kühne, S., Laue, R., Nüttgens, M., Rump, F. J., and Storch, A. (2010). bflow* Toolbox – An Open-Source Modeling Tool. In *Proceedings of the Business Process Management Demonstration Track*, BPM’10, pages 46–51.

Carvalheiro, C., Nóbrega, R., da Silva, H., and Rodrigues, R. (2016). User Redirection and Direct Haptics in Virtual Environments. In *Proceedings of the ACM Conference on Multimedia*, MM’16, pages 1146–1155. ACM.

Chan, T.-C. (1995). The Effect of Density and Diameter on Haptic Perception of Rod Length. *Perception & Psychophysics*, 57(6):778–786.

Chen, Y.-S., Han, P.-H., Hsiao, J.-C., Lee, K.-C., Hsieh, C.-E., Lu, K.-Y., Chou, C.-H., and Hung, Y.-P. (2016). SoEs: Attachable Augmented Haptic on Gaming Controller for Immersive Interaction. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST’16 Adjunct, pages 71–72. ACM.

Cheng, C.-H., Chang, C.-C., Chen, Y.-H., Lin, Y.-L., Huang, J.-Y., Han, P.-H., Ko, J.-C., and Lee, L.-C. (2018a). GravityCup: A Liquid-Based Haptics for Simulating Dynamic Weight in Virtual Reality. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST’18, pages 1–2. ACM.

Cheng, L.-P., Chang, L., Marwecki, S., and Baudisch, P. (2018b). iTurk: Turning Passive Haptics into Active Haptics by Making Users Reconfigure Props in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’18, pages 1–10. ACM.

Cheng, L.-P., Lühne, P., Lopes, P., Sterz, C., and Baudisch, P. (2014). Haptic Turk: A Motion Platform Based on People. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’14, pages 3463–3472. ACM.

Cheng, L.-P., Marwecki, S., and Baudisch, P. (2017a). Mutual Human Actuation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’17, pages 797–805. ACM.

Cheng, L.-P., Ofek, E., Holz, C., Benko, H., and Wilson, A. D. (2017b). Sparse Haptic Proxy: Touch Feedback in Virtual Environments Using a General Passive Prop. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’17, pages 3718–3728. ACM.

Cheng, L.-P., Roumen, T., Rantzsch, H., Köhler, S., Schmidt, P., Kovacs, R., Jasper, J., Kemper, J., and Baudisch, P. (2015). TurkDeck: Physical Virtual Reality Based on People. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'15, pages 417–426. ACM.

Choi, I., Culbertson, H., Miller, M. R., Olwal, A., and Follmer, S. (2017). Grability: A Wearable Haptic Interface for Simulating Weight and Grasping in Virtual Reality. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'17, pages 119–130. ACM.

Choi, I., Hawkes, E. W., Christensen, D. L., Ploch, C. J., and Follmer, S. (2016). Wolverine: A Wearable Haptic Interface for Grasping in Virtual Reality. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, IROS'16, pages 986–993. IEEE.

Choi, I., Ofek, E., Benko, H., Sinclair, M., and Holz, C. (2018). CLAW: A Multi-functional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'18, pages 1–13. ACM.

Choi, I., Zhao, Y., Gonzalez, E. J., and Follmer, S. (2021). Augmenting Perceived Softness of Haptic Proxy Objects through Transient Vibration and Visuo-Haptic Illusion in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(12):4387–4400.

Clarence, A., Knibbe, J., Cordeil, M., and Wybrow, M. (2021). Unscripted Retargeting: Reach Prediction for Haptic Retargeting in Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'21, pages 150–159. IEEE.

Clark, D. and Bailey, M. J. (2002). Virtual-Virtual Haptic Feedback and Why it Wasn't Enough. In Erbacher, R. F., Chen, P. C., Groehn, M., Roberts, J. C., and Wittenbrink, C. M., editors, *Proceedings Volume 4665, Visualization and Data Analysis 2002*, pages 308–318. Society of Photo-Optical Instrumentation Engineers.

Clavelin, G., Gugenheimer, J., and Bouhier, M. (2022). 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, pages 1–3.

Conti, F. and Khatib, O. (2005). Spanning Large Workspaces Using Small Haptic Devices. In *Proceedings of the IEEE World Haptics Conference*, WHC'05, pages 183–188. IEEE.

Cools, R. and Simeone, A. L. (2019). Investigating the Effect of Distractor Interactivity for Redirected Walking in Virtual Reality. In *Proceedings of the ACM Symposium on Spatial User Interaction*, SUI'19, pages 1–5. ACM.

Cooper, N., Milella, F., Pinto, C., Cant, I., White, M., and Meyer, G. (2018). The Effects of Substitute Multisensory Feedback on Task Performance and the Sense of Presence in a Virtual Reality Environment. *PLOS ONE*, 13(2):1–25.

Crnovrsanin, T., Wang, Y., and Ma, K.-L. (2014). Stimulating a Blink: Reduction of Eye Fatigue with Visual Stimulus. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’14, pages 2055–2064. ACM.

Daiber, F., Degraen, D., Zenner, A., Döring, T., Steinicke, F., Ariza, O., and Simeone, A. L. (2021). Everyday Proxy Objects for Virtual Reality. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA’21, pages 1–6. ACM.

de Jesus Oliveira, V. A., Brayda, L., Nedel, L., and Maciel, A. (2017). Designing a Vibrotactile Head-Mounted Display for Spatial Awareness in 3D Spaces. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1409–1417.

de Tinguy, X., Howard, T., Pacchierotti, C., Marchal, M., and Lécuyer, A. (2020). WeATaViX: WEarable ACTuated TAngibles for VIrtual Reality eXperiences. In Nisky, I., Hartcher-O’Brien, J., Wiertlewski, M., and Smeets, J., editors, *Proceedings of the EuroHaptics Conference*, EuroHaptics’20, pages 262–270. Springer.

de Tinguy, X., Pacchierotti, C., Emily, M., Chevalier, M., Guignardat, A., Guillaudeux, M., Six, C., Lécuyer, A., and Marchal, M. (2019a). How Different Tangible and Virtual Objects Can Be While Still Feeling the Same? In *Proceedings of the IEEE World Haptics Conference*, WHC’19, pages 580–585. IEEE.

de Tinguy, X., Pacchierotti, C., Lécuyer, A., and Marchal, M. (2021). Capacitive Sensing for Improving Contact Rendering With Tangible Objects in VR. *IEEE Transactions on Visualization and Computer Graphics*, 27(4):2481–2487.

de Tinguy, X., Pacchierotti, C., Marchal, M., and Lécuyer, A. (2019b). Toward Universal Tangible Objects: Optimizing Haptic Pinching Sensations in 3D Interaction. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’19, pages 321–330. IEEE.

Debarba, H. G., Boulic, R., Salomon, R., Blanke, O., and Herbelin, B. (2018a). Self-Attribution of Distorted Reaching Movements in Immersive Virtual Reality. *Computers & Graphics*, 76:142–152.

Debarba, H. G., Khoury, J.-N., Perrin, S., Herbelin, B., and Boulic, R. (2018b). Perception of Redirected Pointing Precision in Immersive Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’18, pages 341–346. IEEE.

Degraen, D., Fruchard, B., Smolders, F., Potetsianakis, E., Güngör, S., Krüger, A., and Steimle, J. (2021a). Weirding Haptics: In-Situ Prototyping of Vibrotactile Feedback in Virtual Reality through Vocalization. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’21, pages 936—953. ACM.

Degraen, D., Piovarči, M., Bickel, B., and Krüger, A. (2021b). Capturing Tactile Properties of Real Surfaces for Haptic Reproduction. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'21, pages 954—971. ACM.

Degraen, D., Reindl, A., Makhsadov, A., Zenner, A., and Krüger, A. (2020). Envisioning Haptic Design for Immersive Virtual Environments. In *Companion Publication of the ACM Conference on Designing Interactive Systems*, DIS'20 Companion, pages 287–291. ACM.

Degraen, D., Zenner, A., and Krüger, A. (2019). Enhancing Texture Perception in Virtual Reality Using 3D-Printed Hair Structures. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'19, pages 1–12. ACM.

Dennison, M. S., Wisti, A. Z., and D'Zmura, M. (2016). Use of Physiological Signals to Predict Cybersickness. *Displays*, 44:42–52.

Deweze, D., Hoyet, L., Lécuyer, A., and Argelaguet Sanz, F. (2021). Towards “Avatar-Friendly” 3D Manipulation Techniques: Bridging the Gap Between Sense of Embodiment and Interaction in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'21, pages 1–14. ACM.

Dix, A., Finlay, J. E., Abowd, G. D., and Beale, R. (2003). *Human-Computer Interaction (3rd Edition)*. Prentice-Hall, Inc.

Dominjon, L., Lécuyer, A., Burkhardt, J.-M., Richard, P., and Richir, S. (2005). Influence of Control/Display Ratio on the Perception of Mass of Manipulated Objects in Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'05, pages 19–25. IEEE.

Doughty, M. J. (2002). Further Assessment of Gender- and Blink Pattern-Related Differences in the Spontaneous Eyeblink Activity in Primary Gaze in Young Adult Humans. *Optometry and Vision Science*, 79(7):439–447.

Düwel, T. (2021). Generating Optimized Construction Plans for LEGO Duplo®-based Dynamic Passive Haptic Feedback. Master's thesis, Saarland University.

Ernst, M. O. and Banks, M. S. (2002). Humans Integrate Visual and Haptic Information in a Statistically Optimal Fashion. *Nature*, 415(6870):429–433.

Ernst, M. O. and Bülthoff, H. H. (2004). Merging the Senses into a Robust Percept. *Trends in Cognitive Sciences*, 8(4):162–169.

Esmaeili, S., Benda, B., and Ragan, E. D. (2020). Detection of Scaled Hand Interactions in Virtual Reality: The Effects of Motion Direction and Task Complexity. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'20, pages 453–462. IEEE.

Fairchild, M. D. (2005). Human Color Vision. In *Color Appearance Models (2nd Edition)*, chapter 1, pages 1–34. John Wiley & Sons.

Fang, C. M. and Harrison, C. (2021). Retargeted Self-Haptics for Increased Immersion in VR without Instrumentation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’21, pages 1109–1121. ACM.

Feick, M., Bateman, S., Tang, A., Miede, A., and Marquardt, N. (2020a). TanGi: Tangible Proxies For Embodied Object Exploration And Manipulation In Virtual Reality. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*, ISMAR’20, pages 195–206. IEEE.

Feick, M., Kleer, N., Tang, A., and Krüger, A. (2020b). The Virtual Reality Questionnaire Toolkit. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST’20 Adjunct, pages 68–69. ACM.

Feick, M., Kleer, N., Zenner, A., Tang, A., and Krüger, A. (2021). Visuo-Haptic Illusions for Linear Translation and Stretching Using Physical Proxies in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’21, pages 1–13. ACM.

Feick, M., Regitz, K. P., Tang, A., and Krüger, A. (2022). Designing Visuo-Haptic Illusions with Proxies in Virtual Reality: Exploration of Grasp, Movement Trajectory and Object Mass. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’22, pages 1–15. ACM.

Feuchtnner, T. and Müller, J. (2018). Ownershift: Facilitating Overhead Interaction in Virtual Reality with an Ownership-Preserving Hand Space Shift. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’18, pages 31–43. ACM.

Figueredo, P., Guo, R., Takashima, K., and Weyers, B. (2019). Escape Room in Mixed Reality: 10th Annual 3DUI Contest. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’19, pages 1407–1408. IEEE.

Franzluebbers, A. and Johnsen, K. (2018). Performance Benefits of High-Fidelity Passive Haptic Feedback in Virtual Reality Training. In *Proceedings of the ACM Symposium on Spatial User Interaction*, SUI’18, pages 16–24. ACM.

Fremerey, S., Suleman, M. S., Paracha, A. H. A., and Raake, A. (2020). Development and Evaluation of a Test Setup to Investigate Distance Differences in Immersive Virtual Environments. In *Proceedings of the International Conference on Quality of Multimedia Experience*, QoMEx’20, pages 1–4. IEEE.

Frommel, J., Sonntag, S., and Weber, M. (2017). Effects of Controller-Based Locomotion on Player Experience in a Virtual Reality Exploration Game. In *Proceedings of the International Conference on the Foundations of Digital Games*, FDG’17, pages 1–6. ACM.

Fuchs, A., Sturdee, M., and Schöning, J. (2018). Foldwatch: Using Origami-Inspired Paper Prototypes to Explore the Extension of Output Space in Smartwatches. In *Proceedings of the Nordic Conference on Human-Computer Interaction, NordiCHI'18*, pages 47–59. ACM.

Fujinawa, E., Yoshida, S., Koyama, Y., Narumi, T., Tanikawa, T., and Hirose, M. (2017). Computational Design of Hand-Held VR Controllers Using Haptic Shape Illusion. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST'17*, pages 1–10. ACM.

Gall, D. and Latoschik, M. E. (2018). The Effect of Haptic Prediction Accuracy on Presence. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'18*, pages 73–80. IEEE.

Garcia, J. F., Simeone, A. L., Higgins, M., Powell, W., and Powell, V. (2018). Inside Looking Out or Outside Looking In? An Evaluation of Visualisation Modalities to Support the Creation of a Substitutional Virtual Environment. In *Proceedings of the International Conference on Advanced Visual Interfaces, AVI'18*, pages 1–8. ACM.

Gaucher, P., Argelaguet, F., Royan, J., and Lécuyer, A. (2013). A Novel 3D Carousel based on Pseudo-Haptic Feedback and Gestural Interaction for Virtual Showcasing. In *Proceedings of the IEEE Symposium on 3D User Interfaces, 3DUI'13*, pages 55–58. IEEE.

Geslain, B., Bailly, G., Haliyo, S. D., and Duboc, C. (2021). Visuo-Haptic Illusions for Motor Skill Acquisition in Virtual Reality. In *Proceedings of the ACM Symposium on Spatial User Interaction, SUI'21*, pages 1–9. ACM.

Gibson, J. J. (1933). Adaptation, After-Effect and Contrast in the Perception of Curved Lines. *Journal of Experimental Psychology*, 16(1):1–31.

Gonzalez, E. J., Abtahi, P., and Follmer, S. (2019). Evaluating the Minimum Jerk Motion Model for Redirected Reach in Virtual Reality. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology, UIST'19 Adjunct*, pages 4–6. ACM.

Gonzalez, E. J., Abtahi, P., and Follmer, S. (2020). REACH+: Extending the Reachability of Encountered-Type Haptics Devices through Dynamic Redirection in VR. In *Proceedings of the ACM Symposium on User Interface Software and Technology, UIST'20*, pages 236–248. ACM.

Gonzalez, E. J., Chase, E. D. Z., Kotipalli, P., and Follmer, S. (2022). A Model Predictive Control Approach for Reach Redirection in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'22*, pages 1–15. ACM.

Gonzalez, E. J. and Follmer, S. (2019). Investigating the Detection of Bimanual Haptic Retargeting in Virtual Reality. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST'19, pages 1–5. ACM.

Gonzalez, E. J., Ofek, E., Gonzalez-Franco, M., and Sinclair, M. (2021a). X-Rings: A Hand-Mounted 360°Shape Display for Grasping in Virtual Reality. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'21, pages 732–742. ACM.

Gonzalez, J. F., McClelland, J. C., Teather, R. J., Figueroa, P., and Girouard, A. (2021b). Adaptic: A Shape Changing Prop with Haptic Retargeting. In *Proceedings of the ACM Symposium on Spatial User Interaction*, SUI'21, pages 1–13. ACM.

Gonzalez-Franco, M. and Lanier, J. (2017). Model of Illusions and Virtual Reality. *Frontiers in Psychology*, 8:1125.

Gray, R. (2017). Transfer of Training from Virtual to Real Baseball Batting. *Frontiers in Psychology*, 8:2183.

Grechkin, T., Thomas, J., Azmandian, M., Bolas, M., and Suma, E. (2016). Revisiting Detection Thresholds for Redirected Walking: Combining Translation and Curvature Gains. In *Proceedings of the ACM Symposium on Applied Perception*, SAP'16, pages 113–120. ACM.

Gu, X., Zhang, Y., Sun, W., Bian, Y., Zhou, D., and Kristensson, P. O. (2016). Dexmo: An Inexpensive and Lightweight Mechanical Exoskeleton for Motion Capture and Force Feedback in VR. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'16, pages 1991–1995. ACM.

Gugenheimer, J., Wolf, D., Eiriksson, E. R., Maes, P., and Rukzio, E. (2016). GyroVR: Simulating Inertia in Virtual Reality Using Head Worn Flywheels. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'16, pages 227–232. ACM.

Gurocak, H., Jayaram, S., Parrish, B., and Jayaram, U. (2003). Weight Sensation in Virtual Environments Using a Haptic Device With Air Jets. *Journal of Computing and Information Science in Engineering*, 3(2):130–135.

Han, D. T., Suhail, M., and Ragan, E. D. (2018). Evaluating Remapped Physical Reach for Hand Interactions with Passive Haptics in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1467–1476.

Han, T., Wang, S., Wang, S., Fan, X., Liu, J., Tian, F., and Fan, M. (2020). Mouillé: Exploring Wetness Illusion on Fingertips to Enhance Immersive Experience in VR. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'20, pages 1–10. ACM.

Harley, D., Tarun, A. P., Germinario, D., and Mazalek, A. (2017). Tangible VR: Diegetic Tangible Objects for Virtual Reality Narratives. In *Proceedings of the ACM Conference on Designing Interactive Systems*, DIS'17, pages 1253–1263. ACM.

Hart, S. G. and Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In Hancock, P. A. and Meshkati, N., editors, *Human Mental Workload*, volume 52 of *Advances in Psychology*, pages 139–183. North-Holland.

Hartfill, J., Gabel, J., Kruse, L., Schmidt, S., Riebandt, K., Kühn, S., and Steinicke, F. (2021). Analysis of Detection Thresholds for Hand Redirection during Mid-Air Interactions in Virtual Reality. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST'21, pages 1–10. ACM.

Hasanzadeh, S. and de la Garza, J. M. (2019). Understanding Roofer's Risk Compensatory Behavior through Passive Haptics Mixed-Reality System. In Cho, Y. K., Leite, F., Behzadan, A., and Wang, C., editors, *Computing in Civil Engineering 2019: Visualization, Information Modeling, and Simulation*, pages 137–145. American Society of Civil Engineers.

He, Z., Zhu, F., and Perlin, K. (2017). PhyShare: Sharing Physical Interaction in Virtual Reality. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST'17 Adjunct, pages 17–19. ACM.

Heilig, M. L. (1962). Sensorama Simulator. *US Patent No. 3,050,870*. August 28, 1962.

Hemmert, F., Hamann, S., Löwe, M., Zeipelt, J., and Joost, G. (2010). Weight-Shifting Mobiles: Two-Dimensional Gravitational Displays in Mobile Phones. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'10, pages 3087–3092. ACM.

Heo, S., Chung, C., Lee, G., and Wigdor, D. (2018). Thor's Hammer: An Un-grounded Force Feedback Device Utilizing Propeller-Induced Propulsive Force. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'18, pages 1–11. ACM.

Hettiarachchi, A. and Wigdor, D. (2016). Annexing Reality: Enabling Opportunistic Use of Everyday Objects as Tangible Proxies in Augmented Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'16, pages 1957–1967. ACM.

Hinchet, R., Vechev, V., Shea, H., and Hilliges, O. (2018). DextrES: Wearable Haptic Feedback for Grasping in VR via a Thin Form-Factor Electrostatic Brake. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'18, pages 901–912. ACM.

Hinckley, K., Pausch, R., Goble, J. C., and Kassell, N. F. (1994). Passive Real-World Interface Props for Neurosurgical Visualization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'94*, pages 452–458. ACM.

Hirano, Y., Kimura, A., Shibata, F., and Tamura, H. (2011). Psychophysical Influence of Mixed-Reality Visual Stimulation on Sense of Hardness. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'11*, pages 51–54. IEEE.

Hirota, K. and Hirose, M. (1993). Development of Surface Display. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'93*, pages 256–262. IEEE.

Hoffman, H. G. (1998). Physically Touching Virtual Objects Using Tactile Augmentation Enhances the Realism of Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'98*, pages 59–63. IEEE.

Hoffman, H. G. (2004). Virtual-Reality Therapy. *Scientific American*, 291(2):58–65.

Hoppe, M., Knierim, P., Kosch, T., Funk, M., Futami, L., Schneegass, S., Henze, N., Schmidt, A., and Machulla, T. (2018). VRHapticDrones: Providing Haptics in Virtual Reality through Quadcopters. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia, MUM'18*, pages 7–18. ACM.

Horie, A., Saraiji, M. Y., Kashino, Z., and Inami, M. (2021). EncounteredLimbs: A Room-Scale Encountered-Type Haptic Presentation using Wearable Robotic Arms. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'21*, pages 260–269. IEEE.

Houy, C., Fettke, P., and Loos, P. (2012). Understanding Understandability of Conceptual Models – What Are We Actually Talking about? In Atzeni, P., Cheung, D., and Ram, S., editors, *Conceptual Modeling*, pages 64–77. Springer.

Hsu, W.-T. and Lin, I.-C. (2021). Associating Real Objects with Virtual Models for VR Interaction. In *ACM SIGGRAPH Asia Posters, SA'21 Posters*, pages 1–2. ACM.

Huang, H.-Y., Ning, C.-W., Wang, P.-Y., Cheng, J.-H., and Cheng, L.-P. (2020). Haptic-Go-Round: A Surrounding Platform for Encounter-Type Haptics in Virtual Reality Experiences. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'20*, pages 1–10. ACM.

Hummel, J., Dodiya, J., Eckardt, L., Wolff, R., Gerndt, A., and Kuhlen, T. W. (2016). A Lightweight Electrotactile Feedback Device for Grasp Improvement in Immersive Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'16*, pages 39–48. IEEE.

Hummel, J., Dodiya, J., Wolff, R., Gerndt, A., and Kuhlen, T. (2013). An Evaluation of Two Simple Methods for Representing Heaviness in Immersive Virtual Environments. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, 3DUI'13, pages 87–94. IEEE.

Hösch, A. (2018). *Simulator Sickness in Fahrsimulationsumgebungen - Drei Studien zu Human Factors*. PhD thesis, Technische Universität Ilmenau.

Iesaki, A., Somada, A., Kimura, A., Shibata, F., and Tamura, H. (2008). Psychophysical Influence on Tactual Impression by Mixed-Reality Visual Stimulation. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'08, pages 265–266. IEEE.

Insko, B. E. (2001). *Passive Haptics Significantly Enhances Virtual Environments*. PhD thesis, University of North Carolina at Chapel Hill.

Issartel, P., Guéniat, F., Coquillart, S., and Ammi, M. (2015). Perceiving Mass in Mixed Reality Through Pseudo-Haptic Rendering of Newton's Third Law. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'15, pages 41–46. IEEE.

Jackson, B., Lau, T. Y., Schroeder, D., Toussaint, K. C., and Keefe, D. F. (2013). A Lightweight Tangible 3D Interface for Interactive Visualization of Thin Fiber Structures. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2802–2809.

Jauregui, D. A. G., Argelaguet, F., Olivier, A.-H., Marchal, M., Multon, F., and Lécuyer, A. (2014). Toward “Pseudo-Haptic Avatars”: Modifying the Visual Animation of Self-Avatar Can Simulate the Perception of Weight Lifting. *IEEE Transactions on Visualization and Computer Graphics*, 20(4):654–661.

Je, S., Kim, M. J., Lee, W., Lee, B., Yang, X.-D., Lopes, P., and Bianchi, A. (2019). Aero-Plane: A Handheld Force-Feedback Device That Renders Weight Motion Illusion on a Virtual 2D Plane. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'19, pages 763–775. ACM.

Je, S., Lee, H., Kim, M. J., and Bianchi, A. (2018). Wind-Blaster: A Wearable Propeller-Based Prototype That Provides Ungrounded Force-Feedback. In *ACM SIGGRAPH Emerging Technologies*, SIGGRAPH'18, pages 1–2. ACM.

Jeon, S. and Choi, S. (2009). Haptic Augmented Reality: Taxonomy and an Example of Stiffness Modulation. *Presence: Teleoperators and Virtual Environments*, 18(5):387–408.

Jerald, J. (2015). *The VR Book: Human-Centered Design for Virtual Reality*. ACM and Morgan & Claypool.

Jiang, L., Girotra, R., Cutkosky, M. R., and Ullrich, C. (2005). Reducing Error Rates with Low-Cost Haptic Feedback in Virtual Reality-Based Training Applications. In *Proceedings of the IEEE World Haptics Conference, WHC'05*, pages 420–425. IEEE.

Jones, B., Sodhi, R., Murdock, M., Mehra, R., Benko, H., Wilson, A., Ofek, E., MacIntyre, B., Raghuvanshi, N., and Shapira, L. (2014). RoomAlive: Magical Experiences Enabled by Scalable, Adaptive Projector-Camera Units. In *Proceedings of the ACM Symposium on User Interface Software and Technology, UIST'14*, pages 637–644. ACM.

Jones, B. R., Benko, H., Ofek, E., and Wilson, A. D. (2013). IllumiRoom: Peripheral Projected Illusions for Interactive Experiences. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'13*, pages 869–878. ACM.

Jones, L. A. (2000). Kinesthetic Sensing. In *Human and Machine Haptics*, pages 1–10. MIT Press.

Karunananayaka, K., Johari, N., Hariri, S., Camelia, H., Bielawski, K. S., and Cheok, A. D. (2018). New Thermal Taste Actuation Technology for Future Multisensory Virtual Reality and Internet. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1496–1505.

Kataoka, K., Yamamoto, T., Otsuki, M., Shibata, F., and Kimura, A. (2019). A New Interactive Haptic Device for Getting Physical Contact Feeling of Virtual Objects. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces, VR'19*, pages 1323–1324. IEEE.

Keller, G., Nüttgens, M., and Scheer, A.-W. (1992). Semantische Prozeßmodellierung auf der Grundlage “Ereignisgesteuerter Prozeßketten (EPK)”. In *Veröffentlichungen des Instituts für Wirtschaftsinformatik (IWI) – Heft 89 (in German)*. Saarland University.

Kennedy, R. S., Lane, N. E., Berbaum, K. S., and Lilienthal, M. G. (1993). Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness. *International Journal of Aviation Psychology*, 3(3):203–220.

Khosravi, H., Etemad, K., and Samavati, F. F. (2021). Simulating Mass in Virtual Reality using Physically-Based Hand-Object Interactions with Vibration Feedback. In *Proceedings of Graphics Interface, GI'21*, pages 241–248. Canadian Information Processing Society.

Kim, H. and Baek, M. (2021). Moment Controller: VR Controller raises Awareness of the Difference in Weight between Virtual Objects. *Archives of Design Research*, 34(2):133–151.

Kim, J., Kim, S., and Lee, J. (2022). The Effect of Multisensory Pseudo-Haptic Feedback on Perception of Virtual Weight. *IEEE Access*, 10:5129–5140.

Kingdom, F. A. A. and Prins, N. (2016a). Adaptive Methods. In *Psychophysics – A Practical Introduction (2nd Edition)*, chapter 5, pages 119–148. Academic Press.

Kingdom, F. A. A. and Prins, N. (2016b). Classifying Psychophysical Experiments. In *Psychophysics – A Practical Introduction (2nd Edition)*, chapter 2, pages 11–35. Academic Press.

Kingdom, F. A. A. and Prins, N. (2016c). Varieties of Psychophysical Procedures. In *Psychophysics – A Practical Introduction (2nd Edition)*, chapter 3, pages 37–54. Academic Press.

Kingma, I., Beek, P. J., and van Dieën, J. H. (2002). The Inertia Tensor Versus Static Moment and Mass in Perceiving Length and Heaviness of Hand-Wielded Rods. *Journal of Experimental Psychology: Human Perception and Performance*, 28(1):180–191.

Kingma, I., Van De Langenberg, R., and Beek, P. J. (2004). Which Mechanical Invariants Are Associated With the Perception of Length and Heaviness of a Nonvisible Handheld Rod? Testing the Inertia Tensor Hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, 30(2):346–354.

Kitahara, I., Nakahara, M., and Ohta, Y. (2010). Sensory Properties in Fusion of Visual/Haptic Stimuli Using Mixed Reality. In Zadeh, M. H., editor, *Advances in Haptics*, chapter 30, pages 565–582. IntechOpen.

Klatzky, R. L., Lederman, S. J., and Reed, C. (1987). There's More to Touch than Meets the Eye: The Salience of Object Attributes for Haptics with and Without Vision. *Journal of Experimental Psychology: General*, 116(4):356—369.

Knierim, P., Kosch, T., Schwind, V., Funk, M., Kiss, F., Schneegass, S., and Henze, N. (2017). Tactile Drones – Providing Immersive Tactile Feedback in Virtual Reality through Quadcopters. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'17, pages 433–436. ACM.

Knoch, S., Herbig, N., Ponpathirkoottam, S., Kosmalla, F., Staudt, P., Fettke, P., and Loos, P. (2019). Enhancing Process Data in Manual Assembly Workflows. In Daniel, F., Sheng, Q. Z., and Motahari, H., editors, *Proceedings of the Business Process Management Workshops*, BPM'18, pages 269–280. Springer.

Kohli, L. (2009). Exploiting Perceptual Illusions to Enhance Passive Haptics. In Steinicke, F. and Willemse, P., editors, *Proceedings of the Workshop on Perceptual Illusions in Virtual Environments at the IEEE Conference on Virtual Reality*, PIVE'09, pages 22–24.

Kohli, L. (2010). Redirected Touching: Warping Space to Remap Passive Haptics. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, 3DUI'10, pages 129–130. IEEE.

Kohli, L. (2013a). *Redirected Touching*. PhD thesis, University of North Carolina at Chapel Hill.

Kohli, L. (2013b). Warping Virtual Space for Low-Cost Haptic Feedback. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D'13, page 195. ACM.

Kohli, L., Burns, E., Miller, D., and Fuchs, H. (2005). Combining Passive Haptics with Redirected Walking. In *Proceedings of the International Conference on Artificial Reality and Telexistence*, ICAT'05, pages 253–254. ACM.

Kohli, L., Whitton, M. C., and Brooks, F. P. (2012). Redirected Touching: The Effect of Warping Space on Task Performance. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, 3DUI'12, pages 105–112. IEEE.

Kosmalla, F., Daiber, F., and Krüger, A. (2022). InfinityWall — Vertical Locomotion in Virtual Reality Using a Rock Climbing Treadmill. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'22, pages 1–6. ACM.

Kosmalla, F., Zenner, A., Speicher, M., Daiber, F., Herbig, N., and Krüger, A. (2017). Exploring Rock Climbing in Mixed Reality Environments. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'17, pages 1787–1793. ACM.

Kosmalla, F., Zenner, A., Tasch, C., Daiber, F., and Krüger, A. (2020). The Importance of Virtual Hands and Feet for Virtual Reality Climbing. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'20, pages 1–8. ACM.

Kovacs, R., Ofek, E., Gonzalez Franco, M., Siu, A. F., Marwecki, S., Holz, C., and Sinclair, M. (2020). Haptic PIVOT: On-Demand Handhelds in VR. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'20, pages 1046–1059. ACM.

Krekhov, A., Emmerich, K., Bergmann, P., Cmentowski, S., and Krüger, J. (2017). Self-Transforming Controllers for Virtual Reality First Person Shooters. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, CHI PLAY'17, pages 517–529. ACM.

Kriegler, H. M. (2020). A Toolkit for Hand Redirection in Virtual Reality. Bachelor's thesis, Saarland University.

Kuhlen, T. W. and Hentschel, B. (2014). Quo Vadis CAVE: Does Immersive Visualization Still Matter? *IEEE Computer Graphics and Applications*, 34(5):14–21.

Kwon, E., Kim, G. J., and Lee, S. (2009). Effects of Sizes and Shapes of Props in Tangible Augmented Reality. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*, ISMAR'09, pages 201–202. IEEE.

Laha, B., Sensharma, K., Schiffbauer, J. D., and Bowman, D. A. (2012). Effects of Immersion on Visual Analysis of Volume Data. *IEEE Transactions on Visualization and Computer Graphics*, 18(4):597–606.

Langbehn, E., Lubos, P., and Steinicke, F. (2018a). Redirected Spaces: Going Beyond Borders. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’18, pages 767–768. IEEE.

Langbehn, E. and Steinicke, F. (2018). Redirected Walking in Virtual Reality. In Lee, N., editor, *Encyclopedia of Computer Graphics and Games*, pages 1–11. Springer.

Langbehn, E., Steinicke, F., Lappe, M., Welch, G. F., and Bruder, G. (2018b). In the Blink of an Eye: Leveraging Blink-Induced Suppression for Imperceptible Position and Orientation Redirection in Virtual Reality. *ACM Transactions on Graphics*, 37(4):66:1–66:11.

Laugwitz, B., Held, T., and Schrepp, M. (2008). Construction and Evaluation of a User Experience Questionnaire. In Holzinger, A., editor, *Proceedings of the Symposium of the Austrian HCI and Usability Engineering Group*, USAB’08, pages 63–76. Springer.

Le Goc, M., Kim, L. H., Parsaei, A., Fekete, J.-D., Dragicevic, P., and Follmer, S. (2016). Zoooids: Building Blocks for Swarm User Interfaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’16, pages 97–109. ACM.

Lebrun, F., Haliyo, S., and Baily, G. (2021). A Trajectory Model for Desktop-Scale Hand Redirection in Virtual Reality. In Ardito, C., Lanzilotti, R., Malizia, A., Petrie, H., Piccinno, A., Desolda, G., and Inkpen, K., editors, *Proceedings of the IFIP Conference on Human-Computer Interaction*, INTERACT’21, pages 105–124. Springer.

Lederman, S. J. and Klatzky, R. L. (2009). Haptic Perception: A Tutorial. *Attention, Perception, & Psychophysics*, 71(7):1439–1459.

Lee, C.-G., Dunn, G. L., Oakley, I., and Ryu, J. (2020). Visual Guidance for a Spatial Discrepancy Problem of in Encountered-Type Haptic Display. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 50(4):1384–1394.

Lee, C.-J., Tsai, H.-R., and Chen, B.-Y. (2021). HairTouch: Providing Stiffness, Roughness and Surface Height Differences Using Reconfigurable Brush Hairs on a VR Controller. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’21, pages 1–13. ACM.

Lee, J. C., Hudson, S. E., and Tse, E. (2008). Foldable Interactive Displays. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’08, pages 287–290. ACM.

Lee, Y., Jang, I., and Lee, D. (2015). Enlarging Just Noticeable Differences of Visual-Proprioceptive Conflict in VR using Haptic Feedback. In *Proceedings of the IEEE World Haptics Conference*, WHC'15, pages 19–24. IEEE.

Leek, M. R. (2001). Adaptive Procedures in Psychophysical Research. *Perception & Psychophysics*, 63(8):1279–1292.

Leigh, R. J. and Zee, D. S. (2015). *The Neurology of Eye Movements (5th Edition)*. Oxford University Press.

Lim, W. N., Yap, K. M., Lee, Y., Wee, C., and Yen, C. C. (2021). A Systematic Review of Weight Perception in Virtual Reality: Techniques, Challenges, and Road Ahead. *IEEE Access*, 9:163253–163283.

Lindeman, R. W. (1999). *Bimanual Interaction, Passive-Haptic Feedback, 3D Widget Representation, and Simulated Surface Constraints for Interaction In Immersive Virtual Environments*. PhD thesis, The George Washington University.

Lindeman, R. W., Sibert, J. L., and Hahn, J. K. (1999). Hand-Held Windows: Towards Effective 2D Interaction in Immersive Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'99, pages 205–212. IEEE.

Liu, Y., Hashimoto, T., Yoshida, S., Narumi, T., Tanikawa, T., and Hirose, M. (2019). ShapeSense: A 2D Shape Rendering VR Device with Moving Surfaces That Controls Mass Properties and Air Resistance. In *ACM SIGGRAPH Emerging Technologies*, SIGGRAPH'19, pages 1–2. ACM.

Lohse, A. L., Kjær, C. K., Hamulic, E., Lima, I. G. A., Jensen, T. H., Bruni, L. E., and Nilsson, N. C. (2019). Leveraging Change Blindness for Haptic Remapping in Virtual Environments. In *Proceedings of the IEEE Workshop on Everyday Virtual Reality*, WEVR'19, pages 1–5. IEEE.

Lok, B., Naik, S., Whittom, M., and Brooks, F. P. (2003). Effects of Handling Real Objects and Self-Avatar Fidelity on Cognitive Task Performance and Sense of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 12(6):615–628.

Long, B., Seah, S. A., Carter, T., and Subramanian, S. (2014). Rendering Volumetric Haptic Shapes in Mid-Air Using Ultrasound. *ACM Transactions on Graphics*, 33(6):181:1–181:10.

Lopes, P., Ion, A., and Baudisch, P. (2015). Impacto: Simulating Physical Impact by Combining Tactile Stimulation with Electrical Muscle Stimulation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'15, pages 11–19. ACM.

Lopes, P., You, S., Cheng, L.-P., Marwecki, S., and Baudisch, P. (2017). Providing Haptics to Walls & Heavy Objects in Virtual Reality by Means of Electrical Muscle Stimulation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'17, pages 1471–1482. ACM.

Lécuyer, A. (2009). Simulating Haptic Feedback Using Vision: A Survey of Research and Applications of Pseudo-Haptic Feedback. *Presence: Teleoperators and Virtual Environments*, 18(1):39–53.

Lécuyer, A. (2017). Playing with Senses in VR: Alternate Perceptions Combining Vision and Touch. *IEEE Computer Graphics and Applications*, 37(1):20–26.

Lécuyer, A., Burkhardt, J.-M., and Etienne, L. (2004). Feeling Bumps and Holes without a Haptic Interface: The Perception of Pseudo-Haptic Textures. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'04, pages 239–246. ACM.

Lécuyer, A., Coquillart, S., Kheddar, A., Richard, P., and Coiffet, P. (2000). Pseudo-Haptic Feedback: Can Isometric Input Devices Simulate Force Feedback? In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'00, pages 83–90. IEEE.

MacLean, K. E. (2000). Designing with Haptic Feedback. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 1 of *ICRA'00*, pages 783–788. IEEE.

Maehigashi, A., Sasada, A., Matsumuro, M., Shibata, F., Kimura, A., and Niida, S. (2021). Virtual Weight Illusion: Weight Perception of Virtual Objects Using Weight Illusions. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'21, pages 1–6. ACM.

Makhsadov, A., Degraen, D., Zenner, A., Kosmalla, F., Mushkina, K., and Krüger, A. (2022). VRySmart: A Framework for Embedding Smart Devices in Virtual Reality. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'22, pages 1–8. ACM.

Martin, N. A. A., Mittelstädt, V., Prieur, M., Stark, R., and Bär, T. (2013). Passive Haptic Feedback for Manual Assembly Simulation. *Procedia CIRP*, 7:509–514.

Marwecki, S., Wilson, A. D., Ofek, E., Gonzalez Franco, M., and Holz, C. (2019). Mise-Unseen: Using Eye Tracking to Hide Virtual Reality Scene Changes in Plain Sight. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'19, pages 777–789. ACM.

Massie, T. H. and Salisbury, J. K. (1994). 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*, pages 295–301.

Matsumoto, K., Narumi, T., Ban, Y., Yanase, Y., Tanikawa, T., and Hirose, M. (2019). Unlimited Corridor: A Visuo-Haptic Redirection System. In *Proceedings of the ACM SIGGRAPH International Conference on Virtual Reality Continuum and Its Applications in Industry*, VRCAI'19, pages 1–9. ACM.

Matsuoka, Y., Allin, S. J., and Klatzky, R. L. (2002). The Tolerance for Visual Feedback Distortions in a Virtual Environment. *Physiology & Behavior*, 77(4):651–655.

Matthews, B. J. and Smith, R. T. (2019). Head Gaze Target Selection for Redirected Interaction. In *ACM SIGGRAPH Asia XR*, SA'19, pages 13–14. ACM.

Matthews, B. J., Thomas, B. H., Von Itzstein, S., and Smith, R. T. (2019). Remapped Physical-Virtual Interfaces with Bimanual Haptic Retargeting. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'19, pages 19–27. IEEE.

Matthews, B. J., Thomas, B. H., Von Itzstein, S., and Smith, R. T. (2021). Adaptive Reset Techniques for Haptic Retargeted Interaction. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–13. Early Access.

Matthews, B. J., Thomas, B. H., Von Itzstein, S., and Smith, R. T. (2022). Shape Aware Haptic Retargeting for Accurate Hand Interactions. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'22, pages 625–634. IEEE.

McClelland, J. C., Teather, R. J., and Girouard, A. (2017). Haptobend: Shape-Changing Passive Haptic Feedback in Virtual Reality. In *Proceedings of the ACM Symposium on Spatial User Interaction*, SUI'17, pages 82–90. ACM.

McNeely, W. A. (1993). Robotic Graphics: A New Approach to Force Feedback for Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'93, pages 336–341. IEEE.

Meehan, M., Insko, B., Whitton, M., and Brooks, F. P. (2002). Physiological Measures of Presence in Stressful Virtual Environments. *ACM Transactions on Graphics*, 21(3):645–652.

Mercado, V. R., Howard, T., Si-Mohammed, H., Argelaguet, F., and Lécuyer, A. (2021a). Alfred: The Haptic Butler – On-Demand Tangibles for Object Manipulation in Virtual Reality using an ETHD. In *Proceedings of the IEEE World Haptics Conference*, WHC'21, pages 373–378. IEEE.

Mercado, V. R., Marchal, M., and Lécuyer, A. (2020). Design and Evaluation of Interaction Techniques Dedicated to Integrate Encountered-Type Haptic Displays in Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'20, pages 230–238. IEEE.

Mercado, V. R., Marchal, M., and Lécuyer, A. (2021b). ENTROPiA: Towards Infinite Surface Haptic Displays in Virtual Reality Using Encountered-Type Rotating Props. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2237–2243.

Mercado, V. R., Marchal, M., and Lécuyer, A. (2021c). “Haptics On-Demand”: A Survey on Encountered-Type Haptic Displays. *IEEE Transactions on Haptics*, 14(3):449–464.

Milgram, P. and Kishino, F. (1994). A Taxonomy of Mixed Reality Visual Displays. *IEICE Transactions on Information and Systems*, E77-D(12):1321–1329.

Minamizawa, K., Fukamachi, S., Kajimoto, H., Kawakami, N., and Tachi, S. (2007). Gravity Grabber: Wearable Haptic Display to Present Virtual Mass Sensation. In *ACM SIGGRAPH Emerging Technologies*, SIGGRAPH’07, pages 8–11. ACM.

Montano Murillo, R. A., Subramanian, S., and Martinez Plasencia, D. (2017). Erg-O: Ergonomic Optimization of Immersive Virtual Environments. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’17, pages 759–771. ACM.

Murakami, T., Person, T., Fernando, C. L., and Minamizawa, K. (2017). Altered Touch: Miniature Haptic Display with Force, Thermal and Tactile Feedback for Augmented Haptics. In *ACM SIGGRAPH Emerging Technologies*, SIGGRAPH’17, pages 1–2. ACM.

Murray, B. C. M., Peele, B. N., Xu, P., Spjut, J., Shapira, O., Luebke, D., and Shepherd, R. F. (2018). A Variable Shape and Variable Stiffness Controller for Haptic Virtual Interactions. In *Proceedings of the IEEE International Conference on Soft Robotics*, RoboSoft’18, pages 264–269. IEEE.

Myles, K. and Kalb, J. T. (2010). Guidelines for Head Tactile Communication. Technical report, Defense Technical Information Center.

Mündler, T., Bonfert, M., Reinschlüssel, A. V., Malaka, R., and Döring, T. (2022). Haptic Fidelity Framework: Defining the Factors of Realistic Haptic Feedback for Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’22, pages 1–17. ACM.

Mündler, T., Reinschlüssel, A. V., Drewes, S., Wenig, D., Döring, T., and Malaka, R. (2019). Does It Feel Real? Using Tangibles with Different Fidelities to Build and Explore Scenes in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’19, pages 1–12. ACM.

Nakahara, M., Kitahara, I., and Ohta, Y. (2007). Sensory Property in Fusion of Visual/Haptic Cues by Using Mixed Reality. In *Proceedings of the IEEE World Haptics Conference*, WHC’07, pages 565–566. IEEE.

National Research Council (1995). *Virtual Reality: Scientific and Technological Challenges*. National Academies Press.

Nguyen, A. and Kunz, A. (2018). Discrete Scene Rotation during Blinks and Its Effect on Redirected Walking Algorithms. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST'18, pages 1–10. ACM.

Niiyama, R., Yao, L., and Ishii, H. (2014). Weight and Volume Changing Device with Liquid Metal Transfer. In *Proceedings of the ACM Conference on Tangible, Embedded and Embodied Interaction*, TEI'14, pages 49–52. ACM.

Nilsson, N. C., Peck, T., Bruder, G., Hodgson, E., Serafin, S., Whitton, M., Steinicke, F., and Rosenberg, E. S. (2018). 15 Years of Research on Redirected Walking in Immersive Virtual Environments. *IEEE Computer Graphics and Applications*, 38(2):44–56.

Nilsson, N. C., Zenner, A., and Simeone, A. L. (2021a). Propping Up Virtual Reality With Haptic Proxies. *IEEE Computer Graphics and Applications*, 41(05):104–112.

Nilsson, N. C., Zenner, A., Simeone, A. L., Degraen, D., and Daiber, F. (2021b). Haptic Proxies for Virtual Reality: Success Criteria and Taxonomy. In *Proceedings of the Workshop on Everyday Proxy Objects for Virtual Reality at the ACM Conference on Human Factors in Computing Systems*, EPO4VR'21, pages 1–5.

Object Management Group (2013). Business Process Model and Notation (BPMN) – Specification Version 2.0.2. Technical report, Object Management Group, Inc.

Ogawa, N., Narumi, T., and Hirose, M. (2021). Effect of Avatar Appearance on Detection Thresholds for Remapped Hand Movements. *IEEE Transactions on Visualization and Computer Graphics*, 27(7):3182–3197.

O'Regan, J. K., Deubel, H., Clark, J. J., and Rensink, R. A. (2000). Picture Changes During Blinks: Looking Without Seeing and Seeing Without Looking. *Visual Cognition*, 7(1-3):191–211.

Orlosky, J., Itoh, Y., Ranchet, M., Kiyokawa, K., Morgan, J., and Devos, H. (2017). Emulation of Physician Tasks in Eye-tracked Virtual Reality for Remote Diagnosis of Neurodegenerative Disease. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1302–1311.

Pair, J., Neumann, U., Piepol, D., and Swartout, B. (2003). FlatWorld: Combining Hollywood Set-Design Techniques with VR. *IEEE Computer Graphics and Applications*, 23(1):12–15.

Palmerius, K. L., Johansson, D., Höst, G., and Schönborn, K. (2014). An Analysis of the Influence of a Pseudo-haptic Cue on the Haptic Perception of Weight. In Auvray, M. and Duriez, C., editors, *Proceedings of the EuroHaptics Conference*, EuroHaptics'14, pages 117–125. Springer.

Park, C., Kim, J., and Choi, S. (2021). Length Perception Model for Handheld Controllers: The Effects of Diameter and Inertia. *IEEE Transactions on Haptics*, 14(2):310–315.

Patel, S., Henderson, R., Bradley, L., Galloway, B., and Hunter, L. (1991). Effect of Visual Display Unit Use on Blink Rate and Tear Stability. *Optometry and Vision Science*, 68(11):888–892.

Patras, C., Cibulskis, M., and Nilsson, N. C. (2022). Body Warping Versus Change Blindness Remapping: A Comparison of Two Approaches to Repurposing Haptic Proxies for Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'22, pages 205–212. IEEE.

Peiris, R. L., Peng, W., Chen, Z., Chan, L., and Minamizawa, K. (2017). ThermoVR: Exploring Integrated Thermal Haptic Feedback with Head Mounted Displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'17, pages 5452–5456. ACM.

Pezent, E., Israr, A., Samad, M., Robinson, S., Agarwal, P., Benko, H., and Colonese, N. (2019). Tasbi: Multisensory Squeeze and Vibrotactile Wrist Haptics for Augmented and Virtual Reality. In *Proceedings of the IEEE World Haptics Conference*, WHC'19, pages 1–6. IEEE.

Proske, U. and Gandevia, S. C. (2012). The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force. *Physiological Reviews*, 92(4):1651–1697.

Provancher, W. R. (2014). Creating Greater VR Immersion by Emulating Force Feedback with Ungrounded Tactile Feedback. *IQT Quarterly*, 6(2):18–21.

Rakkolainen, I., Freeman, E., Sand, A., Raisamo, R., and Brewster, S. (2021). A Survey of Mid-Air Ultrasound Haptics and Its Applications. *IEEE Transactions on Haptics*, 14(1):2–19.

Ranasinghe, N., Jain, P., Karwita, S., Tolley, D., and Do, E. Y.-L. (2017). Ambiotherm: Enhancing Sense of Presence in Virtual Reality by Simulating Real-World Environmental Conditions. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'17, pages 1731–1742. ACM.

Razzaque, S. (2005). *Redirected Walking*. PhD thesis, University of North Carolina at Chapel Hill.

Razzaque, S., Kohn, Z., and Whitton, M. C. (2001). Redirected Walking. In *Proceedings of the Eurographics Conference – Short Presentations*, EG'01, pages 1–6. Eurographics Association.

Recker, J. C. and Dreiling, A. (2011). The Effects of Content Presentation Format and User Characteristics on Novice Developers' Understanding of Process Models. *Communications of the Association for Information Systems*, 28(6):65–84.

Regitz, K. P. (2020). Leveraging Blink Suppression for Hand Redirection in Virtual Reality. Bachelor's thesis, Saarland University.

Rekimoto, J. (2013). Traxion: A Tactile Interaction Device with Virtual Force Sensation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'13, pages 427–432. ACM.

Rietzler, M., Geiselhart, F., Gugenheimer, J., and Rukzio, E. (2018). Breaking the Tracking: Enabling Weight Perception Using Perceivable Tracking Offsets. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'18, pages 1–12. ACM.

Rietzler, M., Plaumann, K., Kränzle, T., Erath, M., Stahl, A., and Rukzio, E. (2017). VaiR: Simulating 3D Airflows in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'17, pages 5669–5677. ACM.

Romano, J. M. and Kuchenbecker, K. J. (2009). The AirWand: Design and Characterization of a Large-Workspace Haptic Device. In *Proceedings of the IEEE International Conference on Robotics and Automation*, ICRA'09, pages 1461–1466. IEEE.

Roudaut, A., Karnik, A., Löchtefeld, M., and Subramanian, S. (2013). Morphées: Toward High “Shape Resolution” in Self-Actuated Flexible Mobile Devices. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'13, pages 593–602. ACM.

Sagheb, S., Liu, F. W., Bahremand, A., Kidane, A., and LiKamWa, R. (2019). SWISH: A Shifting-Weight Interface of Simulated Hydrodynamics for Haptic Perception of Virtual Fluid Vessels. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'19, pages 751–761. ACM.

Sait, M. S. M. Y., Sargunam, S. P., Han, D. T., and Ragan, E. D. (2018). Physical Hand Interaction for Controlling Multiple Virtual Objects in Virtual Reality. In *Proceedings of the International Workshop on Interactive and Spatial Computing*, IWISC'18, pages 64–74. ACM.

Samad, M., Gatti, E., Hermes, A., Benko, H., and Parise, C. (2019). Pseudo-Haptic Weight: Changing the Perceived Weight of Virtual Objects By Manipulating Control-Display Ratio. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'19, pages 1–13. ACM.

Sasaki, T., Hartanto, R. S., Liu, K.-H., Tsuchiya, K., Hiyama, A., and Inami, M. (2018). Leviopole: Mid-Air Haptic Interactions Using Multirotor. In *ACM SIGGRAPH Emerging Technologies*, SIGGRAPH'18, pages 1–2. ACM.

Sato, M. (2002). Development of String-Based Force Display: SPIDAR. In *Proceedings of the International Conference on Virtual Systems and Multimedia*, VSMM'02, pages 1034–1039.

Schneider, H., Scheibe, R., Hochstrate, J., and Fröhlich, B. (2005). Haptic Feedback Through Weight Shifting. In *Proceedings of the Workshop on New Directions in 3D User Interfaces at the IEEE Conference on Virtual Reality*, pages 1–4.

Schrepp, M. (2019). User Experience Questionnaire Handbook – Version 7. Technical report, UEQ Team.

Schrepp, M., Hinderks, A., and Thomaschewski, J. (2017). Design and Evaluation of a Short Version of the User Experience Questionnaire (UEQ-S). *International Journal of Interactive Multimedia and Artificial Intelligence*, 4(6):103–108.

Schulz, P., Alexandrovsky, D., Putze, F., Malaka, R., and Schöning, J. (2019). The Role of Physical Props in VR Climbing Environments. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’19, pages 1–13. ACM.

Schwind, V., Knierim, P., Tasci, C., Franczak, P., Haas, N., and Henze, N. (2017). “These Are Not My Hands!”: Effect of Gender on the Perception of Avatar Hands in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’17, pages 1577–1582. ACM.

Shapira, L. and Freedman, D. (2016). Reality Skins: Creating Immersive and Tactile Virtual Environments. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*, ISMAR’16, pages 115–124. IEEE.

Sharp, H., Preece, J., and Rogers, Y. (2019). *Interaction Design: Beyond Human-Computer Interaction*. John Wiley & Sons.

Shepard, D. (1968). A Two-Dimensional Interpolation Function for Irregularly-Spaced Data. In *Proceedings of the ACM National Conference*, ACM’68, pages 517–524. ACM.

Shigeyama, J., Hashimoto, T., Yoshida, S., Narumi, T., Tanikawa, T., and Hirose, M. (2019). Transcalibur: A Weight Shifting Virtual Reality Controller for 2D Shape Rendering Based on Computational Perception Model. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’19, pages 1–11. ACM.

Shimizu, S., Hashimoto, T., Yoshida, S., Matsumura, R., Narumi, T., and Kuzuoka, H. (2021). Unident: Providing Impact Sensations on Handheld Objects via High-Speed Change of the Rotational Inertia. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’21, pages 11–20. IEEE.

Simeone, A. L., Nilsson, N. C., Zenner, A., Speicher, M., and Daiber, F. (2020). The Space Bender: Supporting Natural Walking via Overt Manipulation of the Virtual Environment. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’20, pages 598–606. IEEE.

Simeone, A. L., Velloso, E., and Gellersen, H. (2015). Substitutional Reality: Using the Physical Environment to Design Virtual Reality Experiences. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'15*, pages 3307–3316. ACM.

Simons, D. J. and Levin, D. T. (1997). Change Blindness. *Trends in Cognitive Sciences*, 1(7):261–267.

Sinclair, M., Ofek, E., Gonzalez-Franco, M., and Holz, C. (2019). CapstanCrunch: A Haptic VR Controller with User-Supplied Force Feedback. In *Proceedings of the ACM Symposium on User Interface Software and Technology, UIST'19*, pages 815–829. ACM.

Siu, A. F., Gonzalez, E. J., Yuan, S., Ginsberg, J. B., and Follmer, S. (2018). ShapeShift: 2D Spatial Manipulation and Self-Actuation of Tabletop Shape Displays for Tangible and Haptic Interaction. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'18*, pages 1–13. ACM.

Slater, M. (2009). Place Illusion and Plausibility Can Lead to Realistic Behaviour in Immersive Virtual Environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1535):3549–3557.

Slater, M. and Steed, A. (2000). A Virtual Presence Counter. *Presence: Teleoperators and Virtual Environments*, 9(5):413–434.

Slater, M., Usoh, M., and Steed, A. (1994). Depth of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 3(2):130–144.

Slater, M., Usoh, M., and Steed, A. (1995). Taking Steps: The Influence of a Walking Technique on Presence in Virtual Reality. *ACM Transactions on Computer-Human Interaction*, 2(3):201–219.

Slater, M. and Wilbur, S. (1997). A Framework for Immersive Virtual Environments (FIVE): Speculations on the Role of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 6(6):603–616.

Sodhi, R., Poupyrev, I., Glisson, M., and Israr, A. (2013). AIREAL: Interactive Tactile Experiences in Free Air. *ACM Transactions on Graphics*, 32(4):134:1–134:10.

Sousa, M., Mendes, D., Paulo, S., Matela, N., Jorge, J., and Lopes, D. S. o. (2017). VRROOM: Virtual Reality for Radiologists in the Reading Room. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI'17*, pages 4057–4062. ACM.

Speicher, M., Ehrlich, J., Gentile, V., Degraen, D., Sorce, S., and Krüger, A. (2019). Pseudo-Haptic Controls for Mid-Air Finger-Based Menu Interaction. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems, CHI EA'19*, pages 1–6. ACM.

Spillmann, J., Tuchschen, S., and Harders, M. (2013). Adaptive Space Warping to Enhance Passive Haptics in an Arthroscopy Surgical Simulator. *IEEE Transactions on Visualization and Computer Graphics*, 19(4):626–633.

Sra, M., Garrido-Jurado, S., Schmandt, C., and Maes, P. (2016). Procedurally Generated Virtual Reality from 3D Reconstructed Physical Space. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST’16, pages 191–200. ACM.

Srinivasan, M. A. and Basdogan, C. (1997). Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges. *Computers & Graphics*, 21(4):393–404.

Stauffert, J.-P., Korwisi, K., Niebling, F., and Latoschik, M. E. (2021). Ka-Boom!!! Visually Exploring Latency Measurements for XR. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA’21, pages 1–9. ACM.

Steed, A., Friston, S., Pawar, V., and Swapp, D. (2020). Docking Haptics: Extending the Reach of Haptics by Dynamic Combinations of Grounded and Worn Devices. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST’20, pages 1–11. ACM.

Steinicke, F., Bruder, G., Hinrichs, K., and Willemse, P. (2010a). Change Blindness Phenomena for Stereoscopic Projection Systems. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’10, pages 187–194. IEEE.

Steinicke, F., Bruder, G., Jerald, J., Frenz, H., and Lappe, M. (2010b). Estimation of Detection Thresholds for Redirected Walking Techniques. *IEEE Transactions on Visualization and Computer Graphics*, 16(1):17–27.

Steinicke, F., Bruder, G., Kohli, L., Jerald, J., and Hinrichs, K. (2008). Taxonomy and Implementation of Redirection Techniques for Ubiquitous Passive Haptic Feedback. In *Proceedings of the International Conference on Cyberworlds*, CW’08, pages 217–223. IEEE.

Stellmacher, C. (2020). Weight Perception in VR: A Novel Haptic VR Controller with Adaptive Trigger Button Resistance. Master’s thesis, University of Bremen.

Stellmacher, C. (2021). Haptic-Enabled Buttons Through Adaptive Trigger Resistance. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, VRW’21, pages 201–204. IEEE.

Stellmacher, C., Bonfert, M., Kruijff, E., and Schöning, J. (2022). Triggermuscle: Exploring Weight Perception for Virtual Reality Through Adaptive Trigger Resistance in a Haptic VR Controller. *Frontiers in Virtual Reality*, 2:1–18.

Steuer, J. (1992). Defining Virtual Reality: Dimensions Determining Telepresence. *Journal of Communication*, 42(4):73–93.

Strandholt, P. L., Dogaru, O. A., Nilsson, N. C., Nordahl, R., and Serafin, S. (2020). Knock on Wood: Combining Redirected Touching and Physical Props for Tool-Based Interaction in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’20, pages 1–13. ACM.

Strasnick, E., Holz, C., Ofek, E., Sinclair, M., and Benko, H. (2018). Haptic Links: Bimanual Haptics for Virtual Reality Using Variable Stiffness Actuation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’18, pages 1–12. ACM.

Strohmeier, P., Güngör, S., Herres, L., Gudea, D., Fruchard, B., and Steimle, J. (2020). BARefoot: Generating Virtual Materials Using Motion Coupled Vibration in Shoes. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST’20, pages 579–593. ACM.

Strohmeier, P. and Hornbæk, K. (2017). Generating Haptic Textures with a Vibrotactile Actuator. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’17, pages 4994–5005. ACM.

Suhail, M., Sargunam, S. P., Han, D. T., and Ragan, E. D. (2017). Redirected Reach in Virtual Reality: Enabling Natural Hand Interaction at Multiple Virtual Locations with Passive Haptics. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, 3DUI’17, pages 245–246. IEEE.

Suma, E. A., Clark, S., Finkelstein, S. L., and Wartell, Z. (2010). Exploiting Change Blindness to Expand Walkable Space in a Virtual Environment. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’10, pages 305–306. IEEE.

Suma, E. A., Clark, S., Krum, D., Finkelstein, S., Bolas, M., and Warte, Z. (2011a). Leveraging Change Blindness for Redirection in Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR’11, pages 159–166. IEEE.

Suma, E. A., Krum, D. M., and Bolas, M. (2011b). Redirection on Mixed Reality Walking Surfaces. In Steinicke, F. and Willemsen, P., editors, *Proceedings of the Workshop on Perceptual Illusions in Virtual Environments at the IEEE Conference on Virtual Reality*, PIVE’11, pages 33–35.

Suma, E. A., Lipps, Z., Finkelstein, S., Krum, D. M., and Bolas, M. (2012). Impossible Spaces: Maximizing Natural Walking in Virtual Environments with Self-Overlapping Architecture. *IEEE Transactions on Visualization and Computer Graphics*, 18(4):555–564.

Sun, Q., Patney, A., Wei, L.-Y., Shapira, O., Lu, J., Asente, P., Zhu, S., McGuire, M., Luebke, D., and Kaufman, A. (2018). Towards Virtual Reality Infinite Walking: Dynamic Saccadic Redirection. *ACM Transactions on Graphics*, 37(4):67:1–67:13.

Sun, Y., Yoshida, S., Narumi, T., and Hirose, M. (2019). PaCaPa: A Handheld VR Device for Rendering Size, Shape, and Stiffness of Virtual Objects in Tool-Based Interactions. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’19, pages 1–12. ACM.

Sutherland, I. E. (1965). The Ultimate Display. In *Proceedings of the Congress of the International Federation for Information Processing (IFIP)*, volume 2, pages 506–508.

Sutherland, I. E. (1968). A Head-Mounted Three Dimensional Display. In *Proceedings of the AFIPS Fall Joint Computing Conference*, AFIPS’68, pages 757–764. ACM.

Suzuki, R., Hedayati, H., Zheng, C., Bohn, J. L., Szafir, D., Do, E. Y.-L., Gross, M. D., and Leithinger, D. (2020). RoomShift: Room-Scale Dynamic Haptics for VR with Furniture-Moving Swarm Robots. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI’20, pages 1–11. ACM.

Suzuki, Y. and Kobayashi, M. (2005). Air Jet Driven Force Feedback in Virtual Reality. *IEEE Computer Graphics and Applications*, 25(1):44–47.

Suzuki, Y., Kobayashi, M., and Ishibashi, S. (2002). Design of Force Feedback Utilizing Air Pressure toward Untethered Human Interface. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA’02, pages 808–809. ACM.

Swanson, D. K. (2003). *Implementation of Arbitrary Path Constraints using Dissipative Passive Haptic Displays*. PhD thesis, Georgia Institute of Technology.

Swindells, C., Uden, A., and Sang, T. (2003). TorqueBAR: An Ungrounded Haptic Feedback Device. In *Proceedings of the International Conference on Multimodal Interfaces*, ICMI’03, pages 52–59. ACM.

Tachi, S., Maeda, T., Hirata, R., and Hoshino, H. (1994). A Construction Method of Virtual Haptic Space. In *Proceedings of the International Conference on Artificial Reality and Telexistence*, ICAT’94, pages 131–138.

Tanaka, Y., Horie, A., and Chen, X. A. (2020). DualVib: Simulating Haptic Sensation of Dynamic Mass by Combining Pseudo-Force and Texture Feedback. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, VRST’20, pages 1–10. ACM.

Taylor, C. and Cosker, D. (2020). Interacting with Real Objects in Virtual Worlds. In Magnor, M. and Sorkine-Hornung, A., editors, *Real VR – Immersive Digital Reality: How to Import the Real World into Head-Mounted Immersive Displays*, pages 337–353. Springer.

Teng, S.-Y., Kuo, T.-S., Wang, C., Chiang, C.-h., Huang, D.-Y., Chan, L., and Chen, B.-Y. (2018). PuPoP: Pop-up Prop on Palm for Virtual Reality. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'18, pages 5–17. ACM.

Teng, S.-Y., Lin, C.-L., Chiang, C.-h., Kuo, T.-S., Chan, L., Huang, D.-Y., and Chen, B.-Y. (2019). TilePoP: Tile-Type Pop-up Prop for Virtual Reality. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'19, pages 639–649. ACM.

Tiator, M., Geiger, C., Dewitz, B., Fischer, B., Gerhardt, L., Nowottnik, D., and Preu, H. (2018). Venga! Climbing in Mixed Reality. In *Proceedings of the Superhuman Sports Design Challenge*, SHS'18, pages 1–8. ACM.

Tsai, C.-Y., Tsai, I.-L., Lai, C.-J., Chow, D., Wei, L., Cheng, L.-P., and Chen, M. Y. (2022a). AirRacket: Perceptual Design of Ungrounded, Directional Force Feedback to Improve Virtual Racket Sports Experiences. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'22, pages 1–15. ACM.

Tsai, H.-R., Hung, C.-W., Wu, T.-C., and Chen, B.-Y. (2020). ElastOscillation: 3D Multilevel Force Feedback for Damped Oscillation on VR Controllers. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'20, pages 1–12. ACM.

Tsai, H.-R., Tsai, C., Liao, Y.-S., Chiang, Y.-T., and Zhang, Z.-Y. (2022b). FingerX: Rendering Haptic Shapes of Virtual Objects Augmented by Real Objects Using Extendable and Withdrawable Supports on Fingers. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'22, pages 1–14. ACM.

Tseng, W.-J., Bonnail, E., McGill, M., Khamis, M., Lecolinet, E., Huron, S., and Gugenheimer, J. (2022). The Dark Side of Perceptual Manipulations in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'22, pages 1–15. ACM.

Turvey, M. T. (1996). Dynamic Touch. *American Psychologist*, 51(11):1134–1152.

Usoh, M., Arthur, K., Whitton, M. C., Bastos, R., Steed, A., Slater, M., and Brooks, F. P. (1999). Walking > Walking-in-Place > Flying, in Virtual Environments. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH'99, pages 359–364. ACM Press/Addison-Wesley Publishing Co.

Van Baren, J. and IJsselsteijn, W. (2004). Measuring Presence: A Guide to Current Measurement Approaches. Technical report, Deliverable of the OmniPres project IST-2001-39237.

van Beers, R. J., Wolpert, D. M., and Haggard, P. (2002). When Feeling Is More Important Than Seeing in Sensorimotor Adaptation. *Current Biology*, 12(10):834–837.

Van der Linde, R. Q., Lammertse, P., Frederiksen, E., and Ruiter, B. (2002). The HapticMaster, a New High-Performance Haptic Interface. In *Proceedings of the EuroHaptics Conference*, EuroHaptics'02, pages 1–5. EuroHaptics Society.

VanderWerf, F., Brassinga, P., Reits, D., Aramideh, M., and Ongerboer de Visser, B. (2003). Eyelid Movements: Behavioral Studies of Blinking in Humans Under Different Stimulus Conditions. *Journal of Neurophysiology*, 89(5):2784–2796.

Volkmann, F. C. (1986). Human Visual Suppression. *Vision Research*, 26(9):1401–1416.

Vonach, E., Gatterer, C., and Kaufmann, H. (2017). VRRobot: Robot Actuated Props in an Infinite Virtual Environment. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'17, pages 74–83. IEEE.

Wang, C.-H., Hsieh, C.-Y., Yu, N.-H., Bianchi, A., and Chan, L. (2019a). HapticSphere: Physical Support To Enable Precision Touch Interaction in Mobile Mixed-Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'19, pages 331–339. IEEE.

Wang, D., Guo, Y., Liu, S., Zhang, Y., Xu, W., and Xiao, J. (2019b). Haptic Display for Virtual Reality: Progress and Challenges. *Virtual Reality & Intelligent Hardware*, 1(2):136–162.

Wang, D., Ohnishi, K., and Xu, W. (2020a). Multimodal Haptic Display for Virtual Reality: A Survey. *IEEE Transactions on Industrial Electronics*, 67(1):610–623.

Wang, X., Monteiro, D., Lee, L.-H., Hui, P., and Liang, H.-N. (2022). VibroWeight: Simulating Weight and Center of Gravity Changes of Objects in Virtual Reality for Enhanced Realism. In *Proceedings of the IEEE Haptics Symposium*, HAPTICS'22, pages 1–7. IEEE.

Wang, Y., Chen, Z., Li, H., Cao, Z., Luo, H., Zhang, T., Ou, K., Raiti, J., Yu, C., Patel, S., and Shi, Y. (2020b). MoveVR: Enabling Multiform Force Feedback in Virtual Reality Using Household Cleaning Robot. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'20, pages 1–12. ACM.

Wang, Y.-W., Lin, Y.-H., Ku, P.-S., Miyatake, Y., Mao, Y.-H., Chen, P. Y., Tseng, C.-M., and Chen, M. Y. (2021). JetController: High-Speed Ungrounded 3-DoF Force Feedback Controllers Using Air Propulsion Jets. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'21, pages 1–12. ACM.

Ware, C., Arthur, K., and Booth, K. S. (1993). Fish Tank Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'93, pages 37–42. ACM.

Wentzel, J., d'Eon, G., and Vogel, D. (2020). Improving Virtual Reality Ergonomics Through Reach-Bounded Non-Linear Input Amplification. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'20, pages 1—12. ACM.

White, M., Gain, J., Vimont, U., and Lochner, D. (2019). The Case for Haptic Props: Shape, Weight and Vibro-Tactile Feedback. In *Proceedings of the ACM SIGGRAPH Conference on Motion, Interaction and Games*, MIG'19, pages 1–10. ACM.

Whitmire, E., Benko, H., Holz, C., Ofek, E., and Sinclair, M. (2018). Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'18, pages 1–12. ACM.

Wiehr, F., Kosmalla, F., Daiber, F., and Krüger, A. (2016). BetaCube: Enhancing Training for Climbing by a Self-Calibrating Camera-Projection Unit. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'16, pages 1998–2004. ACM.

Williams, B., Narasimham, G., Rump, B., McNamara, T. P., Carr, T. H., Rieser, J., and Bodenheimer, B. (2007). Exploring Large Virtual Environments with an HMD When Physical Space is Limited. In *Proceedings of the ACM Symposium on Applied Perception in Graphics and Visualization*, APGV'07, pages 41–48. ACM.

Williams, E. J. (1949). Experimental Designs Balanced for the Estimation of Residual Effects of Treatments. *Australian Journal of Chemistry*, 2(2):149–168.

Williams, S. R., Suchoski, J. M., Chua, Z., and Okamura, A. M. (2022). A 4-Degree-of-Freedom Parallel Origami Haptic Device for Normal, Shear, and Torsion Feedback. *IEEE Robotics and Automation Letters*, 7(2):3310–3317.

Winfrey, K. N., Gewirtz, J., Mather, T., Fiene, J., and Kuchenbecker, K. J. (2009). A High Fidelity Ungrounded Torque Feedback Device: The iTorqU 2.0. In *Proceedings of the IEEE World Haptics Conference*, WHC'09, pages 261–266. IEEE.

Withana, A., Groeger, D., and Steimle, J. (2018). Tacttoo: A Thin and Feel-Through Tattoo for On-Skin Tactile Output. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'18, pages 365–378. ACM.

Yamaguchi, K., Kato, G., Kuroda, Y., Kiyokawa, K., and Takemura, H. (2016). A Non-Grounded and Encountered-Type Haptic Display Using a Drone. In *Proceedings of the ACM Symposium on Spatial User Interaction*, SUI'16, pages 43–46. ACM.

Yanagida, Y., Kawato, S., Noma, H., Tomono, A., and Tesutani, N. (2004). Projection-Based Olfactory Display with Nose Tracking. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'04, pages 43–50. IEEE.

Yang, H., Johnson, T., Zhong, K., Patel, D., Olson, G., Majidi, C., Islam, M., and Yao, L. (2022). ReCompFig: Designing Dynamically Reconfigurable Kinematic Devices Using Compliant Mechanisms and Tensioning Cables. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'22, pages 1–14. ACM.

Yang, J., Horii, H., Thayer, A., and Ballagas, R. (2018). VR Grabbers: Ungrounded Haptic Retargeting for Precision Grabbing Tools. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, UIST'18, pages 889–899. ACM.

Yem, V., Okazaki, R., and Kajimoto, H. (2016). FinGAR: Combination of Electrical and Mechanical Stimulation for High-Fidelity Tactile Presentation. In *ACM SIGGRAPH Emerging Technologies*, SIGGRAPH'16, pages 1–2. ACM.

Yokokohji, Y., Hollis, R. L., and Kanade, T. (1996). “What You Can See Is What You Can Feel” – Development of a Visual/Haptic Interface to Virtual Environment. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'96, pages 46–53. IEEE.

Yoshida, S., Sun, Y., and Kuzuoka, H. (2020). PoCoPo: Handheld Pin-Based Shape Display for Haptic Rendering in Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'20, pages 1–13. ACM.

Yu, R. and Bowman, D. A. (2020). Pseudo-Haptic Display of Mass and Mass Distribution During Object Rotation in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(5):2094–2103.

Zenner, A. (2016). Investigating Weight Distribution in Virtual Reality Proxy Interaction. Master’s thesis, Saarland University.

Zenner, A. (2020). Enhancing Proxy-Based Haptics in Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, VRW'20, pages 549–550. IEEE.

Zenner, A., Degraen, D., Daiber, F., and Krüger, A. (2020a). Demonstration of Drag:on – A VR Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'20, pages 1–4. ACM.

Zenner, A., Degraen, D., and Krüger, A. (2019a). Addressing Bystander Exclusion in Shared Spaces During Immersive Virtual Experiences. In *Proceedings of the Workshop on Challenges Using Head-Mounted Displays in Shared and Social Spaces at the ACM Conference on Human Factors in Computing Systems*, SHMD'19, pages 1–5.

Zenner, A., Klingner, S., Liebemann, D., Makhsadov, A., and Krüger, A. (2019b). Immersive Process Models. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'19, pages 1–6. ACM.

Zenner, A., Kosmalla, F., Ehrlich, J., Hell, P., Kahl, G., Murlowski, C., Speicher, M., Daiber, F., Heinrich, D., and Krüger, A. (2020b). A Virtual Reality Couch Configurator Leveraging Passive Haptic Feedback. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'20, pages 1–8. ACM.

Zenner, A., Kosmalla, F., Speicher, M., Daiber, F., and Krüger, A. (2018a). A Projection-Based Interface to Involve Semi-Immersed Users in Substitutional Realities. In *Proceedings of the IEEE Workshop on Everyday Virtual Reality*, WEVR'18, pages 1–6. IEEE.

Zenner, A., Kriegler, H. M., and Krüger, A. (2021a). HaRT – The Virtual Reality Hand Redirection Toolkit. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'21, pages 1–7. ACM.

Zenner, A. and Krüger, A. (2017). Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1285–1294.

Zenner, A. and Krüger, A. (2019a). Drag:on – A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'19, pages 1–12. ACM.

Zenner, A. and Krüger, A. (2019b). Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'19, pages 47–55. IEEE.

Zenner, A. and Krüger, A. (2020). Shifting & Warping: A Case for the Combined Use of Dynamic Passive Haptics and Haptic Retargeting in VR. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST'20 Adjunct, pages 1–3. ACM.

Zenner, A. and Krüger, A. (2021). Hand-Virtual Reality/Augmented Reality-Steuergerät, Virtual Reality/Augmented Reality-System mit demselben sowie Verfahren zur Simulation der Haptik. *German Patent 10 2019 105 854*. June 05, 2021.

Zenner, A., Makhsadov, A., Klingner, S., Liebemann, D., and Krüger, A. (2020c). Immersive Process Model Exploration in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(5):2104–2114.

Zenner, A., Regitz, K. P., and Krüger, A. (2021b). Blink-Suppressed Hand Redirection. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, VR'21, pages 75–84. IEEE.

Zenner, A., Speicher, M., Klingner, S., Degraen, D., Daiber, F., and Krüger, A. (2018b). Immersive Notification Framework: Adaptive & Plausible Notifications in Virtual Reality. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, CHI EA'18, pages 1–6. ACM.

Zenner, A., Ullmann, K., and Krüger, A. (2021c). Combining Dynamic Passive Haptics and Haptic Retargeting for Enhanced Haptic Feedback in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2627–2637.

Zhao, Y., Bennett, C. L., Benko, H., Cutrell, E., Holz, C., Morris, M. R., and Sinclair, M. (2018). Enabling People with Visual Impairments to Navigate Virtual Reality with a Haptic and Auditory Cane Simulation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'18, pages 1–14. ACM.

Zhao, Y. and Follmer, S. (2018). A Functional Optimization Based Approach for Continuous 3D Retargeted Touch of Arbitrary, Complex Boundaries In Haptic Virtual Reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'18, pages 1–12. ACM.

Zhao, Y., Kim, L. H., Wang, Y., Le Goc, M., and Follmer, S. (2017). 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, pages 82–91. ACM.

Zhu, K., Chen, T., Han, F., and Wu, Y.-S. (2019). HapTwist: Creating Interactive Haptic Proxies in Virtual Reality Using Low-Cost Twistable Artefacts. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, CHI'19, pages 1–13. ACM.

Ziat, M., Rolison, T., Shirtz, A., Wilbern, D., and Balcer, C. A. (2014). Enhancing Virtual Immersion Through Tactile Feedback. In *Adjunct Publication of the ACM Symposium on User Interface Software and Technology*, UIST'14 Adjunct, pages 65–66. ACM.

Zielasko, D., Bellgardt, M., Meißner, A., Haghgoo, M., Hentschel, B., and Kuhlen, T. (2017). buenoSDIAs: Supporting Desktop Immersive Analytics While Actively Preventing Cybersickness. In *Proceedings of the Workshop on Immersive Analytics at IEEE VIS: Visualization and Visual Analytics*, pages 1–4.

Zielasko, D., Horn, S., Freitag, S., Weyers, B., and Kuhlen, T. W. (2016). Evaluation of Hands-Free HMD-Based Navigation Techniques for Immersive Data Analysis. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, 3DUI'16, pages 113–119. IEEE.

Zielasko, D., Meißner, A., Freitag, S., Weyers, B., and Kuhlen, T. W. (2018). Dynamic Field of View Reduction Related to Subjective Sickness Measures in an HMD-Based Data Analysis Task. In *Proceedings of the IEEE Workshop on Everyday Virtual Reality*, WEVR'18, pages 1–6. IEEE.