UNIVERSITY OF COPENHAGEN
FACULTY OF SCIENCE

BUILDING BLOCKS OF FUTURE INTERFACES

CONCEPT-DRIVEN INTERACTION DESIGN IN VIRTUAL REALITY

ANDREEA MURESAN

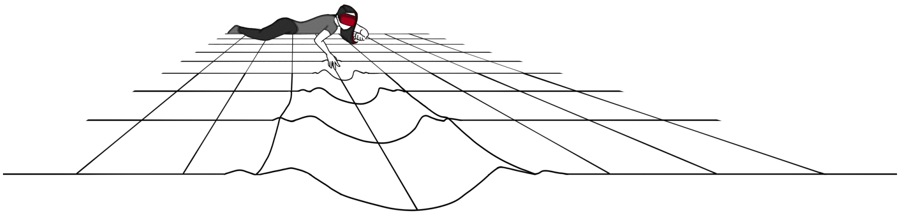

Supervised by KASPER HORNBÆK

Submitted: NOVEMBER 2023

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

Andreea Muresan

Building Blocks of Future Interfaces,
Concept-Driven Interaction Design in Virtual Reality

Supervised by
Kasper Hornbæk,
Professor, Human-Centred Computing,
University of Copenhagen

Assessment Committee:
Enrico Rukzio, *Institute for Media Informatics, University of Ulm*
Michael Nebeling, *Information Interaction Lab, University of Michigan*
Pernille Bjørn, Chair, *University of Copenhagen*

November 2023

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*".

Cover figure credits: Ioana Pop

Final Version as of November 5, 2023 (`classicthesis` version 4.2).

Memento Mori

ABSTRACT

Concepts like embodiment, presence, and immersion can be used to generate novel virtual reality (VR) interaction models. Many recent VR works highlight their *conceptual* contribution and its concrete instantiation — the prototype —- by describing what the prototype may achieve in *real-world* scenarios. This thesis adopts a conceptual lens of interaction design, considering *concepts* as the *building blocks of novel, future interfaces* and proposes design spaces as their default representation. Establishing a more principled way to represent novel interactions may help researchers situate their work and give a snapshot overview of interaction design. In turn, practitioners may use these design spaces as a real-world tool for design.

Chapter 2 combines several frameworks, tracing the historical context for concepts, design spaces, and prototyping and presenting design spaces as an evolving conceptualization tool within Human-Computer Interaction (HCI). Chapter 3 presents a research paper containing an analysis of 233 YouTube VR fail videos and develops the concept of *fail* as opportunities or as breakdowns. This research paper presents a classification of fails and their causes. Based on design implications, the paper sketches interactions meant to prevent fails or harness their positive implications.

Chapter 4 develops the traditional concept of *feedforward* for VR through theoretical grounding, prototyping, and an expert evaluation. The outcome of this research is a feedforward design space that embeds both practical and theoretical concerns and may serve as a tool for developers of tutorial and training VR applications of the future. Chapter 5 presents a research work (pending peer review) that develops the concept of *interacting through multiple avatars*. This work develops a design space from brainstorming workshops where experts generated ideas for using this concept in the real world. This work also instantiates a prototype, which includes a user-friendly interface for recording interactions. This artifact may serve as an authoring tool, making VR tutorials and feedforward development more accessible for non-technical designers.

Chapter 6 takes the overarching conceptual perspective of these works and provides additional context for how to generate and structure ideas, how to build on existing concepts, and how to embed usefulness and derive practical guidelines from concept-driven design through design-space-making. Chapter 7 presents a critical reflection of this work and provides an avenue to develop the works presented within this thesis further.

Lastly, in Chapter 8, I summarize the findings of this thesis and present a set of implications derived from the discussion in previous chapters. This work highlights the path *from theory to artifact* and *from artifact to theory* as a design space-making process. This process, upon dutifully recorded, ultimately serves to systematize knowledge and inform the design of future interfaces.

DANSK RESUMÉ

Udnyttelse af begreber som embodiment, tilstedeværelse og fordybelse muliggør nye virtual reality (VR) interaktionsmodeller. Mange nutidige VR-artikler fremhæver deres konceptuelle bidrag og koncpetets konkrete instansiering — prototypen — ved at beskrive, hvad prototypen kan opnå i real-world scenarier. Denne afhandling benytter sig af en konceptuel vinkel af interaktionsdesign, der betragter koncepter som byggestenene i nye, fremtidige brugergrænseflader og foreslår design spaces som deres standardrepræsentation. Etablering af en sådanne principiel måde at repræsentere nye interaktioner på kan hjælpe designere med at placere deres arbejde og hjælpe praktikere med at navigere mellem teori og praksis ved at give dem et øjebliksbillede af interaktionsdesign.

Chapter 2 kombinerer flere frameworks, for at spore den historiske kontekst for koncepter, designrum og prototyper og præsenterer designrum som et udviklende konceptualiseringsværktøj inden for Menneske-Maskin Interaction (engl. Human-Computer Interaction, HCI). Chapter 3 præsenterer en forskningsartikel, der indeholder en analyse af 233 YouTube VR-failvideoer, udvikler konceptet *fail* ud fejl eller sammenbrud og foreslår designimplikationer beregnet til at forhindre fejl og hvordan de kan designes. Denne forskning fremlægger en klassificering af fails og deres årsager og skitserer interaktionsteknikker for at informere det fremtidige design af VR i hjemmet.

Chapter 4 udvikler det traditionelle koncept *feedforward* til VR ved teoretisk forankring, prototyping og en ekspertevaluering. Resultatet af denne forskning er et feedforward-design space, der indlejrer både praktiske og teoretiske aspekter og kan tjene som et værktøj for udviklere af fremtidens tutorial- og trænings-VR-applikationer.

Chapter 5 præsenterer en forskningsartikel, (afventer peer review) som udvikler konceptet om *interaktion gennem flere avatarer*. Dette arbejde udvikler et design space ud fra brainstorming-workshops, hvor eksperter genererer ideer til at bruge dette koncept i den virkelige verden. Dette arbejde instansierer også en prototype, som inkluderer en brugervenlig grænseflade til optagelse af interaktioner. Dette artefakt kan tjene som et forfatter-værktøj, der gør VR-tutorial og feedforward-udvikling mere tilgængelig for ikke-tekniske designere.

Chapter 6 tager det overordnede konceptuelle perspektiv af disse værker og giver yderligere kontekst for, hvordan man genererer og strukturerer ideer, hvordan man bygger videre på eksisterende koncepter, og hvordan man indlejrer anvendelighed og udleder praktiske retningslinjer fra konceptdrevet design. Chapter 7 præsenterer en kritisk afspejling af dette arbejde og giver en informeret diskussion om design af fremtidige grænseflader i VR.

Til sidst, i Chapter 8, opsummerer jeg resultaterne af denne afhandling og præsenterer en opsummering af implikationerne afledt af den gennemtænkte diskussion i det foregående kapitel. Jeg fremhæver vejen fra teori til artefakt og *fra artefakt til teori* som designprocessen til at skabe rum. Denne proces, efter pligtopfyldende registreret, tjener i sidste ende til at systematisere viden og informere udformningen af fremtidige grænseflader.

Ever the dull alchemist
I have before me all the necessary elements.
It is their combination that eludes me.

I Trawl the Megahertz — Paddy McAloon [69]

ACKNOWLEDGMENTS

This Ph.D. thesis is a mountain of effort and knowledge produced during uncertain, unprecedented times. Many people have helped me along the way, and I am profoundly grateful for these people.

This work could not have been achieved without my supervisor, Kasper Hornbæk, who mentored me and allowed me the necessary freedom and flexibility to pursue my research interests. Kasper always found a way to improve any draft I had. I am grateful to my co-authors and their efforts in realizing the research work: Emily Dao, Jarrod Knibbe, Jess McIntosh, Aske Mottelson, Guido Makransky, and, more recently, Teresa Hirzle, who has been my mentor and my rubber duck. I extend my gratitude to all the experts who found time to participate in my studies. I also thank the researchers who turned my research visit into a pleasant stay, in particular Mel Slater, for the opportunity to do research at the Event Lab.

I extend my gratitude to my family, who have supported and encouraged me during this time and welcomed my ambitions. Doresc sa multumesc familiei mele pentru ca m-a sprijinit si incurajat in efortul meu de a obtine studii superioare, multumesc parintilor, surorii mele, bunicilor mei, matusa si unchiul meu. I thank my partner Martin, my supporter and cook in times of need, and my friends for their empathy, who were my cheerleaders, my video and figure models.

I am thankful for my colleagues, many of whom are friends, in the Human-Centered Computing Section from the University of Copenhagen, who commiserated and celebrated with me, made crunches a team effort, and effectively realized one of the best places to work in research: Arpit Bhatia, Hubert Zajac, Thomas van Gemert, Tor-Salve Dalsgaard, Martin Maunsbach, Joanna Bergström, Difeng Yu, Mantas Cibulskis, Valeria Borsotti, Pernille Bjørn, Irina Shklovski, and all others making up the HCC collective during my time there. Special thanks to ex-HCC members Carlos Tejada for being my concert buddy, Klemen Lilija, Jonas Schjerlund, and Hyunyounng Kim.

PREFACE

In 2019, I attended my first CHI conference after publishing an extended abstract as a Master's project [226] and co-authoring a couple of papers while working as a student worker [213, 228]. While I have conducted research for some time, this thesis captures my research output since starting my Ph.D. program in May 2020.

In 2020, the global COVID-19 pandemic started overtaking the world. Soon after, in the fall of 2020, everyone would go remote, and onsite research activities would stop. While I enjoyed digging into my research topics, every day, the news from home and the world dragged my daily life into a kind of global perspective that I did not need. Some struggles put other struggles in perspective. On the one hand, the PhD required much research work and implementation to determine exciting insights for prospective VR projects. On the other hand, people were dying and losing their jobs en masse. In this global context, researching virtual reality seemed prescient. There was no other option. Virtual reality was the new reality, and going offline was not an option. The VR headsets were a novel way to be online for some [255].

I started my PhD with a different topic than I ended the PhD with. My colleagues and I were single-minded and focused on one concept –– *body ownership*. It was the perfect opportunity to run remote studies with everyone stuck indoors. I implemented a robust VR application where people sorted cubes with a green hand (low body ownership condition) and a skin-matched hand (high body ownership condition). A pilot, 157 participants, and 37 gigabytes later, the machine learning algorithms had poor performance. People's homes are messy, lighting conditions vary, and using movements to determine body ownership levels was more complicated than I thought. The outcome? Project successfully on hiatus. But wait, people's homes are messy? That sounds familiar.

Cue the abstract for *Bad breakdowns, useful seams, and face slapping* [263]. People's homes did have some of these qualities, some of which made it hard to run studies remotely. The pandemic might have kickstarted VR sales according to some [255], and yet the most significant barrier to mass adoption seemed to be user experience[1]. A timely topic emerged — how do people use VR in the home, and what can we learn from that? The global context and this research work changed my interests. There were still ample opportunities to discover simple concepts that would lead to novel VR design, the kind that, in hindsight, may elicit reactions of the type — *"Of course, that's obvious! Why didn't I come up with that?"* — reaction. The videos we analyzed gave us some clues, based on which we suggested a few designs. This analysis and suggestions would shed light on why VR lacked in user experience.

So, for the rest of my research, I embarked on a quest for clues. In particular, I searched for crucial information describing the building block of a technique, an interface, and an interaction. The main problems tackled within this thesis put forward solutions that are novel designs based on different types of information or concepts. Each paper has the stand-alone contribution of putting forward a technique that answers a need, fills a gap,

---

[1] https://techjury.net/blog/virtual-reality-statistics

or prevents a problem. In this thesis, I will take the opportunity to reflect on the process of revealing these needs or concerns — a conceptual process. Every researcher has done this type of work — looking for a gap to fill with their work. However, reflecting on this process is critical to determining a fruitful method. In my work, I used design spaces to formalize searching for gaps. The research papers within this thesis —include novel VR interaction designs based on YouTube videos, — present a design space for VR interactions that show people what to do and how to do it, and — generate a design space for acting through multiple avatars in VR.

PUBLICATIONS

The main part of this Ph.D. thesis includes two peer-reviewed articles and one in manuscript (CoreX), which represent the core research contributions. I contributed to other works throughout my Ph.D. (OtherX), which I will touch upon but not discuss to maintain the coherence of this work. The contributions are:

** = Shared first authorship*

[Core1] Andreea Muresan*, Emily Dao*, Kasper Hornbæk, and Jarrod Knibbe. “Bad Breakdowns, Useful Seams, and Face Slapping: Analysis of VR Fails on YouTube.” In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI ’21: CHI Conference on Human Factors in Computing Systems. New York, NY, USA: ACM, May 6, 2021, pp. 1–14. ISBN: 978-1-4503-8096-6. DOI: 10.1145/3411764.3445435. URL: https://doi.org/10.1145/3411764.3445435.

[Core2] Andreea Muresan, Jess McIntosh, and Kasper Hornbæk. “Using Feedforward to Reveal Interaction Possibilities in Virtual Reality.” In: *ACM Trans. Comput.-Hum. Interact.* (2023). ISSN: 1073-0516. DOI: 10.1145/3603623. URL: https://doi.org/10.1145/3603623.

[Core3] Andreea Muresan, Teresa Hirzle, and Kasper Hornbæk. “How and Why To Act Through Multiple Avatars in Virtual Reality.” In: Pending Peer-Review Process. Association for Computing Machinery, 2024.

[Other1] Aske Mottelson, Andreea Muresan, Kasper Hornbæk, and Guido Makransky. “A Systematic Review and Meta-Analysis of the Effectiveness of Body Ownership Illusions in Virtual Reality.” In: *ACM Trans. Comput.-Hum. Interact.* (2023). ISSN: 1073-0516. DOI: 10.1145/3590767. URL: https://doi.org/10.1145/3590767.

[Other2] Andreea Muresan and Mel Slater. “Shared and Individual Experiences of a Virtual Reality Concert.” In: Pending Ethics Approval Process.

CONTENTS

1

INTRODUCTION

In HCI (Human-Computer Interaction), concepts and artifacts often co-occur. Once in a while, paradigm-shifting concepts emerge that are demonstrated through smart artifacts. *Dynabook* [123] and *Sketchpad* [45] are great examples of visionary ideas within HCI. Alan Kay's *Dynabook* vision and the technological transformation that occurred after that culminated in the personal computing revolution decades later [84]. Kay's graduate supervisor was Ivan Sutherland, famous for *Sketchpad* [45], considered the first instance of interactive graphics. Kay referred to *Sketchpad* as *"the first personal computer"* [84] and considered Sutherland highly influential to his object-oriented thinking [84]. Both artifacts embody concepts that would later become fundamental for everyday life.

Maxwell [84] gives an account of the digital transformation started by Dynabook in his thesis.

Sutherland later pondered on the *"ultimate display"*, a device that could immerse users in the virtual world. This conceptual pondering resulted in the *"Sword of Damocles"*, now considered the first head-mounted display, or virtual reality (VR) display. Another significant VR development emerged from Jaron Lanier and the team at JPL Research, who further modernized the VR headset [183]. Recently, a critical year in VR development was the Oculus Headset Launch in 2016, and its improved version in 2020 [183]. These developments popularized VR as a technology in the home, making it vastly more accessible to practitioners and users alike, perhaps benefiting from the context of the pandemic [255]

Within HCI, artifacts like prototypes are considered research contributions [184] and concrete manifestations or concepts [97]. Prototypes *implicitly* account for the concept-generation process. Recent VR works highlight the *conceptual* contribution to VR design. The contribution of these works is the concept and its instantiation — the prototype — which gives a scenario-based account of what the concept may do in the *real-world* [282, 203, 298]. Their evaluations mostly involve expert interviews or feasibility studies [244, 245, 237], in some cases, evaluations of performance [291, 277].

To *explicitly* account for the concept generation process, some researchers start by mapping the design space of interaction techniques from literature [261] or preliminary expert interviews [237]. Design spaces aid in idea generation by enabling a systematic categorization of related works [220] to reveal gaps [261]. The researchers typically fill these gaps by instantiating parts of the design space as prototypes or interaction techniques. This approach is meant to account for the *novelty* claims put forward by researchers regarding their prototypes. In addition, design spaces provide valuable abstracted overviews of artifacts and reveal different lenses of interactions [220].

The theoretical implication of the thesis contribution.

To help bridge the gap between the theoretical aspects of design spaces and their practical applications as real-world design tools, we approach concept generation systematically and anchor it to usefulness. This research borrows the tenets of Stolterman and Wiberg [119]'s *concept-driven design*, namely, a view that (1) interaction design evolves from futuristic use scenarios, (2) reasons from theory, (3) and has an explorative nature, purposefully manifesting visions in tangible designs. The work in this thesis presents design spaces not just as points in space but as representations of models of generative interactions, following Beaudouin-Lafon, Bødker, and Mackay [257]'s path from theory to artifact. This thesis attempts to instantiate design spaces both as *paths from theory to artifact* and also as *from artifact to theory*. This thesis aims to establish concept-driven design in virtual reality as a design-space-making task, which reveals the underlying interaction models of prototypes.

The practical implication of the thesis contribution.

Currently, virtual reality is undergoing a reinvention of itself as the race for the ultimate display continues. With the advent of so many novel systems, prototypes and interaction techniques, navigating this mountain of knowledge requires substantial work. Researchers and practitioners must understand the interaction models proposed to situate the work and build on it with novel contributions. Design spaces offer semantic accounts of systematic literature reviews and represent a research contribution themselves [261, 220]. To aid in comparing and structuring this mountain of effort, we suggest using design spaces not just for idea generation but as semantic tools for representing the parameters that bound the interaction model underlying an interaction technique.

This systematic account would enable the designer to situate a prototype in the sea of concepts. Accounting for the idea-generation process would allow researchers to make more informed design decisions before prototyping without pursuing a systematic review themselves. VR is a far cry from being prevalent as a consumer product, but the estimates are positive[1]. As the VR finds its footing as a large-scale consumer product, researchers strive more and more to tie VR use to real-world situations to hasten this process of consumer VR adoption.

This thesis highlights another conceptual avenue for design spaces —- not just as paths *from theory to artifact*, but *from artifact to theory*. This thesis proposes two design spaces that model whole classes of generative interactions in VR. The design space dimensions serve as models of interactions, whereas the parameters reveal practical design opportunities for exploration. This work informs the future design of VR interactions by analyzing VR in the home and proposing design implications through the conceptual lens of *seamful* design. This work informs the design and implementation of future tutorial and training applications by revealing a model, a design space, and an authoring tool for feedforward interactions in VR. This thesis also gives a first systematic account of how to approach designing interactions through multiple avatars and in which application scenarios this might be useful.

---

[1] https://www.fortunebusinessinsights.com/industry-reports/virtual-reality-market-101378

1.1 ABSTRACTS AND OUTLINE

From the research work I have done during my Ph.D., this thesis covers topics from three research papers. The papers are presented as published or in manuscript, with their associated related work to help contextualize their contribution. The research works cover the following topics, presented as their respective abstracts:

1. *Bad Breakdowns, Useful Seams, and Face Slapping: Analysis of VR Fails on YouTube*
   Virtual reality (VR) is increasingly used in complex social and physical settings outside of the lab. However, not much is known about how these settings influence use, nor how to design for them. We analyse 233 YouTube videos of VR Fails to: (1) understand when breakdowns occur, and (2) reveal how the seams between VR use and the social and physical setting emerge. The videos show a variety of fails, including users flailing, colliding with surroundings, and hitting spectators. They also suggest causes of the fails, including fear, sensorimotor mismatches, and spectator participation. We use the videos as inspiration to generate design ideas. For example, we discuss more flexible boundaries between the real and virtual world, ways of involving spectators, and interaction designs to help overcome fear. Based on the findings, we further discuss the 'moment of breakdown' as an opportunity for designing engaging and enhanced VR experiences.

   When referring to each paper, I will use a shorthand, for this one VRFails, which is [Core1] or [263].

2. *Using Feedforward to Reveal Interaction Possibilities in Virtual Reality* In virtual reality (VR), interactions may fail when users encounter new, unknown, or unexpected objects. We propose using feedforward[2] in VR to help users interact with objects by revealing how such objects work. Feedforward lets users know what to do and how to do it by showing the available actions and outcomes before an interaction. In this paper, we first chart the design space of feedforward in VR and illustrate how to design feedforward for specific VR interactions. We discuss starting the feedforward, previewing actions and outcomes, and returning the virtual world to its state before the feedforward. Second, we implement three real-world VR applications to show how feedforward can be applied to multistep interactions, perceived interactivity, and discoverability. Third, we conduct an evaluation of the design space with 14 VR experts to understand its usefulness. Finally, we summarize the findings of our work on VR feedforward in 15 guidelines.

   I will refer to this work as Feedforward, which is [Core2] or [305].

3. *Why and How to Act Through Multiple Avatars in Virtual Reality*
   In virtual reality (VR), users typically control one virtual body — their avatar. Previous works enable users to act through multiple avatars simultaneously. These works, though, lack a systematic account of *why* users want to interact with multiple avatars and do not explain *how* users can manipulate and generate avatars. To address this, we run six workshops with 12 VR experts and develop a design space that captures four fundamental dimensions for acting through multiple avatars in VR: *Appearance*, *Context*, *Input/Output*, and *Control*.

   I will refer to this work as MultipleAvatars, which is [Core3] or [309], in manuscript.

---

[2] Feedforward is an HCI concept that involves showing people what to do before they do it, like the iPhone unlock screen. Feedback shows the result of an action after performing it.

Researchers can use the design space to generate novel interaction opportunities involving multiple avatars or analyze existing work. We then run a usability study with 17 participants to understand the practicalities of an interface that integrates parts of the design space, which reveals conceptual and technical challenges that we address through design recommendations.

Chapter 2 brings together several frameworks, tracing the historical context for concepts, design spaces, and prototyping. The purpose of this chapter is to present design spaces as an evolving conceptualization tool within HCI. This chapter presents the design space as an implicit product of prototyping, which, upon being made explicit, may serve as the default conceptualization tool for scenario-based artifact contributions within HCI.

While the papers have their individual contribution, in this thesis, I will reflect on the methods, implications, and contribution of each paper from a conceptual lens. Section 1.2 gives an overview of the individual contributions of the research work and the thesis contribution, which presents design-space-making as semantic representations for concept generation in VR.

Chapter 3 presents VRFails and is based on [263]. Chapter 4 presents Feedforward and is based on [305]. Chapter 5 presents my latest work, submitted to a conference and pending a peer-review process.

In Chapter 6, I provide some additional context for the research work and discuss implications for design, novelty, usefulness of concepts and design spaces. This chapter explains and draws theoretical implications for approaching artifact creation in a more systematic manner through design spaces and various types of knowledge that determine the design space.

In Chapter 7, I address the limitations of the research work and of design spaces in general and provide an informed discussion about the future of interaction design in virtual reality.

Lastly, in Chapter 8 I provide closing remarks, and present a summary of this research work and its implications for the creation of the VR interfaces and interactions of tomorrow.

1.2 CONTRIBUTIONS

The main purpose of this work has been to advance the design of virtual reality interfaces by assuming a concept-driven lens. To this end, the thesis has several HCI research contributions and implications [184, 295] in line with the concept-driven interaction design framework [119].

1.2.1 *Interaction design is explorative in nature, purposefully manifesting visions in tangible designs*

This thesis gives a systematic account of idea generation and how it may be represented through design-space-making and prototyping. The research works explore concepts and manifest artifacts that serve as an *artifact contribution* and *design implications*.

1. VRFails analyzed 233 YouTube videos of VR fails using qualitative content analysis and suggested a systematic classification for them. This research presents design implications as comic figures that capture imagined scenarios of VR use.

2. Feedforward developed the concept of feedforward systematically as an interaction model, a design space, and a prototype which enables the exploration of its design space. This research puts forward an authoring tool prototype that generates feedforward interactions, instantiating various parts of the design space.
3. MultipleAvatars developed the concept of acting through multiple avatars by brainstorming use cases for it and systematically generating a design space that captures its key characteristics. This research generated a prototype that instantiated aspects of the design space and presented design implications and empirical findings from a user study.

1.2.2 *Interaction design evolves from futuristic use scenarios*

To bridge the gap between concepts as abstractions and concepts as prototypes in the real world, we embed usefulness in various stages of the interaction design. The work in this thesis highlights the paths of concepts not only from theory to artifact, but from artifact to theory. The research contributions are grounded in futuristic use cases as follows, involving *method contributions* and *implications*:

1. VRFails imagines future use cases of VR in the home, and based on an analysis of YouTube videos, presents design implications as low-fidelity artifacts (sketches). This research informs the interaction design of future VR headset use in the home.
2. Feedforward develops a design space with experts who imagined contexts of use for feedforward interactions. Feedforward is evaluated *practically* — as a prototype for generating feedforward interaction, and *theoretically* — as a way of reasoning about the design of feedforward. This research informs the interaction design of tutorial and training VR applications with theoretical and practical implications.
3. MultipleAvatars approaches the interaction modeling problem starting from *the why* in an attempt to embed real-world usefulness in the parameters of the interaction. During these workshops, the experts imagined various real-world use cases for using multiple avatars in virtual reality, some of them in the context of existing VR applications. This research offers a user-friendly interface for authoring feedforward interaction through motion capture, making it more accessible for non-technical designers.

1.2.3 *Interaction design reasons from theory*

This thesis builds on theoretical concepts to inform novel interaction design and propose new models of interaction, as *theoretical contributions* and *method implications*:

1. VRFails introduces the concept of *fail* from a VR social perspective and uses the conceptual lens of *seamful* design to generate novel interaction based on the concept of *fail* as *breakdown*, or *fail* as *opportunity*.
2. Feedforward builds on the theoretical concept of feedforward in HCI, and further develops it for VR. The feedforward design space is a

semantic or theoretical representation of feedforward in VR, but may also be used as a practical tool to help practitioners generate novel designs.

3. MultipleAvatars develops a model for using multiple avatars in virtual reality using concepts from generative models of interaction [257], semantically represented as a design space. This space is used to categorize related work, both in VR and then AR, showcasing its *generative*, *comparative*, and *descriptive* power. This research has shown that extending the scope of related work has the potential to fill in design-space gaps and, thus, has implications for methods to generate and validate design spaces.

1.2.4 *Summary*

This thesis highlights the conceptual work needed to generate novel interaction methods in VR. Apart from providing a path from theory to practice as a design space, this thesis highlights another path — *from artifact to theory*. For Feedforward, the expert study led to extending the feedforward design space by embedded practical concerns, not just theoretical ones. VRFails introduces the concept of *VR fails* and informs future design by analyzing the landscape of VR interaction in the home, developing design implications based around the concept of *seamfulness*. MultipleAvatars uses general morphological analysis to obtain a design space that highlights its *descriptive, generative, and evaluative* power by mapping related works beyond VR. Table 1 shows an overview of the research contributions of each paper. This thesis concludes with additional implications based on the related work presented in Chapter 2, and a discussion in Chapter 6, which takes the overarching conceptual perspective of these works.

Research	*Contribution*	HCI Contribution
VRFails *published*	Classification Low-fi Artifacts (Sketches) as design ideas	theoretical contribution (defining VR fails concept), artifact contribution, and design implications
Feedforward	Design Space	theoretical and method implications
published	Guidelines	empirical and design implications
	Artifact	artifact implications
	Model	theoretical contribution
MultipleAvatars	Design Space	theoretical and method implications
pending review	Guidelines	empirical and design implications
	Artifact	artifact implications

Table 1: Overview of the contributions and implications of the research from an HCI and practical perspective.

1.3 OTHER WORK

This section contains excerpts of my work in the Event Lab, which is pending an ethics approval process.

At the beginning of my PhD, I researched body ownership with colleagues from the Department of Psychology. An outcome of this research has been Aske's systematic review (Other1), which I co-authored. In another related project, we ran an online study with a between-subjects setup where people used green hands or skin-matched human hands to perform some hand movements, sort blocks according to color, and fill in a follow-up embodiment questionnaire[3].The accuracy and precision of the machine-learning models were subpar, however. Pending improving the algorithm, a lab study could have yielded more functional data.

In September 2022, I had the fantastic opportunity to visit Mel Slater's Event Lab at the University of Barcelona. I worked on a project there extending their work in realizing virtual reality concerts (Other2). My work in the Event Lab combined my fondness for concerts and my interest in shared VR experiences. More specifically, this study would use the same setup as one of their current projects — a replica of Dire Straits performing *"Sultans of Swing"* in 1983. In this follow-up study, we will concentrate on shared versus individual experiences of a virtual rock concert. The qualitative outcome from their previous research was interesting because it revealed that people may conceptualize the same setup completely differently. For example, research showed females showed more negative sentiments toward the experience of virtually attending rock concerts than males within the same virtual setup [307]. This perceptual difference of the same event reveals a gap between what VR communicates versus what users perceive, which relates to the plausibility aspect of presence. With this research, we aimed to explore what causes this gap, specifically by looking at uncanny valley and paranoia measures, together with past concert-going experiences of users. To assess these measures, we gather responses to the modified Godspeed uncanny valley questionnaire [116], the Green Paranoia scale [93], and other demographic data and collect surveys about the participants' experiences, which will be analyzed through sentiment analysis.

This study would advance scientific understanding of the mechanism that helps establish shared realities in virtual space and whether friends' presence may alter such perceptions. Just as in [307], we expect that participants observe behavior from virtual agents that were not implemented. This research will reveal whether such observations are enhanced or lessened in a shared environment. Establishing what enhances realism from events could help to elicit realistic behavior for psychological training (e.g., treating phobias, anxieties, and paranoia similar to [169, 185] to develop behavior change mechanisms). This work remains ongoing.

1.4 METHODS

This section gives an account of the methods used for the research work presented in this thesis. Table 2 shows an overview. Each research work began with source materials and involved a method to develop the concept and a method to generate and/or evaluate artifacts. The outcome of the research is a contribution that involves theoretical aspects — in terms of interaction models or concepts, and practical aspects — in terms of prototypes, design

---

[3] A video of this project may be found on YouTube here if the reader is interested in the implementation.

guidelines, and design implications based on usability studies. This research highlights the cyclical, iterative process of designing interaction in VR.

	VRFails	Feedforward	MultipleAvatars
Methods	Video content analysis seamfulness lens Sketching	Design space Literature review Prototyping Expert usability study	Workshop Affinity diagram Design space User walkthrough usability study
Theoretical Contributions and Implications	VR Fails concept VR Fails classification	Feedforward concept Design space Design space	Multi-avatar interactions concept Design Space
Practical Contributions and Implications	Design implications Scenario-based sketches	Design implications Authoring tool prototype	Design implications Prototype

Table 2: Overview of the methods used across the three research papers.

Situated in the field of human-centered design, the works presented in this thesis adopt a qualitative perspective [90] and a co-design, context-situated approach. Qualitative methods are often used in design to highlight experience and "reflections or intuitions about a design" [94]. ISO [221] defines *context of use* as *"the combination of users, goals and tasks, resources, and environment"*.

With this view, VRFails describes a qualitative content analysis of videos of VR users to make informed recommendations about designing for VR at home. Feedforward includes an expert study where a prototype instantiating various parts of a design space is evaluated from a theoretical model and a usability perspective.

MultipleAvatars includes a workshop where experts brainstormed application scenarios and technical requirements of a hypothetical system. A later prototype of this design process was evaluated with users to gather qualitative data related to the usefulness and usability grounded by practical tasks. In this study, Likert questions serve as tools for reflection and preliminary results. Qualitative methods highlight experience and *"reflections or intuitions about a design"* that reveal complexities and hypotheses beyond subjective questionnaires [94].

Next, we present an overview of the methods used in this research work.

1.4.1 *Literature review*

Each research paper presents its background and related work, which was determined by a literature review. Xiao and Watson [234] give an overview of the types of literature reviews, while the PRISMA guidelines have been adopted by the HCI community, including VR [251], to report on systematic reviews and scoping reviews [110]. Literature reviews are commonly performed to derive design spaces, but as discussed above, this research underscores the importance of involving experts in the decision-making process. For Feedforward, we relied heavily on literature and took a deductive approach similar to Mayring [74]'s deductive (top-down) and inductive (bottom-up) view of thematic analysis. We first derived the general model from the literature for the concept, starting with Norman's *"stages of action model"* [144]. For MultipleAvatars, we performed a review to identify related works and classify them using the design space obtained via a brainstorming exercise.

1.4.2 *Workshops and Interview*

We follow a co-design approach based on *requirements elicitation and analysis* process [147, 44, 68, 73]. For MultipleAvatars, to tease out the concept, we ran formative studies with expert users who were interviewed and then performed a written brainstorming exercise [16]. We showed participants a video mock-up as a *concept demonstrator* [44]. We conduct the study by running workshops between pairs of VR experts, who are involved as *potential users* and aid the design and development of a system. Open coding was most often used to report on the interview and relate the findings to the literature. Following a usability study evaluating a prototype of MultipleAvatars we interviewed participants using prompts or questions from well-established qualitative surveys [200, 235, 102]. In this breadth-first analysis, we focused on meaning rather than the scores given by participants to the questions and further asked them to explain themselves. This enabled us to identify very broad use cases for the multiple avatar concepts based on how participants related to the avatars.

Prototype evaluations with experts are common in VR and they usually follow interviews where design guidelines are presented [237, 150, 275, 245]. Following a co-design approach, for Feedforward involves a walkthrough usability study with experts, followed by an interview.

1.4.3 *Open coding and thematic analysis*

Open coding refers to the process of inductively forming categories Mayring [74], which underlies the first step of grounded theory [107]. Thematic analysis is often used for grounded theory and concept generation and involves six steps as put forward by Braun and Clarke [80]: getting an overview of the data, generating codes, identifying and reviewing codes, defining the themes and summarizing the process in terms of a report which relates to the literature, the themes, the codes with examples. This procedure was used to analyze interview data in Feedforward. We used open coding to report on interview data and followed a process similar to thematic analysis to derive the design space from the affinity diagram in MultipleAvatars.

1.4.4 *Affinity diagramming*

Affinity diagramming helps instantiate, organize, and make sense of unstructured, dissimilar qualitative data. We used the KJ method for affinity diagramming [165]. This method involved making a diagram bottom-up one by one, taking a note, labeling it, and placing it on the table until all notes were labeled. During each labeling, "lone wolves" were placed on the side to be re-coded or discarded. We iterated over the diagram to re-code, rename, or collapse codes if necessary. In the end, we re-drew the diagram and either coded or discarded the remaining lone wolves.

Affinity diagramming is often used to interpret and analyze prototypes [165]. In MultipleAvatars, we used an affinity diagram and a thematic analysis process to generate the design space. The affinity diagram produced clusters of ideas from the brainstorming workshop with the experts. We derived dimensions and parameters for a design space that could instantiate all ideas using a thematic analysis process.

1.4.5 *Video content analysis*

YouTube videos have been used as a data source in HCI to provide insights into user interaction with certain technologies [101, 223] as well as real-world use of devices by specific groups of users [138, 163]. Anthony, Kim, and Findlater [138] conducted a content analysis of 187 YouTube videos to explore touchscreen use by people with motor impairments. As a resource to examine a social internet phenomenon, however, online videos have not been widely used, and we seek to do this in our paper. For VRFails, coding was then done in iterative steps, with all authors going over the videos independently and discussing disagreements before updating the coding scheme. We followed the steps of the qualitative content analysis to perform the video analysis, frequently used in mass-media studies [74]. The procedure is as follows: problem formulation, assembling sample video, determining category or coding scheme, defining the codes/categories, identifying the prototype of each category and defining its limits, and lastly, coding the data, comparing and determining frequencies, and providing a thoughtful interpretation of the results. This type of analysis results in a classification. In the case of VRFails, we identified types of fails and types of causes of fails.

Content analysis has specific quality criteria, according to [74]. The quality criteria involve semantic validity, sampling validity, correlational validity, predictive validity, and construct validity. The inter-coder agreement contributes to the reliability criteria of reproducibility, stability, and accuracy. We ensured coding validity by calculating Fleiss' kappa [4] to report agreement across the raters following the inductive category procedure. The outcome of the analyses represents frequency analyses.

1.4.6 *Prototyping and sketching tools*

Prototyping and sketching are common tools for obtaining artifacts in HCI [86] and VR [241]. For this research work, I used Blender, Figma and ClipStudio Paint were used for sketching and ideation, whereas the artifacts were implemented with the Unity Engine and Oculus Quest headsets, later purchased by Meta and rebranded as Meta Quest. Prototyping is often used in VR to evaluate and instantiate parts of the design space in the process of design space filtering [237, 150]. To generate the affinity diagram notes for MultipleAvatars from Excel, I used Figma and two plugins, Google Sheets Sync and CopyDoc Text Kit.

1.4.7 *Usability study*

This research mostly involved walkthrough studies as the aim was to understand the user's experience of the system. Walkthrough studies are a particular type of usability studies [36], common to conceptual VR research questions [266] and VR system with a wide array of interactions like in *Spacetime* [211], *Poros*, and *VRSketchIn* [237]. Tasks may be introduced in order of complexity to aid in learning [241]. Following ethics guidelines from the University of Copenhagen, we used an informed consent form specifying the risks and the rewards of the experiments and anonymized and stored data in accordance with GDPR rules.

1.4.8 *Design spaces*

Section 2.1 provides an in-depth overview of design spaces and the methods to uncover them. Feedforward and MultipleAvatars both generated design spaces assembled through different methods. Feedforward constructed a design space based on HCI theory, prototyping, and an expert evaluation. MultipleAvatars generated a design space based on brainstorming ideas from a workshop with VR experts and users. Conversely, VRFails provides a classification of VR fails instead of a design space.

1.4.9 *Ethics*

To perform the studies, we followed the ethical guidelines of the University of Copenhagen. The university follows Data protection legislation to prevent abuse and maintain informed consent for data processing and storage. To adhere to these principles, participants were given informed consent and explained the study beforehand. Participants were allowed to withdraw from the study or take a break at any moment, and they were under the constant supervision of the experimenter. We followed the data storage guidelines from the university, which confirmed the GDPR standards. More information can be found on the university's website: here.

2

BACKGROUND AND RELATED WORK

Concepts have a long-standing history in HCI — *Dynabook* predicted portable computers, building upon notions from *Sketchpad*, which represented the modern graphical user interface. *"Concept-driven interaction"* is a methodology to draw out theoretically grounded advancements focused on the future [119]. This method begins with theory, evolves with artifact development, and results in a design that embodies a concept rather than answers a problem. The concept-driven approach thus settles *theory* as a building block that abstracts knowledge related to "*fundamental entities at the core of a discipline*" [119]. The concept-driven approach incorporates *"disciplined imagination"*[23] as a core tenant, a view that emphasizes imagination, mapping, conceptualizing, and speculating. These activities are conducive to developing theories, which are seen as ordered *"assertions about a generic behavior"* [23] that hold across instances [23, 119, 8]. Design spaces are a type of HCI theoretical contribution that instantiates these characteristics of concepts: they provide an ordered way to characterize the qualities of a whole. This chapter gives an account of how concepts, prototypes and design spaces are related and a reasoning for why design spaces may be a suitable representation for concepts in HCI.

2.1 CONCEPTS AND KNOWLEDGE IN HCI

Most researchers fundamentally agree that design spaces are "*conceptual spaces*" [173, 86, 149, 97, 79]. However, concepts themselves are distinct from design spaces and capture a different type of knowledge than theory. The following section gives an overview of the types of knowledge produced within HCI and how design spaces serve as a unifying representation.

HCI has been framed as problem-solving research[180], sometimes targeting *wicked problems* with speculative design approaches [278]. Oulasvirta and Hornbæk [180] classified HCI research problems as *empirical* — real-world descriptions of HCI phenomena, *conceptual* — explaining unrelated phenomena, and *constructive* — understanding interactive artifacts. The researchers highlight that *"without conceptual 'glue' to anchor them, constructive contributions readily remain point designs and empirical studies point studies."*[180]. Gaver [128] reflects on the "*scientific*" nature of theories under the *research through design* umbrella and suggests having moderate expectations with respect to the creation of *verifiable theory*. He stresses the value of exploring, speculating, and manifesting tangible results in terms of *"conceptually rich artefacts"* [128], as opposed to standardization. HCI design methods often tackle *wicked problems* [128]. Coined by Horst Rittel, wicked problems [5] refers to socially situated, complex problems that cannot be solved linearly — solutions to these problems change their understanding. Tom Ritchey found that *wicked problems* were particularly well-suited for morphological analysis, having successfully applied it to complex policy issues. Morphological analysis has been used as a method within HCI to generate design spaces, such as for VR gaze interaction [220].

In a paper report, Wobbrock and Kientz [184] identifies seven types of research contribution for HCI: empirical, artifact, method, theory, dataset, survey, and opinion. More recently, Berkel and Hornbæk [295] discuss seven types of implications for HCI research, which serve to present a short summary of the research and highlight important parts. These implications inform various aspects of methodology, theory, design, practice, community, policy, and society. Rogers [290] provides an overview of knowledge types within HCI, which are not necessarily mutually exclusive and may overlap. Relevant for this thesis, *models* serve as shorthands for actionable theory basics, *frameworks* represent area-specific knowledge scoped into concepts, questions, challenges, and principles related [290]. It is easier to exemplify concepts in HCI rather than define them. *Dynabook*, affordances, and feedforward are some examples of concepts [290]. Fallman [64] presents the romantic, pragmatic, and conservative way of idea generation and put forward a distinction *"between the knowledge-generating Design-oriented Research and the artifact-generating conduct of Research-oriented Design."*

2.1.1 *Concepts*

Attempts to define concepts have led to two contradicting [67] schools of thought. Outside of HCI, concepts are strongly related to semantics and meaning and are highly debated in fields like psychology and philosophy. In cognitive psychology, concepts refer to categories and describe mental models, where membership is non-arbitrary [134]. Each field may have its own view of concepts, with various degrees of similarity. For cognitive science, "*concepts are the building blocks of thought*" [50] and represent the *"mental glue"* [75] that underlies reasoning.

The main models explaining concepts refer to exemplars and prototypes (for a detailed overview of concepts, refer to [67, 134]). Exemplar theories [15, 11] consider that categories are determined by exemplars, or remembered instances. Prototype theory [7, 12] posits that common or typical features rather than instances determine concepts. According to this model, prototypes can emerge empirically by determining features of several category instances [67]. Exemplars and prototypes are also concepts within HCI, which perhaps best exemplify how domain-interdependent concept formation is.

Concept blending refers to constructing meaning from different sets of knowledge. This process has been analyzed by cognitive semanticists from the perspective of an integration model, whereby two inputs from different knowledge domains are projected in a mental space called the blended space [65, 56]. Metaphors are related to conceptual blends and may be defined as referencing one domain with language from a distinct domain [56], which have been used in HCI and VR. For example, the *the virtual hand* or *pointer metaphors* refer to groups of similar interactions that may be instantiated by different implementations [85].

2.1.2 *Classification*

Prototype and exemplar theories have implications regarding the process of classification within HCI. Mayring [74] provided a six-step approach in type-building content analysis, which results in a typology. The steps are: defining the dimensions, logic i.e., extreme types and frequent types, inductively developing categories from types, revising, choosing type representatives (*exemplars*), and describing results. For inductive category development, Mayring [74] added a coder agreement check as an empirical reliability validity measure.

Another related field here is set theory, a branch of mathematics solely devoted to establishing strict rules in combinatorics. Goguen [3]'s attempts to mathematically model inexact concepts have resulted in fuzzy logic.

In semantics and social science, categories are represented by typologies and taxonomies[1]. Typologies, seen as conceptual classifications, are multidimensional, non-hierarchical structures that are fully overlapping. Combining 2x2 conceptually non-overlapping category dimensions results in a typology that often has names or types in its cells [31]. Taxonomies are strictly hierarchical and non-overlapping structures (think of a tree structure) that classify empirical entities [31]. Typologies begin from conceptualization, whereas taxonomies classify empirical cases by observed phenomena. These definitions have been given for the sake of rigour, however, in practice, these semantic structures are often interchanged and may result in intermediary structures [31] —- like design spaces in HCI. Design spaces began initially as non-heirarchical classifications, introduced as taxonomies [24]. However, in time, they developed hierarchies, like dimensions, and have been defined as typologies. Typologies seem to be a more accurate description, as the process of generating design spaces involves morphological analysis [2], which results in a typology.

2.1.3 *Concepts in HCI*

Concepts are strongly related to categories and pattern identification, which represent a fundamental aspect of HCI. Categorization spans various methodologies, from qualitative analysis to systematic surveys and usability studies reporting. Researchers are often explicit with the outcomes of this research, defining taxonomies or typologies [13].

In HCI, Höök and Löwgren [129] introduced *strong concepts* as *intermediate-level knowledge*, as a type of knowledge that *"carries a core design idea which has the potential to cut across particular use situations and perhaps even application domains."* Strong concepts abstract instances and behave like generative designs (think of interaction metaphors)[2]. Dalsgaard and Dindler [152] introduce *bridging concepts* as a different form of intermediate knowledge that brings together both theory and practice. Bridging concepts explain the theory, practical *exemplars* parameters of the concept, and theoretical background. Exemplars here refer to artifacts that illustrate critical or salient aspects of a concept.

Exemplar and prototype theories have influenced many faces of HCI design.

2.2 DESIGN SPACES

Design spaces are highly debated within HCI, but generally refer to structures obtained by way of deriving categories from different types of data to inform future generations of applications. Eriksson et al. [265] highlight that design spaces define intermediate-level knowledge **hook1012strong**. Classification, conceptualization, and conception blending are operations performed by HCI researchers when exploring design spaces. The outcome of design space is a semantic structure — a design space schema (in the design-thinking view) or a morphological field (in the semanticists view).

---

[1] The reader may refer to Tom Ritchey's ongoing work for the epistemology of these structures, which have been overly simplified in this thesis: here

[2] Höök and Löwgren [129] give the touch-as-slingshot interaction from the game *Angry Birds* as given as an example of how to abstract design.

Some research uses the term design space to refer to the hypothetical space of design and not the semantic structure, like, for example, Kim et al. [285]'s work with VR and autism.

2.2.1 *A brief historical context for design spaces*

The late 80s and early 90s marked a turn towards design as a way, a process, or a methodology. Researchers sought to understand and formalize *design rationale* during this time — which essentially meant revealing the design decisions, design options, and how these influenced design [22]. An account of this research is given by [247]. Various frameworks were produced to aid in explicitly representing the design rationale like Lee and Lai [243]'s *Decision Representation Language*, and MacLean et al. [29]'s *Questions, Options, and Criteria*, the latter of which formalized representing design rationale through a process called design space analysis (shown in Figure 1, C). For design space analysis, the authors viewed an artifact's features as options within a set, posing questions about the options in design it determined.

In 1990, Mackinlay, Card, and Robertson [25] published work instantiating *The Design Space of Input Devices*, an effort towards *"systematizing knowledge about input devices"* [24]. This initial iteration of the work used Baecker and Buxton [18]'s semantic analysis to build a parametrized space based on existing taxonomies, comparing *"the design of human-machine interfaces"* to *the design of artificial languages for communicating between human and machine*[3]*."* This work argued that such taxonomies may be analyzed in terms of *expressiveness* — abilities to describe a device's semantics referring to the *formal language theory in computer science* at the time [25]; — and *effectiveness* of deductive reasoning, which relates the result of taxonomy to theories of human performance (footprint, bandwidth) as *points* of the design space, and not the whole [24]. The design space was represented as a multidimensional Cartesian matrix with connected graphs that marked the properties of an input device.

Or descriptive power.

In the following year, Card et al. [26] (design space shown in Figure 1, A) formalized the methodology in terms of *morphological design space analysis* following Zwicky's morphological approach to idea generation [2] (shown in Figure 1, B). The purpose of this analysis was to provide an understanding of input devices as points in a space that was parametrically defined, leading to generating designs to be further tested. This work moves toward using the term *design space* to refer to the space of parameters. The parameters and values themselves were derived from taxonomies in the previous work [24]. The design space remained as a multidimensional matrix that represented points in the design space as connected graphs, between values of parameters. Later, in 1993, Nigay and Coutaz [30] put forward a design space for multimodal systems, which captured three dimensions[4] (vision, levels of abstraction, and use of modalities) and continued with a representation similar to Card's [26] mentioning it as a typology. Later, this representation would be known as a Zwicky box.

---

[3] As a branch of linguistics, semantics deals with investigating the meaning of language, and, more recently, semantic analysis refers to the process of establishing meaning in natural language processing, e.g., sentiment analysis; [249] provides a review.

[4] These works use dimensions and parameters interchangeably.

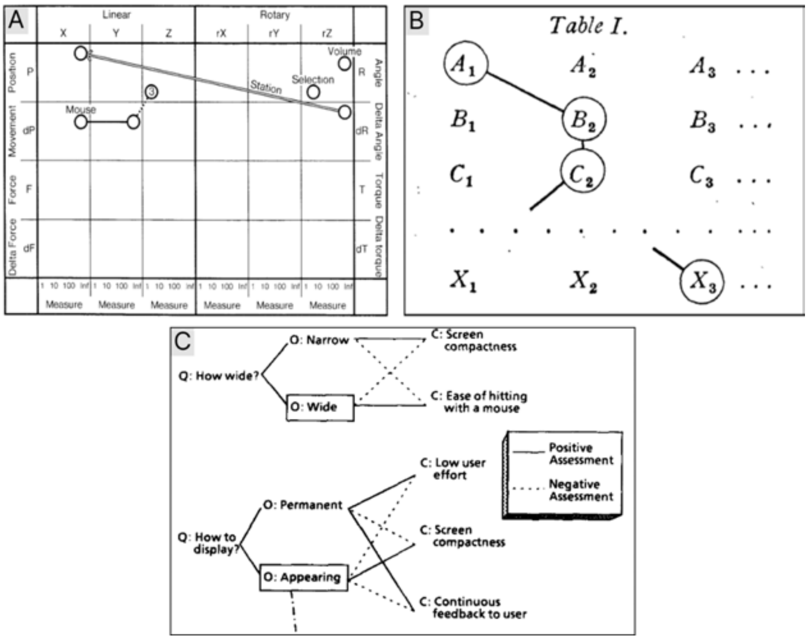

Figure 1: Subfigure A shows *Figure 3* from Card, Mackinlay, and Robertson [26]'s representation of the design space of input devices. B shows *Table 1* from Zwicky [136] original representation of the morphological box. C shows part of *Figure 3* from MacLean et al. [29], showing the questions, options, criteria representation of the design space.

The work underlying the design space of input devices is based on Mackinlay [17]'s graphical representations of relational data. He later applied the problem of generating and encoding such designs to graphical user interface generation [19]. He introduced the view that underlies the design space inception — that graphical language may be expressed as sentences encoding semantic meaning. He formalized this into criteria for expressiveness and effectiveness within a composition algebra to make alternative designs. This particular work is very influential in the field of information visualization. Mackinlay [19] concludes: *"Perhaps the ultimate lesson of this research should be that creativity and theory go hand in hand, one to inspire and initiate and the other to refine, test, and extend."*[19]

2.2.2 *Morphological Analysis*

In *New Methods of Thought and Procedure* [137] Fritz Zwicky introduced the morphological approach to invention [2]. This approach involved creating a morphological box containing all parameters of a given problem, later known as *Zwicky Box* [159] (shown in Figure 1 B). Zwicky, an astronomer and physicist himself [83], received various patents for discovering jet engines and propulsion mechanisms using this method. The morphological approach (later morphological analysis) was essentially an *extended typology analysis* [124] and provided an effective way to generalize ideas, bound problems, which could have resulted in "*systematic field coverage*"[2]. Zwicky considered the morphological approach to be a significant contribution to epistemology [6], enabling a paradigm which *"consisted in the replacement of*

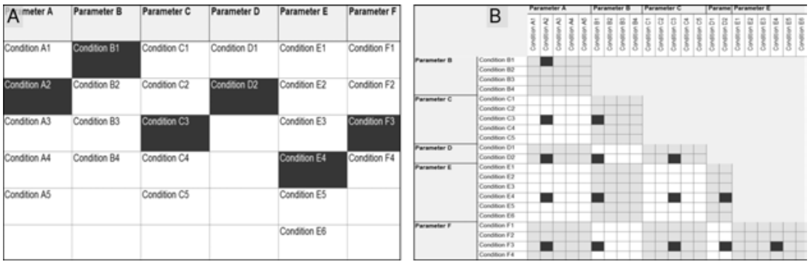

meter A	Parameter B	Parameter C	Parameter D	Parameter E	Parameter F
Condition A1	Condition B1	Condition C1	Condition D1	Condition E1	Condition F1
Condition A2	Condition B2	Condition C2	Condition D2	Condition E2	Condition F2
Condition A3	Condition B3	Condition C3		Condition E3	Condition F3
Condition A4	Condition B4	Condition C4		Condition E4	Condition F4
Condition A5		Condition C5		Condition E5	
				Condition E6	

Figure 2: Subfigure A shows the morphological field (*Figure 1* from [159]) and B shows the corresponding cross-consistency matrix (*Figure 2* from [159]).

one solution by all solutions, one path by all paths, one system by all systems"[6]. He describes the following steps to arrive at a morphological box:

1. *First Step.* Set the goal and problem.
2. *Second Step.* Find and describe the parameters of the problem.
3. *Third Step.* Make the morphological box/ multidimensional matrix of all solutions.
4. *Fourth Step.* Evaluate and analyze all solutions relating to the goal.
5. *Sixth Step.* Select and execute the best solution; this may require further morphological study.

Presently, the morphological approach has been formalized into *modern* General Morphological Analysis by Tom Ritchey [159, 124] as *"a general method for non-quantified modelling."* General morphological analysis results in a morphological field that categorizes the values of a given problem. *Parameters* capture qualitative and quantitative faces of the problems (not always mathematically), broadly within systems science nomenclature, like factors determining a system's variable behavior [124, 159]. General Morphological Analysis introduces a method to validate the field called a cross-consistency assessment. Its outcome, the cross-consistency matrix, approaches the modern design space representation (Figure 2 shows the matrix and the field). While the cross-consistency matrix cells contain types, design spaces cells map prototypes, systems, related work, etc. The parameters/values indicate the constraints of the mapped applications. With successful applications for *wicked problems* [124], Álvarez and Ritchey [159] describe general-morphological analysis as follows (paraphrased from their work):

The morphological field is virtually the design space schema in HCI terms.

1. *First Step.* Identify and define crucial parameters of the problem
2. *Second Step.* In natural language, assign each parameter conditions or ranges of relevant values.
3. *Third Step.* Construct the morphological field: create an n-dimensional matrix by putting all parameter values against each other on an X and Y axis (shown in Figure 2. Solutions are defined as a combination of values of each parameter.
4. *Fourth Step.* Examine all combinations of the parameter values to establish whether they are possible or not. This means analyzing their viability and assessing whether they are practical and, optionally, interesting. This viable combination of parameters represents the solution space of the problem within the morphological field.

5. *Fifth Step.* Perform cross-consistency assessment: all relationships between parameters are examined, and the contradictory relationships are removed. The idea here is to determine whether the pair of parameters can coexist and under which conditions. This reduced the problem space and the morphological field. The cross-impact or cross-consistency matrix captures this process, shown in Figure 2 B.
6. *Sixth Step.* Identify and mark the logical contradictions, empirical constraints, and normative constraints revealed by the assessment.

2.2.3 *From concept to prototype*

Prototypes are defined as *"a tangible artifact, not an abstract description that requires interpretation"* [86]. Beaudouin-Lafon and Mackay [86] explain design spaces as a set of ideas and constraints, which a designer may alter by changing the ideas based on the constraints or changing the constraints. Prototyping makes this process tangible and involves approaches like horizontal, vertical, task-oriented, and scenario-based, the last of which follows real-world scenarios. Particular to early prototyping stages, *"the goal is to generate a wide range of ideas and expand the design space, not determine the final solution"*[86]. Higher fidelity prototyping has a different goal *"video brainstorming expands the design space, by creating a number of unconnected collections of individual ideas, whereas video prototyping contracts the design space, by showing how a specific collection of design choices work together"* [86].

Fallman [64] discuss the limit of using design and prototyping tools and argue that design should be a dialogue rather than a series of events, whereby sketching, and by extension prototyping, hide a complexity. Lim, Stolterman, and Tenenberg [97] introduce a framework for conceptualizing prototypes and remark that prototypes go beyond enabling evaluations towards idea generation, particularly when exploring design spaces. They remark that prototypes may be used as filters to traverse design spaces to manifest concrete conceptual ideas. As filters, prototypes serve to generate and evaluate ideas to ultimately make better choices in design, grounded in dimensions like appearance, data, functionality, interactivity, and spatial structure. Lim, Stolterman, and Tenenberg [97] view design spaces as infinite and refer to their manifestations in terms of material, resolution, and scope.

Dove, Hansen, and Halskov [173] remark that the power of prototypes can be revealed by design-space thinking. Prototypes may instantiate a subset of dimensions or parameters for further evaluation. Keeping a writen record of these prototype changes in terms of a design space enables novel ways of reflecting on the design process [173].

2.2.4 *Design spaces in HCI*

Following a research-through-design approach, Halskov and Lundqvist [269] give an overview of the types of design spaces and the various knowledge they are based on. They propose *design-space thinking* to understand the creative evolution of the design process in time, refocused on the artifact and the *primordial soup* of knowledge that enables it. They categorize design spaces as representing a class of technologies, accumulated knowledge, or a set of ideas, designs, or sketches. They categorize the physical manifestations of these representations as a Cartesian space, a networked graph, or a conceptual space.

Biskjaer, Dalsgaard, and Halskov [149] define design spaces from the point of view of creative constraints, *"that govern what the outcome of the design process might (and might not) be."* The researchers propose a design space schema containing options nested within aspects to represent design spaces. The design space schema is the *design-thinking* nomenclature for the morphological field; it refers to the various dimensions, parameters, and values the design space contains. Design spaces are conducive to reasoning because they reflect the reasoning process [29]. As idea-generation tools, design spaces are facilitated by brainstorming, a common component of participatory design and group-work [86].

In HCI, design spaces have two main functions: (1) they have a *theoretical and conceptual function*, since design spaces aid in generating new ideas by identifying gaps in knowledge and proposing ways to fill them [261, 220, 196]; (2) design space have *practical function* because they reveal options necessary to by implementing and conceptualize novel applications, which are then typically evaluated for performance [146, 39] and usability [237, 150]. The process of using design spaces to classify related work and uncover gaps is sometimes referred to as design space *exploration.* For example, Danyluk et al. [261] generate a design space of worlds-in-miniature for VR based on open coding from 25 papers. Drey et al. [237] create a design space to explore the combinations of 2D and 3D sketching, categorizing related work and generating novel interaction metaphors. Eriksson et al. [265] generate a design space based on the collaborative efforts of 50 developers that implemented *4in1 games*. Eriksson et al. [265] generated the design space in four phases: horizontal analysis (concept synthesis and generation), vertical analysis (relating to theory), inductive analysis where independent coders clustered concepts, and testing the generated design space by re-coding and re-generating designs from the initial games. MultipleAvatars and Feedforward apply similar steps to design space generation, namely theory triangulation and re-coding.

Note that design space exploration is also a term used in systems design and architecture to synthesize model solutions based on an algorithm for further prototyping and testing [42]. These works contribute with novel algorithms that optimize the design space and involve mathematical concepts like automated theory proving and combinatorial optimization [122]. In this field, design spaces have mathematical constraints that can be specified in arithmetic or boolean operations.

In HCI, however, design spaces capture concepts. The exploration of the design space is more often understood as an inherent process of design-space-making. For example, Hirzle et al. [220] generates a design space for 3D gaze interaction for head-mounted displays and argues for its feasibility by using it as an idea generation tool and to classify various applications and interaction techniques. Halskov and Fischel [217] specify other methods for design space analysis in six broad steps:

1. *Step one.* Defining the scope.
2. *Step two.* Getting the data from the empirical materials.
3. *Step three.* Making a coding schema and coding the data.
4. *Step four.* Analyzing the data.
5. *Step five.* Reporting on the result.

2.2.5 *Generative interactions*

Many of these works borrow from Beaudouin-Lafon [70]'s framework for generative interactions. In his initial work describing this model, Beaudouin-Lafon [52] proposes that post-WIMP[5] interfaces should be expressed in terms of interaction models represented as design spaces that reveal their properties for further analysis and comparison. Revealing this design space enables other designers to *"make an informed choice"*. To inform these novel interfaces, the models should be *descriptive*, *generative*, *comparative* or *evaluative*. These models should (1) incorporate both novel and existing techniques, (2) reveal properties to compare alternatives, and (3) enable the generation of novel interaction techniques. He provides *temporal and spatial offset* and examples of properties that enable comparing different techniques. In later work, Beaudouin-Lafon [70] reframes this research as a call to arms to design for interaction. He adds the caveat that good interaction models may still evoke *"terrible interfaces"*, and design spaces do *"not guarantee"* quality in the resulting design [70].

Later, Beaudouin-Lafon, Bødker, and Mackay [257] propose *generative theories of interaction* to analyze and generate novel technological artifacts, defining concepts and actionable principles. The main purpose of such theories would be to *spark innovation*, guide novel designs to meet users' needs. Generative theories underlie a design space and provide *"a path from theory to artifact and a principled method for exploring the research design space."* [257]. The researchers define generative theories as follows:

1. Underlying human activity and technology theories,
2. Grounded in critical, constructive, and analytical lenses,
3. Actionable by way of the concepts and generative principles that specify the theories.

While the discussion refers to models and theories, design-space makers have adopted the evaluation framework because design spaces represent the parameters of the model underlying a generative interaction. Therefore, interaction generation evaluative criteria apply. Beaudouin-Lafon [70] proposes three criteria for generative interactions: descriptive, evaluative, and generative. Some works refer to the evaluative criteria as *comparative* [261] to prevent confusion. Beaudouin-Lafon [70] does not refer to implementation, technical view of evaluation (e.g., usability study), rather this criteria refers to the *theoretical* comparative power of interaction models.

2.3 CONCEPTUALIZING IN VIRTUAL REALITY

Several central concepts border the virtual reality part of HCI: embodiment, immersion, and presence. A large body of work incrementally builds on these concepts and blends them unexpectedly. Exploring the boundaries of these concepts has given rise to novel interaction metaphors, like virtual hand metaphor for manipulation [46], virtual pointer metaphor [85], or combinations of gaze and gesture metaphors like gaze and pinch [192]. These interactions leverage *embodiment* — sense that another body (e.g., avatar) is processed as if it is the user's real body [131]. *Presence* refers to the illusion of *being there* in the virtual environment, while *immersion* refers to the

---

[5] WIMP: windows, icons, menus, pointer.

technical qualities of a system [292]. Expanding the limits of these concepts, researchers have studied how people may embody animals in VR [224], or blend real and virtual spaces for physical meetings [298].

Even in the 90s, researchers found VR useful as a concept design tool. The 'unconstructive' capabilities of VR were more appealing than Computer-Aided Design (CAD) tools at the time because of the lack of design restrictions (see *CORVIDS* [40]). VR served as a test-bed for interactions by enhancing conceptualization through prototyping. Authoring tools are a type of prototype to enable quick design deployment for testing before production. Despite this, Nebeling and Speicher [205] remark that VR and AR development poses many technical challenges, particularly from tools that are not suitable for non-technical designers. Researchers have developed friendlier authoring tools for XR to close this technical knowledge gap and create faster prototypes in VR. For example, Jetter et al. [241] implemented and evaluated a prototyping environment where designers may interact with digital twins before making tangible prototypes. Other XR works enabled video prototyping with sketching and generated animations that prototype interactions and real-world scenarios [244, 270]. In a recent review of how VR aids in design-thinking, Lyu et al. [303] highlights that defining design-thinking would leverage the full potential of VR as a creativity tool.

Concepts are also relevant for VR design. Following the path of related work reveals how concept blending generates novel designs and prototypes in VR. For example, building sense-making through video recordings in VR [245] and the concept of worlds-in-minuiature [37] resulted in spatial design queries with direct manipulation [304]. Since design spaces serve to generate and categorize concepts, they are suitable as the default semantic representation for prototyping in VR. Design spaces have enabled concept generation and exploration for tangible devices and XR [308], supported visualization tasks with multiple displays in AR and VR [299], generated designs for 3D sketching with pen and tablet interactions in XR [237], enabled exploring data transformation in mixed reality [288].

Recent research has introduced various conceptually-driven interaction techniques and interfaces like asynchronous reality [282], blended reality [298], and remixed reality [203]. Among these, some works highlight their novelty while building on WIMP[6] interface concepts. For example, *OVRLap* [291] prototypes interacting from multiple places using multiplexing viewpoints derived from CHI'95' *Transparent Layered User Interfaces* [34]. Danyluk et al. [261] generate the design space of the worlds-in-miniature in VR to uncover novel interaction possibilities and designs. First introduced at CHI'95 [37], the world-in-miniature metaphor is an immersive technique that augments the user with a miniature version of the virtual environment.

2.4 CONCLUSION

In this chapter, we traced the origins of concepts and design spaces. We presented the prototype and exemplar models of concepts from cognitive psychology and tied them to HCI *prototypes*. *Design space filtering* [97] bridges the gap between the *disciplined imagination* and tangible artifacts — prototypes are developed to investigate particular concepts as dimensions of design. In this context, the design process refers to instantiating the prototype,

---

[6] WIMP: windows, icons, menus, pointer.

and, as Fallman [64] decries, the process *"seems often obliterated from descriptions of research projects; research prototypes just seem to happen."* In this work, we formalize this process in terms of design spaces and further explore the design space schema's usefulness as a model for generative interactions, following Beaudouin-Lafon et al.'s path from theory to artifact [257]. Together, design spaces and prototypes define a feedback loop of design: *"what the prototype states is subject to evaluation; what the prototype leaves open is subject to more discussion and design space exploration"*[86].

While the process of prototyping nests a design space itself, the tangible artifact instantiates the unique dimensions, parameters, and values of the design space. Prototyping as a process allows for concept exploration. However, authoring tools enable design space exploration in practice. Design spaces tell a story of how prototypes happen because they reveal the boundaries of design and the concepts underlying it. Authoring tools reveal design spaces in practice. Prototypes enable the exploration of concepts by filtering and informing design spaces to manifest design ideas concretely [269]. Design spaces also allow researchers and practitioners to record their design process and communicate the basic blocks of their interaction designs at a glance [**dove**].

3

ANALYSIS OF VR FAILS ON YOUTUBE

This chapter is based on and reproduced from VRFails, entitled *Bad Breakdowns, Useful Seams, and Face Slapping*, published as [262]: *Andreea Muresan*, Emily Dao*, Kasper Hornbæk, and Jarrod Knibbe. "Bad Breakdowns, Useful Seams, and Face Slapping: Analysis of VR Fails on YouTube." In: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. CHI '21: CHI Conference on Human Factors in Computing Systems. New York, NY, USA: ACM, May 6, 2021, pp. 1–14. isbn: 978-1-4503-8096-6. doi: 10 . 1145/3411764.3445435. url: https://doi.org/10.1145/3411764.3445435*[1].

Emily and I worked together wonderfully on this research paper, which, for me, serves as a model for how to conduct successful research. This paper received a *Honourable Mention Award*, given to the top 5% of papers[2]. This work reveals the social lens of conceptualizing.

Figure 3: Figure showcasing key moments in three VR fails clips. In comic strip (a), the VR player loses balance while playing a rollercoaster game in a shopping mall (ID157). A spectator tries to hold the player up. In comic (b), a user is strapped into a wired vest that is held by a metal bar from the ceiling (ID104). They are playing a VR shooter game and scream finding themselves overwhelmed by their virtual opponents. The attendee notices their distress and tries to take the headset off. However, the player responds with more screaming and backs away, afraid and surprised by the unexpected touch. In comic (c), a user is playing a VR game at home and throws their controller across the room to reach a virtual object of interest (ID65).

3.1 ABSTRACT

Virtual reality (VR) is increasingly used in complex social and physical settings outside of the lab. However, not much is known about how these settings influence use, nor how to design for them. We analyse 233 YouTube

---

[1] * = Shared first authorship

[2] The process for selecting and awarding papers is explained here: https://chi2021.acm.org/for-attendees/highlights/awards

videos of *VR Fails* to: (1) understand when breakdowns occur, and (2) reveal how the seams between VR use and the social and physical setting emerge. The videos show a variety of fails, including users flailing, colliding with surroundings, and hitting spectators. They also suggest causes of the fails, including fear, sensorimotor mismatches, and spectator participation. We use the videos as inspiration to generate design ideas. For example, we discuss more flexible boundaries between the real and virtual world, ways of involving spectators, and interaction designs to help overcome fear. Based on the findings, we further discuss the 'moment of breakdown' as an opportunity for designing engaging and enhanced VR experiences.

3.2 INTRODUCTION

As Virtual Reality's popularity grows, we need to better understand and design for its use outside of VR laboratories. In contrast to research labs which can be easily controlled with specific furniture or layout configurations [202], domestic spaces are messy, busy, and dynamic. How VR plays out in those spaces and the seams and breakdowns within those interactions remain poorly understood.

With VR's increasing adoption, we are also seeing the rise of community content showcasing its use. YouTube is rife with popular VR gamers completing difficult levels in *BeatSaber*[3], playing haunted-house style games, or showcasing gameplay fails. Prior research has examined YouTube content for insights into technology adoption and interaction design [138, 223]. From this, we suggest YouTube VR videos may provide one lens through which we can better understand how people engage with VR beyond the lab.

We examine *VR fails* on YouTube, as one of the popular emerging themes of videos. These videos capture people in VR colliding with furniture, falling over, hitting spectators, and experiencing fear, joy, nausea, surprise, and more. The videos also capture spectator reactions and participation.

We analyze the VR fails corpus from two perspectives. On the one hand, we think about VR fails as *breakdowns* [33] - moments of stark disruption in the experience. We seek to understand where and why these breakdowns occur, to reveal design opportunities for addressing and solving them. On the other hand, we consider VR fails as Seamful Design [63] - exploring how the technical, experiential, and social aspects of VR are stitched together, and how resulting 'fails' may become a positive feature of this shared experience.

We begin with a thematic analysis of 233 VR fails clips, revealing themes across *types* of fails – such as excessive reaction, colliding, hitting, and covering, *causes* of fails – such as fear, sensorimotor mismatch, and spectator participation, and *spectator reactions* – such as laughter, concern, and support behaviours. We discuss each theme in depth, present examples, and highlight how they co-occur. Based on our analysis, and motivated by the examination of VR fails as both breakdowns and an exposé of seamful design, we describe new design opportunities that VR fails inspire. For example, we speculate about how head-mounted pico-projectors may enable shared play and reduce collisions, and how 'natural' gestures may enable novel peak-through dynamics for scary experiences.

VR fails provides a lens through which we can better understand the implications and opportunities of VR use beyond the lab. We highlight what

---

[3] BeatSaber, BeatGames, May 2019

this media shows and how we may incorporate these insights into our interaction design process in HCI.

3.3 RELATED WORK

As VR emerges commercially, research explores how to further develop the experience. Recent research, for example, has looked at novel controller opportunities to enhance haptic feedback [210, 225], new techniques to improve locomotion through virtual spaces [178, 209], and new opportunities for shared virtual experiences [186, 208]. Simultaneously, research seeks to better understand the existing experience of VR. We draw from several strands of prior work on presence and the moment of exit, spectator engagement, and VR in the real world and YouTube as a data source.

3.3.1 *Presence and Breakdowns in VR*

Prior work has highlighted the central challenges associated with VR, such as the effects of simulator sickness [51], ocular fatigue [219], social anxiety and awareness [202], and collisions [218]. Slater et al. [53] sought to understand the breakdowns in experience that occur as a result of these challenges. They counted moments of breakdown as an inverse measure of presence[4]. This clearly positions breakdowns as something to be minimised and avoided, so as to maximise presence. Conversely, despite the many challenges VR breakdowns may pose, studies have shown that they do not necessarily diminish the player's experience, but may instead enrich it [60]. The popularity of the VR fails phenomena on YouTube may support this finding, highlighting the entertainment generated by the breakdowns.

Knibbe et al. [202] explored the variety of factors that impact the transition between virtual and real environments that may result from a breakdown, with a focus on the final 'moment of exit'. They explored four example applications and found that participants experienced spatial and temporal disorientation, control confusion, and heightened social awareness upon exit. Their work suggested the existence of a social contract of VR play, where players discussed the need to feel safe during play and were especially socially-aware at the moment of exit. Kappen et al. [155] echoed the sentiment around social contracts in play, highlighting how the interaction between players and spectators can increase engagement. This motivates a desire for VR to support awareness of 'others' and their relative 'proximity' during play [166]. While research is exploring ways of including passers-by in the VR single-player experience, [e.g., 233], we expect VR fails to reveal more nuance around the social contract of interactions between players and spectators.

3.3.2 *Spectator Engagement with Play*

Adding to existing work on *social contracts* of play, research has also sought to reveal the different ways in which spectators can engage with play in general [195, 78]. Tekin and Reeves [195] outlined different kinds of spectating in Kinect Play: (1) 'scaffolding' play to seek to display continuous

---

[4] where presence and immersion are a typical goal in VR; providing a technical foundation and experience such that the player believes they are *there* in the virtual world

engagement with the player, such as by providing timely instructions, or encouragements, (2) critiquing play technique and gaming movements by the player, (3) recognising and complimenting 'good play', and (4) reflecting on past play (as a former *player*).

But spectator engagement not only impacts the player's experience, it also directly impacts gameplay. For example, Reeves et al. identified a performative aspect to co-located gaming, where players perform extra gestures that do not directly impact the interface, but instead are for the audience [78]. They present four design approaches to designing public interfaces: (1) 'secretive', where manipulations and effects are mostly hidden, (2) 'expressive', where the performer's actions are revealed to spectators, allowing them to appreciate the performer's interaction, (3) 'magical', where spectators see the effects but not the manipulations that caused them, and (4) 'suspenseful', where spectators only understand the manipulations behind the visible effects when they become the performer themselves. According to this, VR fails sit across different scales of manipulation. The videos are *expressive* when players perform extra movements to be visible. VR fails are *entertaining* when they result in players failing or losing balance, and they are *magical*, as is the case when players fall over despite not making clear movements.

This performative aspect of co-located gaming is mirrored by the *staging effect*, which occurs when technology use in public spaces creates a performance stage for the user [1]. Dalton et al. [161] echoed Reeves' findings that in the presence of onlookers (e.g., at a shopping centre) some users interact with technology for the sole purpose of being noticed by others. However, the public stage can also result in a negative experience, where users may avoid interactions altogether to prevent social embarrassment [171].

While this body of work undoubtedly deepens our knowledge of virtual reality, it is still early days for VR research outside of the lab [190], and we have only limited understanding of how VR experiences come to be enacted between players, spectators, and their environment.

3.3.3 *VR in the Real World*

Recent work has investigated head-mounted display (HMD) use in public settings and the validity of in-the-wild VR studies.

In an in-the-wild deployment of HMDs in a public setting, Mai and Khamis [204] investigated the parallels between interacting with public HMDs and more traditional public displays. They found that recognition of the HMD as an available public display was a challenge. This is partly due to the visual clutter of real spaces, but also the black-box, enclosed, nature of the headset (even while coupled to glowing controllers). Some participants were unclear on the connection between the separate display screen and the HMD, while others searched the demo space for an official authority to allow them to interact with the HMD. Mai and Khamis put forward suggestions for future solutions for accompanying displays to communicate HMD usage and functionality, and invite passers-by to interact without the need for an authority present.

As seen within the context of public settings, HMD use is not only influenced by known VR factors, such as immersion, but also by real-world factors such as perceived boundaries around permission of use. George et al. [216] explored non-experts' mental models and future expectations towards mobile virtual reality (MVR) through a field study using drawing tasks, a storytelling exercise, and the technology acceptance model (TAM). Their work revealed participants' struggle to balance wanting to be immersed in the experience and their concern over possible dangers in the real world, such as falling over cables, etc. Furthermore, participants stated their preference for VR use in the home, rather than in a public setting, due to fear of being observed by strangers and bystanders. However, the public setting was more desirable if participants are accompanied by friends and family. These findings highlight the need to consider factors beyond technical capability when considering designing public HMD experiences.

Mottelson and Hornbæk sought to validate the potential of conducting VR studies in the wild [190], comparing in-lab to out-of-lab VR experiments over three canonical tasks. Their results show that the effects found in the laboratory were comparable to those found in the wild, suggesting that conducting VR studies outside the laboratory is feasible and ecologically valid. Steed et al. [182] pursued a similar angle, running an out-of-lab experiment for mobile app-based VR devices, such as Google Cardboard and Samsung Gear VR, to study presence and embodiment. While their experimental results supported their hypotheses and validated their methodology, they reflect on the additional challenges and work involved in conducting this style of study, such as the required level of polish, ease of use, completeness of experience, etc. As VR headsets become increasingly common-place, these in-the-wild studies will become increasingly practicable.

3.3.4 *Understanding the User on YouTube*

Recently, YouTube has also been used for insights into the real-world use of emerging devices [101, 163] and specific user contexts with technology [223, 138, 139]. Prior work has highlighted the value of publicly available, user generated content in informing input and interaction design in HCI [101, 138]. For example, Anthony et al. [138] conducted a content analysis of 187 YouTube videos to explore touchscreen use by people with motor impairments.

More recently, Komkaite et al. [223] analysed 122 YouTube videos to gain insight into users' interaction with common non-medical insertable technologies. We argue that such a methodology, applied to the growing corpus of online user-generated VR videos, may also provide real-world insights for VR research, and contribute to the ongoing discussion of VR in-the-wild.

3.4 EXPLORING VR FAILS

VR fails videos capture momentary transitions between the virtual and real worlds, and the associated reactions of spectators. These momentary transitions occur as players collide with objects in the real world, interact with non-VR spectators, and experience strong emotions, such as fear. Previously, these momentary transitions have been treated akin to breakdowns [53]. In breakdowns, technology becomes visible to us in a new way, failing to function as we anticipate [33]. Seen as VR breakdowns, fails may indicate usability problems; representing mismatches at an interaction level, which cause the user to express surprise or uncertainty of how to engage with the technology. Thus, fails may reveal errors in the VR experience that should be addressed.

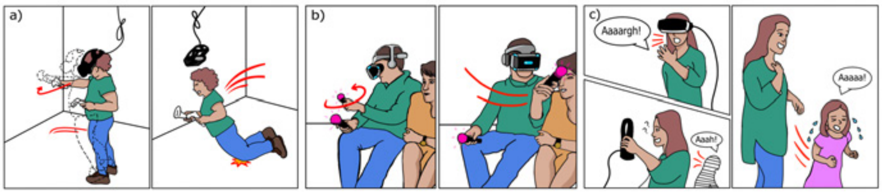

Figure 4: Illustrations of key fails clips. Comic strip *a)* illustrates the video (ID38) in which a young player is slowly rotating their controller to manipulate a virtual object in front of them. Eventually, the player leans far enough to fall forward and, as they fall, the headset partially arrests their fall, but releases just in time for their face to collide with a wall. Comic *b)* shows a fail (ID20) in which the player moves a controller around themselves, eventually hitting the person next to them. The spectator was not paying attention to the controllers. In comic *c)* we capture the fail (ID55) in which the player screams, seemingly scared by something from the VR game. They take off their headset and notice that the child behind them was scared by their yelling and has started crying.

However, even a brief exploration of VR fails highlights that these moments are not solely problematic and that they may provide an insight into the motivation and enjoyment of VR use. As such, we also explore the VR fails corpus as an insight into Seamful Design, whereby breakdowns are seen as a 'resource' rather than a 'system failure' [63]. From this view, we aim to understand how the complex interactions between personal VR, shared spectator experiences and physical settings play out as enjoyable events within a VR experience. By understanding these moments, VR designers and technologists may come to add them to their repertoire as design opportunities and elements for engagement.

From these two perspectives, we conduct a video content analysis of the most popular VR fails clips on YouTube. We do not attempt to create the 'fails' classification, nor reconsider whether videos within the corpus are 'fails' per se. Instead, we accept the fails as a community-driven label and instead seek to understand what we may learn from these videos.

3.4.1 *Methodology*

Our study includes two phases: (1) video searching - finding user-generated VR fails videos on YouTube, and (2) performing video content analysis (VCA) [74] on the final set of videos.

3.4.1.1 *Phase 1: Video searching*

We used YouTube's search to identify relevant videos. We initially defined our search term as 'VR fails', the accepted community-driven term.

We explored synonyms of 'fails' such as 'break', 'failure', 'crash', 'malfunction' and 'stop'. However, these returned results that are not within our scope, such as scenes from various VR games or troubleshooting videos, and as such were not included. Furthermore, we tried to expand our search by including augmented reality ('AR'), another technology close to VR, and searched for 'hilarious' to capture funny moments caused by fails. However, we yielded no additional, nor relevant results. Thus, our final search term remained at 'VR fails'.

Initially, all the videos were screened by one of the researchers. We recorded the basic information about the videos, including the title of the video, the webpage link, frequency of viewing (the number of times a video had been watched) and length of the video. We sorted the search results in descending order by the number of views. The majority of the videos in question are compilations. We treated each clip within a compilation as its own video for analysis. If a clip with identical content appeared multiple times across different compilations, we only included it once.

At the time of searching[5], the most popular video had 3.067 million views. We looked at all returned results above 27,000 views, as the high number of views demonstrates sufficient community acceptance of the video. Additionally, as the number of views decreased, we observed higher frequencies of duplicate videos in the compilations. We excluded clips of streamed gameplay (such as on *Twitch*), where the VR player is not in view, and one 'prank' clip in which someone watched pornography in VR on a subway.

We also excluded videos that are less than 2 seconds in duration, as it is difficult to understand the fail within that short time frame. In total, we found 382 video clips, across 32 compilations, that depicted users experiencing VR and failing. After removing duplicates across the videos, our final data set is 233 videos. This number is similar to the upper range of other YouTube analyses [138, 101, 223, 139, 163]. The clips are 15.83 seconds long on average (std. dev. 10.2s).

3.4.1.2 *Phase 2: Video content analysis*

Video content analysis is a well-known methodological procedure for studying qualitative data and is frequently used in studies of mass media [74]. First, we developed a coding scheme that focused on user interaction with VR and the specifics behind the fails. This was an iterative process, in which all authors viewed the same subset of 16 video clips and refined the coding scheme through discussion. Next, we analysed a further 20 video clips per our coding manual and ensured inter-rater reliability (Fleiss Kappa [4] of .816 for *Types of Fails* and .610 for *Cause of Fails*, indicating high-agreement across four coders [10]). Finally, we coded all video clips in our data set along the coding dimensions (3), as inspired by Anthony et al. [138] and Komkaite et al. [223], with each clip coded by at least two authors. Disagreements were discussed and resolved.

3.5 WHAT DO WE SEE IN VR FAILS VIDEOS?

This section provides an overview of how we think about fails and the types of fails we see in our data set. We also highlight the interaction styles that we observed between the user-in-VR and their spectators.

From the videos, we interpret fails as a clash at the intersection between the virtual and real worlds. This may include physical clashes between the player and real-world objects, or social clashes between the players' actions and spectator expectations. Importantly, fails are determined through the spectator's perspective, and may not always be experienced as a fail by the player. We identify the moment of failure as either (1) the point of greatest

---

[5] We conducted our Youtube search on 15th May 2020

Video characteristics

Type of VR application
Physical context: e.g., home, store
User's interaction with VR: head-track, hand-track or body-track

Fail characteristics

Point of failure (timestamp)
Type of fail
Cause of fail

Spectator involvement

Spectator visible/audible: yes or no
Secondary display for spectators: yes or no
Spectator interaction before fail
Spectator reaction in response to fail

Table 3: The final 10 dimensions used to code the VR fails videos (inspired by Antony et al.[138])

reaction by the audience, or (2) the moment of intersection between the player and their physical surrounds.

We analysed 233 clips of VR fails. The majority of the clips (61%) were filmed in a private setting, featuring sofas, dining tables, bookshelves, and daily clutter. The other 39% were in public spaces such as in conference booths, in the central walkway of a shopping centre, anchored to a sales cabinet in an electronics store, or in a workplace surrounded by colleagues. The vast majority of videos include one or more spectators (including those capturing the video). These spectators are typically stood or sat around actively watching the player and sharing in the experience. The camera is typically focused primarily on the player, but sometimes also includes a VR view (through an additional display - 53%) for the spectators.

The most common headsets were head-tracked only (53%), where the engagement is primarily visual and not embodied, such as with a Google Cardboard or Samsung Gear VR. The second most common type of headsets were full-featured (35%), such as the HTC Vive, followed by hand-tracked (12%), such as the Playstation VR.

One of the fails clips had an educational purpose (a car mechanics training setting), whereas all others focused on game mechanics. The most common types of games were experiential, making the most of VR's immersive nature, such as walk-the-plank and rollercoaster experiences (23%), followed by horror games (18%). We also see other types of games such as action (12%), sport (11%), and others (like racing and puzzle games, 4%). In many of the cases (30%), it is unclear which application is being used due to the camera angle of filming or lack of additional display.

3.5.1 *Types of Fails*

Through our analysis, we identified six types of fail: Colliding (9% of clips), Hitting (10%), Falling Over (18%), Excessive Reaction (53%), Covering (7%), and Other (3%). Some clips result in players 'exiting' the virtual reality experience (i.e., removing the headset), while others demonstrate players recovering or continuing with their experience. In the following, we present the six types of fail, together with highlights of the content that we see across these types and their relation to exiting. We provide example clip IDs as we go, the links for which can be found in the supplementary material.

3.5.1.1 *Colliding*

Twenty videos show people *Colliding* in VR. These videos include people walking into walls (ID179) or furniture (ID177). In one clip, we see the player quickly walk into an upright mattress (ID11), as they dash to interact with a virtual game element. This mattress was positioned as a protective barrier in the space (a room in a home), perhaps suggesting that running towards that position was a common occurrence. In another example, a player stands on the end of a virtual and physical plank (i.e., present in reality and in their game), before taking a large step off the plank and straight into a glass door (ID144) (Figure 5b). Notably here, the player is being encouraged on by a spectator, telling them *"...you should just jump off. Just jump."*, who surely should have preempted the collision with the door. This is indicative of many clips, where the spectator appears more knowledgeable of the experience.

Only a small portion (3/20) of collisions result in players exiting the virtual experience (e.g., ID62, where damage to the headset necessitates exiting). In most cases, players seemingly recover and continue (e.g., ID177), or, at the very least, the video does not show exiting (i.e., ID148).

3.5.1.2 *Hitting*

Twenty-three videos show VR players actively *hitting* things in their environments. (This is different to passively *colliding* their hands with objects in the environment). Players hit walls (e.g., ID21), furniture (ID13), and spectators (ID20). In clip ID218, we see a man in a school hall walk forward approximately 5 meters and hit a wall with his controller - breaking his controller in half. This clip is notable as few commercial VR applications require that range of locomotion, yet players often interpret virtual spaces as fully explorable (i.e., walkable). Further, in this example, the game itself is music-based (possibly *BeatSaber*), and the crowd are participating by singing along. This is a rare example of indirect participation in the data set, where the spectators are deriving enjoyment from a feature of the game, rather than from the player's experience.

In clips ID220 and ID20 (Figure 4b), the VR player inadvertently hits (punches) a nearby spectator. In both of these cases, the spectator is watching a display of the VR players' view. Even with this shared view, the spectators do not preempt the movements the VR player is likely to perform. In ID220, this has quite serious consequences, as the spectator is knocked out by the VR player. Unaware, despite the commotion from the other spectators, the immersed user continues playing. This unawareness of spectators in the shared physical environment is consistent with Mai et al.'s [204] findings,

Figure 5: Illustrations of key VR fails. Comic *a)* shows the clip (ID73) in which a spectator suddenly touches the player's waist in an attempt to scare them. In response, the player turns around immediately and slaps the spectator. In comic strip *b)* we capture the fail (ID144) in which a player is standing on a plank in what appears to be the middle of their living room. The player is looking down and a spectator instructs them to *just jump*. The player takes a step and makes the jump but hits the glass door in front of them. In comic *c)* we illustrate the fail (ID230) in which a user is playing a VR game holding their arms up. As they reach further in the virtual game, they hit the ceiling fan light. Comic strip *d)* illustrates the fail (ID12) in which the player is surprised and scared by a virtual element in the game and steps backwards, colliding with a painting that causes it to fall off the wall. A spectator notices the falling object and scolds the player for damaging it, adding that it is an expensive painting.

where all participants forgot about their real-world surroundings, including spectators, after some period of immersion.

In 6/23 cases, the player exits VR after hitting something.

3.5.1.3 *Falling Over*

In forty-three of the VR fails clips players fall over. This primarily occurs during locomotion (e.g., 58) or because the game does not fully support an intended action (e.g., 40). In 13 of these cases (30%), this results in the player exiting VR.

In clip ID22, ID222, and ID229, the player is performing the required action (playing snooker, climbing, and inspecting an engine, respectively) as they would in the real world. While the applications encourage that action, they do not fully support their execution. For example, in real-world snooker a player will often lean on the table; a climber will use the wall to counterbalance; and a mechanic will lean on the edge of a vehicle for support. In these clips, we see the players attempt to execute this standard co-action in the virtual space, resulting in them losing their balance. We describe this as a *False Signifier* (which we will refer back to when discussing *Causes* of Fails: 3.5.2.5).

Clips ID38 and ID29, show the player falling over during locomotion.

In ID38, the VR headset is anchored to the ceiling. As the player falls, the headset partially arrests their fall, but releases just in time for their face to collide with a wall (figure 4a).

3.5.1.4 *Excessive Reaction*

The most numerous type of fail we see is 'excessive reaction' (123 clips). These clips involved players experiencing heightened reactions to VR, including physical reactions (38 clips) such as flailing (throwing their arms and legs around, e.g., ID77) and falling into an observer for physical support (ID94); vocal reactions (31 clips) such as screaming (through fear, ID142, or joy, ID75), and combinations of these (54 clips)(e.g.,ID46). We also saw 11 instances of players fully engaging with the experience (e.g., fast sprinting on an omni-directional treadmill - ID113, or boxing very enthusiastically - ID89), in a manner that is simultaneously very entertaining for spectators. We suggest these excessive reactions are frequently fails because they fall outside of expected social norms of behaviour and are unexpected by the spectators (including participants' pets, for whom the behaviour is especially bizarre - e.g., ID54). This is described by Simon [111] as *gestural excess*, a form of movement often seen in body-based gaming that would cause ridicule and shame outside of the envelope of gaming, but create its kinaesthetic pleasure to the player in-game. Notable here, these gestural excesses are not usually experienced as fails or breakdowns by the players, but occur as fails only for the spectators.

In 15% of clips, this results in the player exiting VR. This is most frequently a result of fear. We will return to a discussion of this below, in *Causes of Fails*: 3.5.2.1.

3.5.1.5 *Covering*

Our dataset includes 17 examples of people *Covering*. In these clips, players adopt protective poses, whether tucking into a ball (ID103) or covering their face (in ID51, the player covers her head with a blanket and in ID211, the player uses a blanket to constantly hide behind), in response to the in-game content. Again, as a result of fear, three clips result in players exiting VR.

3.5.1.6 *Other*

Our VR fails dataset contains six outlier clips, which do not fit easily into our other categories or are not sufficiently numerous to warrant their own category. These clips include two instances of players accidentally throwing their controller at a wall (ID65 (Figure 3c) and ID153), three clips of spectator fails (ID130 - where a spectator tries to hide in the corner so to not get hit by the player, see Figure 7), and one clip where the VR headset falls off (ID72 - resulting in exiting).

3.5.2 *Causes of Fails*

From a design perspective, we are interested in understanding both *how* VR players fail and *why* players fail. Across our six types of fails, we found seven corresponding causes of fails: Fear (40%), Sensori-motor mismatch

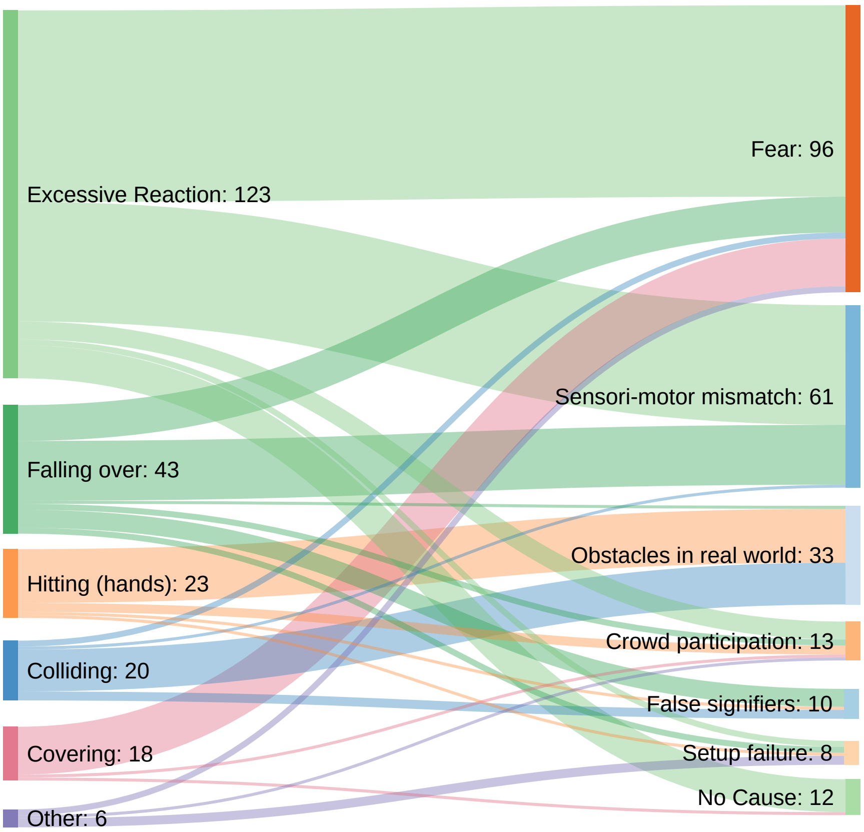

Figure 6: Relationships between Types of VR fails (left) and Causes of VR fails (right).

(26%), Obstacles in the Real World (14%), Crowd Participation (6%), False Signifiers (4%), Setup Failure (3%), and No Cause (6%) (Figure 6).

3.5.2.1 *Fear*

Many VR example applications take advantage of the players' immersion to induce fear, whether jump scares (ID59), collision fear (i.e., objects coming at you, as in ID209), or motion-based fear (as experienced on a rollercoaster, e.g., ID120, ID232). As a result of this, in 94 of the clips fear caused the fail.

Clip ID55, for example, shows someone screaming as a result of a jump scare in virtual reality. Their loud and alarming reaction causes two spectating children to start crying (Figure 4c). In clip ID99, we see a player in an underwater shark cage, experiencing fear as the shark suddenly approaches. In many introductory VR experiences (i.e., experiences that are not intended to be complete games, rather short experiences to demonstrate the capability of VR), such as shark cages, haunted houses, and un-real rollercoasters, fear is targeted as a quick win within an immersive experience.

Fear causes a complete range of VR fails; including Excessive Reactions (ID69, ID74, ID138), Covering (ID25), Falling Over (ID199), and even Colliding (ID9, where the player attempts to run away from their fear, straight into a wall).

3.5.2.2 *Sensori-motor Mismatch*

Sixty-one clips result in fails caused by sensori-motor mismatch, where sensory (primarily visual and proprioceptive) cues are temporally decoupled from movement. This predominantly results in *Falling Over* and *Excessive Reactions*. Across these clips, we suggest there are two dominant forms of mismatch; (a) latency in the headset, and (b) visual-only feedback. These two forms of mismatch are challenging to precisely identify, and require our interpretation as VR experience designers.

A set of clips demonstrate players beginning to lean, stumbling to recover, and then falling over (ID44 and ID210). We would expect the players to take a step as they begin to lean, and so counteract their shifting weight. However, we instead see the players attempting to take this step too late. We suggest this is likely a culmination of latency in both the headset and our sensorimotor (especially proprioceptive) loop.

Another set of clips demonstrate players participating in fast-motion experiences, such as rollercoasters (ID157, ID188, ID224), where the visual-only feedback leads to balance issues (through a lack of vestibular cues, as opposed to proprioceptive cues).

3.5.2.3 *Obstacles in the Real World*

Obstacles in the Real World are the primary causes of *Colliding* and *Hitting* fails. Important here, these real world objects do not have a virtual counterpart. We see players make contact with televisions, tables, walls, mattresses, toys, and stairs, for example. In clip ID230, the player hits a lamp above their head (Figure 5c). This is notable as VR devices require the user to specify their play area prior to starting a game, however, no game requires the player to specify the ceiling boundary.

3.5.2.4 *Crowd Participation*

In 13 clips we see fails as a result of crowd participation. On the one hand, these fails stem from spectators interfering with the player. For example, in clip ID73, we see a spectator tickle the VR player, causing the player to turn around and hit them (Figure 5a). In clip ID45, a spectator attempts to interrupt a player, causing them to scream.

On the other hand, we see spectators being the cause of the fail. For example, in clip ID189, the spectator attempts to retreat and hide from the player in the corner of the room, to prevent getting hit accidentally. As the player nears, their concern grows.

3.5.2.5 *False Signifiers*

False Signifiers cause 4% of the fails. These are opportunities and actions within the game, that cannot fully support the associated real-world action. We have previously discussed climbing, snooker, and engine bay examples of this (3.5.1.3). ID22 shows another such example, where a child attempts to skydive from a plane. They stand gazing out of the door in VR, before throwing themselves out of the door, as they would from a real plane. They quickly exit the camera's view, as they collide with the solid floor.

3.5.2.6 *Setup Failure*

In eight clips, we see people fail as a result of setup failure. Primarily, this stems from players accidentally throwing the controllers (ID65, ID173, ID153), because they are not sufficiently strapped to their hands. Across all clips, we also see the emergence of a range of additional physical simulators to support the VR experience (such as treadmills - ID113, spinning chairs - ID124, and rodeo-style horses - ID116). We also see players breaking or falling out of these simulators (ID116 and ID3).

3.5.2.7 *No Cause*

There is a subset of 14 videos that contain no obvious cause of fail. The majority of these clips relate to the *Excessive Reaction* type, specifically those where the players' actions are odd from a social-spectator perspective. The only possible cause here, for a subset, could be 'enthusiastic participation'.

3.5.3 *Spectator Interaction*

Our corpus showcases a wide range of how spectators engage with players. We define spectators as nearby onlookers who actively follow the in-game experience, but are not wearing the VR headset. We typically see spectators crowded in a group of two or more, such as at home parties or in shopping malls, gathered close and focused solely on the player. In some clips we see multiple people sat around on sofas watching, in others a crowd stands nearby watching over the player's shoulder. We also see clips with only a single spectator filming. In some videos, the observers simply walk past the shared environment and only become spectators at the moment of failure, which is often indicated by a scream from the player, or the loud noise of a collision. In response to these fails, we see an overwhelming amount of laughter and mockery from spectators. These reactions are consistent with Harper and Mentis' work [141] on the social organisation of Kinect play, which emphasise accounts of "fostering [of] laughter" and "ridicule and mockery". Even when the fails result in unfortunate accidents in the home (i.e., to furniture or other persons), the spectators and players are still able to jointly laugh and find shared enjoyment in the absurdity of the VR experience. In more serious accidents, spectators are also seen running to the player to check if "[they] are okay," expressing a grave amount of concern as they help the player recover from the fails.

In 53% of videos, the VR view is made available to spectators through an additional display, such as the living room TV, a monitor or projections on the wall. This enables the spectators to share in the virtual experience and facilitates impromptu participation (ID23, ID190, ID208). We see a range of interaction between the player and spectators before the moment of failure. These interactions include (a) providing gameplay instructions to the player (ID6, ID16), (b) providing contextual cues about the real environment (ID26), (c) general chatter between the player and the spectator(s) (ID11, ID59) , and (d) physical support (such as holding the players to prevent them falling over - ID24, ID156).

In observing spectator reactions to the moment of failure, we see three prominent reactions: laughing and screaming, active help and support, and

Figure 7: Illustrations of key VR fails clips. In comic *a)* we show a video (ID130) in which a VR player nears a spectator. As the player gets closer and closer, the spectator backs into a corner and holds his arms close to their body to avoid being hit. Comic strip *b)* shows a clip (ID131) that captures an example of co-play between a player and a spectator. Here, the user plays a VR skiing game and the spectator is holding their hands, taking control of the player's movements. During this, the player is scared by a cliff that appears in the virtual environment and jumps up and screams, leaning on the spectator for support and making them laugh.

expressing concern. The spectator reactions are not mutually exclusive to each other.

3.5.3.1 *Laughing and screaming*

The most common form of spectator reaction we see is laughing or screaming (59%). We see laughing and screaming occur most with *Excessive Reaction* fails, followed by *Falling over*. These reactions range anywhere from slightly chuckling (ID141) to full hysterics (ID7, ID40, ID95). In clip ID53, a spectator closeby waves his hand as if to say 'hello' to the approaching VR player. Unaware of the spectator's whereabouts, the player slaps the spectator in the face. The oddness of the player's behaviour, and perhaps the contrasting response to the 'hello', results in fits of laughter from the group of spectators, including the spectator who has just been slapped.

We also see spectators screaming and experiencing fear themselves by watching the VR player. In clip ID136, the player's scream in response to losing her balance on a rollercoaster game caused an infant to cry in the background.

3.5.4 *Expressing empathy and concern*

Spectators also expressed empathy and concern in 12% of the videos. This is most frequently seen in the *Falling over* category. We often see *Active help and support* co-occurring with empathy and concern. In clip ID41, the VR player loses balance on his chair and falls backwards, hitting the staircase behind him before falling to the ground. The filming spectator remarks in shock "Oh...no!", while another spectator rushes to help the player up, and we can see others in the background covering their mouth with a concerned expression.

Interestingly, there are some instances where the spectator expresses concern towards the headset or furniture, and *not* the VR player. In clip ID12, the spectator remarks and expresses frustration that the player collides with a painting that is *"really worth some money"* (Figure 5d). Meanwhile, in clip ID215, the spectator runs towards the headset that was dropped on the floor.

In these videos, the spectators' concerns revolve around damages to objects in the real world rather than the player.

3.5.4.1 *Active help and support*

In many videos where we see an expression of sympathy and concern from spectators, this concern often leads to an offer of help and support to the player during or after a fail (29%). We suggest that this occurs when spectators perceive that the fail has caused damage (to the player or the environment). We see this most prevalent with *Excessive Reaction* fails, followed by *Falling over*. In many of the excessive reaction fails where the player physically flails (ID94, ID157 (Figure 3a), and ID166), the spectator provides support by holding the VR player to provide balance (ID18, ID90, ID163), or attempts to *catch* the player before they fall over (ID81, ID93). When the VR player falls over, spectators are often seen rushing towards the player and checking if they have hurt themselves. We see this in clip ID29 when the spectator drops the camera to run to the player who has fallen over to offer help. In some instances, the spectator facilitates the exit on behalf of the player, as seen in clip ID104. Here, a spectator nearby attempts to approach the player to take the headset off (despite not being asked to) upon hearing screams from the player and interpreting that as a call for help (Figure 3b).

3.6 DISCUSSION

The VR fails clips reveal interesting insights into the use of virtual reality for HCI. First, VR as a technology is designed to be immersive and support full-body, lifelike (or 'natural') interaction. As a result of this, the clips reveal players treating the environment as a fully-supported real-world space. For example, we see players seeking to walk long distances (ID62 , ID218), use un-tracked body-parts (such as kicking whilst boxing, ID89), and lean on virtual-only objects (ID222). These actions all fall outside of the design of the game, but yet could be considered 'expected' behaviours. These fails could be interpreted as failures in game design and, from a breakdowns perspective, perhaps game designers need to go further to differentiate between supported and virtual-only objects and stimuli. From an alternate perspective, however, perhaps these fails are part of developing an understanding and expertise in VR, and in turn facilitate some of the joy and entertainment that we witness.

Second, previous work has hinted at the existence of a social contract associated with VR-play [202]. Placing yourself in VR requires a certain level of trust in the people and spaces around you, as you are sensorily decoupled from the space in which you play. This social contract plays out in numerous different ways across the VR fails clips. On the one hand, we see clips of spectators attempting to move out of the way of the player, to not impede their play (ID130). This is a clear instance of prioritising the player. On the other hand, we see the player physically striking spectators (ID171), being reprimanded for risking damaging their environment (ID12), and being interrupted for being too loud (ID45).

These are all instances of play needing to fit within the lived environment, potentially at the expense of the play experience. While in a lab setting, VR participants are 'protected' and 'prioritised' by the researchers-as-spectators, this relationship perhaps has greater equity outside of the lab.

Next, the clips reveal different ways in which fails are experienced. Some fails, such as collisions and hitting, are experienced and/or enjoyed by both the player and the spectators. Other clips, however, such as excessive reactions (whether physical or emotional), are experienced as fails only by the spectators. To the player, these reactions may form a natural part of the interaction or experience. Following Reeves et al. [78], we also see clear instances of players performing for spectators, purposefully overacting their experience for the joy of those around them. We believe that novelty bias is one driver for this type of fail, notably where players participate in scripted VR experiences (e.g., roller coasters, shark cages, etc.). However, for the completeness of our analysis, it is important to understand these fails as a contributor to VR Fails, and the design implications in Section 5 are not directly derived from the novelty fails (i.e., flailing, screaming, etc.). Further, if we view these interactions as breakdowns, we could solve them by revealing more of the player's environment and experience to the spectators. This would serve to create a better-shared understanding between the immersed users and the spectators. However, these unexpected reactions play a role in creating part of the excitement and enjoyment for the spectator, at the seams of the player and spectator experience, and could be further emphasised.

Most VR experiences require the player to specify a boundary grid prior to starting a game. This acts as a virtual wall, revealing itself when you are a short distance from it and preventing you from colliding with the real environment. However, our analysis of VR fails revealed that play spaces are rarely as neat, static, or well-specified as headset manufacturers and researchers design for. For example, people and pets wander in and out, and play occurs in the middle of non-uniform cluttered environments. This introduces many opportunities for collisions and hitting, as the fails clips show. Further, many games require fast movements for interaction (such as swinging lightsabers in *BeatSaber*), that interact poorly with this visual boundary grid metaphor, leaving the player little time to respond and alter their actions. VR experiences reveal many opportunities that the play space cannot fully support, e.g., allowing players to swing a bat when near a boundary and revealing non-player-characters for interaction who are outside the play space. So, while it may be reasonable to require a clear play space for VR, game designers need a more dynamic approach to space configurations, taking on some of the area awareness burdens that are currently placed on players.

Across the fails clips we see spectators observing and interacting in different ways. They offer encouragement and advice, observe quietly, sing along to the background music, and provide physical support. A subset of this interaction directly echoes the sentiment of work by Cheng and colleagues, such as the physical haptics provided by spectators in HapticTurk [151]. In one instance, we see a spectator puppeting the player to perform the necessary skiing actions (ID88). Further, in the out-of-home clips, we see instances of other approaches to additional haptic input, such as tickling players' legs with a feather duster (ID115). In this particular clip, the player is mounted into a spinning cage-like seat, so it is difficult to imagine that the feather duster contributes much to the overall experience.

Many of the spectator interactions, however, do not include physical participation. Instead, the spectators act as an additional expert 'eye', offering advice and instructions to the players. It is rare, however, to see instructions being interpreted correctly, and this is a frequent cause of fails. For example, we see one player throwing their VR controller in response to hearing the spectator saying *"throw the racket"*, or another player attempting to perform the spectator's instruction of *"just jump"* and end up colliding with the wall (ID144 - Figure 5b). To encourage and enhance participation, designers could provide a set of language specific to the game, to create a purposeful inclusion dynamic for spectators to provide information from the real world to players. This is currently performed well in the multi-player, co-located VR game *Keep Talking and Nobody Explodes*[6].

3.7 DESIGN IMPLICATIONS

Having examined the diverse range of fail experiences, we depart from our findings to uncover new design opportunities for VR experiences. We group these design opportunities into categories: (1) Preventing Collisions, (2) New Interactions, and (3) Spectator Engagement.

3.7.1 *Preventing Collisions*

If we consider VR fails as *breakdowns*, then we should explore opportunities to address or prevent them, especially if the implications are more serious, such as with collisions (e.g., ID144) and hitting (ID220).

Changing the Play Space

The fails, especially those caused by Obstacles in the Real World, highlight a need to explore more nuanced approaches to play areas and virtual boundaries (or 'guardians'). On the one hand, this could be as simple as allowing users to specify more complex play area shapes, for example by enabling players to specify boundaries around fixed obstacles within the space (such as with the co-worker sat at their desk within the play area - ID172), and around the ceiling and overhead obstacles (e.g., in ID230). This would serve to reduce collision and hitting fails. Additionally, by reimagining the boundary grid, we may both reduce the chance of fails and present new gameplay opportunities.

Currently, in many commercial headsets (such as the Oculus Quest), as the player passes through the boundary grid, the front-facing cameras turn on, revealing the real world. We suggest expanding this technique as a gameplay opportunity: **a mixed-reality boundary play space** (Figure 8).

Departing from clip ID220: *The boxing game loads, revealing a gloved-up opponent. The player starts jumping from side to side, and throwing punches. As the player approaches the boundary, a virtual grid marks the edge of the virtual play space. Passing through this boundary, the player can 'see-through' the headset to the real world. This 'mixed-reality' space, however, is not necessarily beyond the game - the boxing opponent could follow the player through into this space. This space could also support novel game 'abilities', such as enabling communication with other players, map views, health regeneration, etc. Depending on the experience, this area could also enable the player to 'step out' of the virtual space to interact with spectators. Looking back, the player can still see into the virtual world, and simply steps back in to re-enter the ring and continue play.*

---

[6] Keep Talking and Nobody Explodes, Steel Crate Games, 2015.

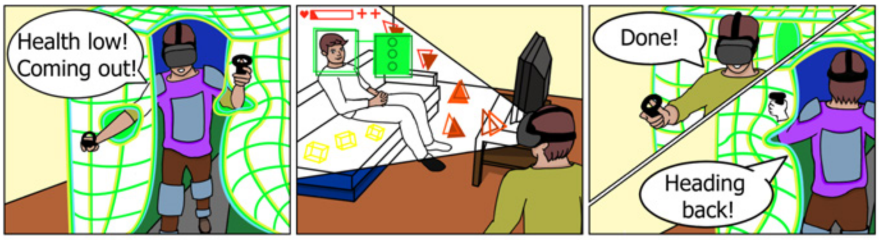

Figure 8: Illustration showing the mixed-reality boundary play space, where the space outside the boundary grid can be used for novel interactions. In this instance, without removing the headset, the player steps out of the virtual game world to regenerate health and interact with the spectator, before stepping back in to the game.

3.7.1.1 *Changing the Game*

A number of clips show players colliding with walls (e.g., ID144), televisions (ID13), and picture frames (ID12). Current VR experiences place the burden of spatial positioning on the player; the player needs to keep track of play boundaries, repositioning themselves as they approach walls, etc. While we previously considered changing the boundary behaviour, there also exist opportunities to change the mechanics of these games to **account for player positioning and context** (Figure 9).

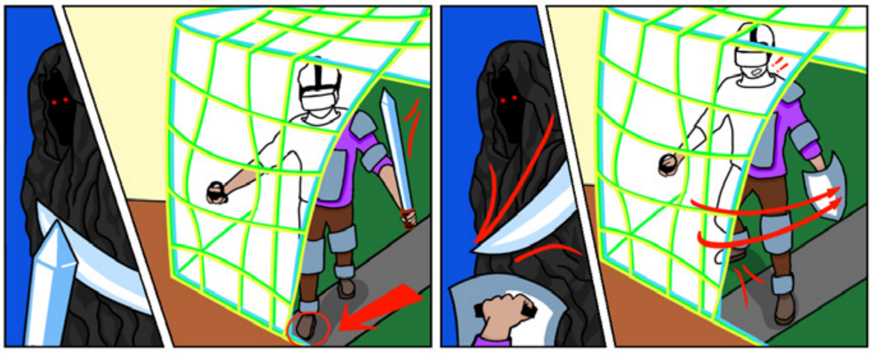

Figure 9: Illustration of changing game elements in order to prevent the player from stepping beyond the boundary grid of their play area. Here, the game switches the player's sword for a shield, to prevent them swinging their arm beyond their play area.

In clip ID11, the player stands in the middle of their play space, sword in hand. During play, the user walks forward, nearing the boundary of their space. *To prevent large swinging actions and potential collisions, the game changes their sword to a shield, constraining the motion the player is likely to perform. At all times, pre-empting larger movements the player might make, the game highlights objects and enemies that are available for interaction vs. approaching but not yet available.*

3.7.2 *New Interactions*

Alongside considering VR fails as breakdowns to be avoided, we can also consider them a constructive lens through which to reveal new interaction opportunities. Across the corpus, we see players performing actions that arise instinctively, for example curling into a ball (ID103), kicking (ID77), covering their face (ID211) and so on. We can incorporate these actions into controls themselves, for example **enabling 'peek-through' or 'scaredy-cat' mode** (Figure 10). These controls can further be context-dependent. In a horror VR game, the screen can go dark once the user covers their face in order to make the fear stimuli disappear. For an action game where users might cover their face because of incoming objects, the game could provide them with a shield once it detects covering.

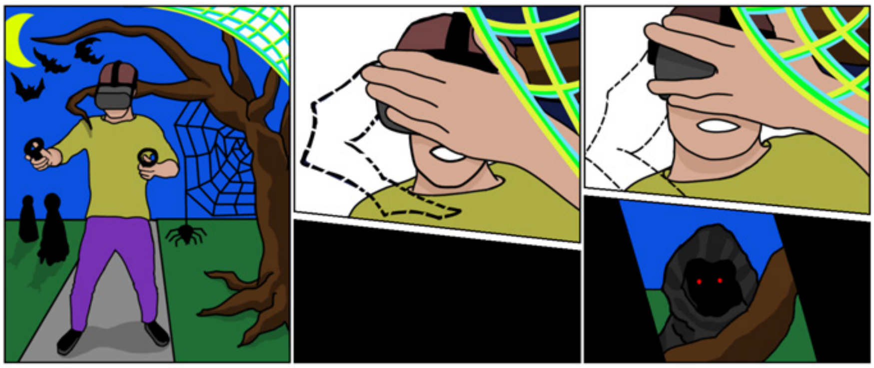

Figure 10: VR could offer new headset interaction opportunities. For example, this illustrations shows a peek-through mode, in which the player covers the headset to disable the display, and parting their fingers reveals thin slithers of view.

The player walks through a haunted house, full of dark shadows and eerie music. Scared, the player hides their head behind their hands. The screen goes dark. The player slowly parts their fingers, enabling peek-through mode - revealing only small slithers of the virtual view.

3.7.3 *Spectator Engagement*

3.7.3.1 *Increasing Spectator Awareness*

VR fails clips reveal a range of techniques for sharing visual insights into the virtual world with the spectator. For example, we see a variety of co-located screens, whether desktop displays (e.g., ID223, ID224), televisions (ID9) or wall-scale projections (ID81).

While this supports a level of understanding of the experience, it does not appear to support spatial understanding of the play area. As a result of this, it can be hard for the spectator to predict the space requirements or likely actions of the player. In ID20, for instance, we see the spectator getting hit by the player though they are simultaneously watching the game view on television.

Current work on shared VR experiences [e.g., 252, 240, 238] use projectors to share a spatially co-located view. These explorations examine only modestly dynamic settings, where revealing objects and environment details enable the spectator to understand the space. As the applications become more dynamic and the view angle changes more frequently, these projections will become increasingly hard to interpret and render. In these scenarios, these projectors could be re-purposed to **project simple visual movement predictions** (Figure 11).

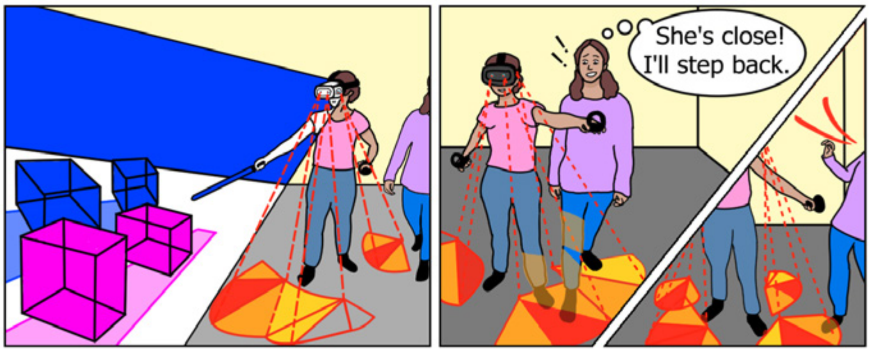

Figure 11: Illustration demonstrating a head-mounted projector, to notify onlookers about the likelihood of the player's movements. This abstract display supports spectator participation by enabling them to move furniture and people to prevent collisions.

The player loads their favorite song in Beatsaber - a game requiring the player to slash flying cubes with a lightsaber in time with the music. Their head-mounted pico-projector projects red polygons onto the ground in front, showing spectators where the player is likely to swing their arms next. Seeing their legs light up red, one of the spectators quickly moves out of the area.

3.7.3.2 *Increasing Spectator Participation*

Our analysis reveals that the most prevalent spectator reaction in fail videos are laughter (e.g., ID172) and screaming (ID136). More recently, games such as that Acron: Attack of the Squirrels![7] and Ruckus Ridge VR Party[8] enable collocated spectator engagement through asymmetrical gaming, whereby spectators can join the gaming experience on non-HMD devices such as mobile phones and PC. Given the prevalence of spectating VR, we expect many more such participatory experiences to emerge.

One such opportunity could be to enable spectators to simply impact the game using non-digital modalities, such as through **spectator power-ups**.

The player is in the midst of a fighting game and is becoming quickly overwhelmed. The game recognises specific forms of vocal encouragement from the spectators, and gives the player more 'health' points. As encouragement, the spectators could make specific hand gestures towards the immersed user that are then displayed in the game. Additionally, they could interact with objects found in the mixed-reality space which they can see through the player's view of the game on a separate screen.

---

[7] Acron: Attack of the Squirrels!, Resolution Games, 2019

[8] Ruckus Ridge VR Party, ForeignVR, 2016

3.8 DISCUSSION

VR is increasingly used in complex social and physical settings. We investigated how such environments influence VR use by analysing 233 YouTube videos of fails. We have identified typical failures and the reasons they occur. Further, we have identified design opportunities from the fails. Next, we discuss the main findings as well as concerns about validity when working with video as data.

3.8.1 *VR Outside the Lab*

Although we have focused on breakdowns and seams, it is worth reflecting on the videos as a source of data on VR outside the lab. The videos showcased a variety of positive, engaged experiences. Immersion appears to work for users, even in very social settings and cluttered spaces, very different from most VR laboratories. The videos also highlighted the important role of spectators, showing how varied social contracts about spectating play out in VR in front of an audience. In that way, our analysis is more detailed than typical crowdsourced VR studies [e.g., 190, 182] and more large-scale than typical studies of spectator experiences [e.g., 78].

The number and popularity of VR fails videos may be due to the novelty of immersive technology. This will likely change as the technology matures and more people play VR games. Currently, however, videos of VR fails seems to capture both first-time and returning users, alone and together. Fails themselves (beyond just VR fails) are a popular video type with a rich history, emerging from television shows such as *Jackass* and *Takeshi's Castle*, and then broadening to unscripted self-captured moments with the emergence of personal video capture devices and smartphones. This genre, of course, does not cover all VR experiences outside of the lab. Nevertheless, the clips in general rarely seem acted or scripted; some fails, then, should be considered an opportunity for VR designers and developers to hold on to, promote, and design for.

3.8.2 *Avoiding Breakdowns*

We see a range of fails within the VR fails clips. Some of these fails are dangerous, causing harm to spectators (ID220) and players (ID38) alike. As designers, it is our responsibility to prevent people from hurting themselves or others, or unwittingly breaking things within their environment. We should look to design opportunities that prevent some of these dangers. Motivated by the fails that we see, we have presented some initial ideas, such as reimagining the VR boundary grid, as techniques for preventing further collision and hitting type fails.

The specific types of breakdowns that we identified are useful to researchers and designers because they separate some key mechanisms, for instance, between the categories of sensorimotor mismatch and false signifiers. This extends earlier categories of broad breakdowns [53].

3.8.3 *Designing from Fails*

Many of the fails that we analyse result in a shared joyful experience: we see players and spectators laughing together; people deciding to capture these moments on video and sharing them with a wider audience; and these shared video clips receiving millions of views. Consequently, we should not look to replace or fix the aspects of VR design that lead to these fails (i.e., jump scares, vertigo-experiences, fast motion embodied play, spectator engagement and involvement), as they appear to be a central tenet of VR play and its intersection with the real world. Instead, we should design to further promote these features of VR, and some research is actively doing so [e.g., 151, 160]. Following this approach, we further present design ideas, such as providing contextual motion predictions to spectators in order to prevent them becoming unwitting obstacles.

Video gaming itself is often a shared experience, where player and spectator engagement drive and promote enjoyment. VR fails reveal similar traits in VR gaming, albeit with different interaction dynamics. VR, by design, is a private experience within a world hidden in a headset. In turn, it reveals little of the visual cues to the spectator. While some clips feature additional displays for spectators, we see onlookers finding other ways to participate and enjoy the shared gaming experience (e.g., singing along to the soundtrack, physically interacting with the player, or simply enjoying their reactions). The secret elements of the VR players' experience (their visual cues, and spatial understanding and affordances), reveal new kinds of participation and enjoyment that should be harnessed and exploited further. Anecdotally, we have observed this type of enjoyment occur in some popular *BeatSaber* songs, which have been perfectly choreographed to make the player perform the famous associated dance for any spectator (for example, Gangnam Style by Psy). This is an example of Reeves et al.'s *expressive* spectator experience [78], where the player becomes a performer for the audience.

3.8.4 *Limitations*

We analyse VR fails to begin to understand how VR is played in-the-wild, and its interplays with spectators and the lived environment. We believe VR fails provides a good starting point for this kind of analysis. That said, however, the corpus is not representative of broader real-world use as it is specifically collated by the community to showcase clashes at the intersection of VR and the real world. As such, it excludes the mundane, everyday, private play that may yet constitute a large part of VR use. Future research should approach this from various perspectives and incorporate many views in order to accurately capture the breadth of experiences of consumers with virtual reality. For example, Twitch or other gameplay streaming platforms may provide another insightful source of data.

Further, analysing the clips is difficult. The home-video nature of the clips can introduce uncertainty to the coding. For example, some clips end very quickly after a fail (making it hard to determine if it resulted in exiting or how the spectators supported the player), some provide no insight into specifics of the VR experience (i.e., what game they are playing), and some have no single moment of failure. Thus, a part of our coding relied on our own expertise as VR designers and researchers.

The primary use case of virtual reality is currently gaming; most clips concerned gaming. While other application areas are emerging, these are not prevalent across the VR fails corpus. Future work should explore the breakdowns and useful seams in other application areas.

3.9 CONCLUSION

Empirically describing the use of VR remains an important research challenge. In particular, the social and physical factors that shape VR remain underexplored. We have used clips of breakdowns in VR as a source of data to understand those factors and how they may inform design. Through our findings, we propose a range of design ideas that aim towards involving spectators and the physical environment, in order to enhance the VR experience.

4

DEVELOPING FEEDFORWARD FOR VIRTUAL REALITY

> *To be an observer exclusively in VR is to be a phantom, a subordinate ghost who cannot even haunt.*
>
> *Dawn of the new everything: Encounters with reality and virtual reality* — Jaron Lanier [188]

This chapter is based on and reproduced from Feedforward, published as *Using Feedforward to Reveal Interaction Possibilities in VR* [305]: *Andreea Muresan, Jess McIntosh, and Kasper Hornbæk. "Using Feedforward to Reveal Interaction Possibilities in Virtual Reality." In: ACM Trans. Comput.-Hum. Interact. (2023). issn: 1073-0516. doi: 10.1145/3603623. url: https://doi.org/10.1145/3603623.*

This work encapsulates the effort towards introducing the notion of feedforward (showing people what to do and how to do it) in **VR!**. This work highlights the path from concept to artifact through a design space and the feedback loops of prototyping and interaction modeling.

4.1 ABSTRACT

In virtual reality, interactions may fail when users encounter new, unknown, or unexpected objects. We propose using feedforward in VR to help users interact with objects by revealing how such objects work. Feedforward lets users know what to do and how to do it by showing the available actions and outcomes before an interaction. In this paper, we first chart the design space of feedforward in VR and illustrate how to design feedforward for specific VR interactions. We discuss starting the feedforward, previewing actions and outcomes, and returning the virtual world to its state before the feedforward. Second, we implement three real-world VR applications to show how feedforward can be applied to multistep interactions, perceived interactivity, and discoverability. Third, we conduct an evaluation of the design space with 14 VR experts to understand its usefulness. Finally, we summarize the findings of our work on VR feedforward in 15 guidelines.

4.2 INTRODUCTION

A fundamental aspect of interacting in virtual reality (VR) is knowing how to manipulate objects, what interactions are available, and where to navigate. Designers often help users understand this by relying on real-world knowledge. Virtual objects can mimic real-world objects in appearance and functionality. For example, if the user sees a virtual door, they may turn its handle to open the door. If the user sees a virtual button, they may press it to perform a certain action.

Despite its convenience, mimicking real-world interaction brings new challenges to VR. For instance, users may expect that they can act in the same way in VR as in the real world. However, the point of some VR experiences is to provide excitement by going beyond mundane life. Interactions can be purposely different or novel to users.

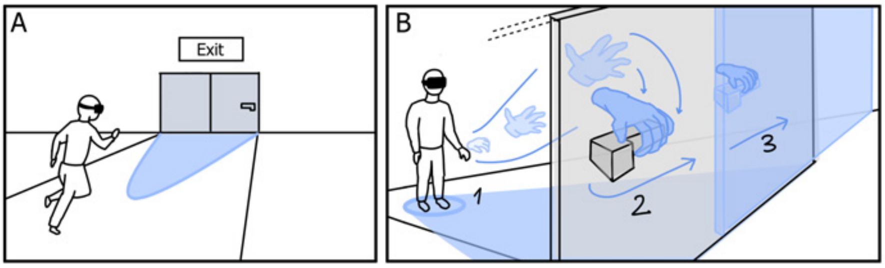

Figure 12: Feedforward helps users understand how to act in virtual reality (VR). The user moves towards the door and, upon proximity, triggers a feedforward in Panel B, Moment 1. Then, the user sees a preview of the actions and outcomes involved in opening the door (Panel B, Moments 1, 2, 3). The preview is made up of a ghosted hand that moves a ghosted copy of the door, thus showing the user how to interact with it. By Moment 3, the feedforward has revealed to the user how to open the door.

Such VR interactions add value by contributing to immersion, presence, and fulfilling experiences. To do so, they need to differ from their real-world counterparts. Thus, mimicking the real world is not always a viable option. In such cases, designers cannot meet some users' expectations of interactivity.

We propose using *feedforward* to show users what they can do in VR and how to do it. Unlike feedback, feedforward shows actions and their results before the user performs them. A canonical example of feedforward in 2D interfaces is the preview of swiping the iPhone lock screen to unlock the phone. Vermeulen et al. discussed feedforward at length for 2D interfaces and as a general concept, and gave examples of real-life use [148]. Although feedforward has been developed and tested successfully for 2D environments, many questions remain about using feedforward within VR. Why would feedforward be useful for VR?

On the one hand, feedforward offers similar benefits as feedforward for non-VR user interfaces [e.g., 148]. This includes bridging Norman's gulf of execution and improving discoverability, helping users plan and execute actions [144]. On the other hand, VR differs from traditional user interfaces. In virtual environments, any object can have many manipulations available, and others may be available for the rest of the environment. Given that examples of feedforward mainly contain one action (e.g., the iPhone's "slide to unlock" function), how can we cater to the higher complexity of VR? Furthermore, actions in VR may vary in granularity, from a single movement to a long sequence of actions. How does feedforward apply to such situations?

We answer these questions by defining feedforward for VR and detailing its design space. Throughout the paper, we focus on physics-based interactions using hand-tracking to exemplify the feedforward design space. We develop the design space as an idea-generation tool for designing and describing VR feedforward aimed at researchers and designers. For example, a feedforward may show users a preview of opening a door with a misleading opening mechanism (see Figure 12). We can design this feedforward by showing users a ghosted version of their avatar performing the movements required to open the door. Ultimately, the feedforward design space may

Figure 13: This figure illustrates the three key stages of feedforward in VR. The user starts in some initial state in the virtual environment (VE). Then, the user executes a (1) trigger, in this case, by turning around and gazing at the virtual door. That leads to (2) previewing the actions involved in opening the door (operating the slider) and their outcomes (opening the door). After the preview is shown once, the feedforward (3) exits from the preview. Finally, the door is returned to its initial state. This hopefully leaves the user better informed about what to do and how to do it.

aid practitioners in solving challenges related to discoverability and understanding of VR interactions.

While VR can easily mimic 2D feedforward, the design space goes beyond traditional feedforward to include VR-specific dimensions. We explore different ways of triggering feedforward in VR, previewing actions and outcomes using virtual objects and avatars, and returning the world to its initial state. We test the feasibility of VR feedforward by implementing a VR feedforward system that enables a subset of the design space. We then create three demos consisting of feedforward variations for three real-world application scenarios using the system. We evaluate the feedforward design space with 14 VR experts focusing on both theoretical and practical aspects of feedforward. Finally, we summarize our observations in 15 general guidelines for applying feedforward in VR.

Morphological analysis results in a multidimensional matrix that consists of problem-defining qualitative or quantitative parameters. This matrix defines the Zwicky box. While initially used in engineering and astrophysics to identify rocket propulsion systems, the Zwicky box has evolved into a useful problem-solving tool for many fields. Researchers have also used this methodology to arrive at design spaces in HCI [196], for example, for input devices [26], and in VR for gaze interaction [220], 3D sketching [237], and, more recently, a design space for worlds-in-miniature [261].

4.3 OVERVIEW OF FEEDFORWARD IN VR

We introduce feedforward to VR to help designers generate ideas for showing users how to act. Feedforward shows users a preview of what they can do in a VR environment. In such an environment, the user needs to (1) *trigger* the preview, for instance, by being in a particular place or looking at a particular object. When triggered, the feedforward shows the user a (2) *preview of the actions and their outcomes* available in a particular place for a particular object. For instance, the preview could show how to grab the wheel of a car and steer it. The user may then (3) *exit the preview* to perform the actions themselves. A preview may be shown only once or in a loop, or the user may interrupt a preview if they wish. One of the main goals of feedforward is to reveal information when the user needs it. Taken together, *triggering*, *previewing*, and *exiting* constitute the three stages of feedforward in VR.

Figure 13 shows an example of applying the stages of feedforward to a VR interaction. Here, the user triggers a preview by gazing at a door from a distance. In the preview, a ghosted hand shows the user how to operate a nearby slider and open the door. Here, moving the slider represents the action, and opening the door is the outcome. The feedforward continues to the third stage, exiting, after which the door slides back to its initial position, and the ghost hand disappears. The states of the environment before and after feedforward are the same because changes during feedforward do not persist. However, the user now has more knowledge about the interaction with the door.

Feedforward in VR is intended as a bridging concept [152], which is more general than a particular design but less general than a theory. Bridging concepts describe theory and practical applications that may reveal potential design opportunities and novel theories. For the theoretical part, we develop a design space for feedforward in VR and a working model for feedforward interactions containing the three stages. For the practical side, we implement a feedforward system and evaluate it with experts to uncover new design ideas and reshape the theory.

4.4 BACKGROUND

Both VR and non-VR applications rely on a large body of work within HCI that informs design and helps maintain usability standards. Here, interactions play a key role in the communication between the user and the system. Norman has argued that the goal of interaction design "is to enhance people's understanding of what can be done, what is happening, and what has just occurred" [144]. Thus, when the technology aims to benefit the user, it should explain itself and its purpose to the user. In this section, we present various human-centered strategies to either reveal purpose or embed purpose into design.

Hornbæk and Oulasvirta [187] considered the types of interaction prominent in HCI, defining them as a "mutual determination" between humans and computers. They identified seven ways to conceptualize interaction: dialogue, transmission, tool use, optimal behavior, embodiment, experience, and control. The authors emphasized the teleological nature of interaction, in which the users navigate technology as a means to an end. The embodiment dimension of interaction highlights this view of intention regarding action and context. For example, moving a cursor toward a target captures the intent of navigating to a specific web page to acquire specific information.

A more general model of system use is Norman's stages of interaction [144]. The model covers the steps of interaction, beginning with the user's goal and covering all actions that lead to accomplishing that goal. The stages can be grouped according to whether they concern execution or evaluation. Execution refers to planning, specifying, and performing the actions required to achieve this goal. Evaluation refers to observing the outcomes of these actions, interpreting them, and comparing them to the desired outcomes. The stages-of-action model offers a simplified but useful look at the chain of actions within an interaction.

Once users engage with the system, they might expect some confirmation of their engagement. Feedback is a fundamental HCI concept that describes such confirmations. Essentially, feedback communicates to users the outcome of an action. This, in turn, helps users understand the state of the system after they engage with it and helps them inform future interactions.

The ability to provide users with "understanding" and "discoverability" represents one way in which systems reveal their purpose [144]. Norman gives a door with a misleading design as an example of poor discoverability in the real world. This design communicates to users that pushing the doors will open them. However, the doors do not work as expected, and users often struggle to find the proper way to open them. *Discoverability* refers to how a system can reveal its available actions to users and how to perform them [144]. *Understanding* reveals to users the meaning of the system's controls and actions [144]. Together, these are the two most essential qualities of good design in Norman's view. However, design itself cannot always provide an answer to complexity. The author also suggests using manuals or instructions for complex devices. VR regularly includes such complexity.

A related notion is suggested interactivity, which concerns how users discover which elements of the UI are interactive [170]. Boy et al. investigated what makes people engage with information visualizations on the web [170]. They also presented a design space for suggested interactivity consisting of 45 cues split into animation, attractor, trigger, visual attributes, and persistence, as well as intended interaction and feedforward. In a study, the researchers evaluated three of the designs of suggested interaction. A particular cue with feedforward was the most successful in attracting users to interact with a figure. The authors hypothesized that feedforward helped users to understand the potential benefits of engaging with the figure. From a theoretical perspective, Sundar et al. presented a framework for interactivity that covers source, medium, and message [120]. They conclude that as society becomes increasingly involved with interactivity, understanding its consequences is crucial. Feedforward has proven to be useful in revealing the outcomes of interactivity.

Apart from explicit attention-getting cues, researchers can use people's knowledge of the real world to suggest interactions. Originating with J. J. Gibson [153], affordances capture relationships between agents and objects. Affordances refer to the information conveyed by objects about how to interact with them. Norman further categorized affordances as *perceived affordances*, capturing perceived action possibilities, and *real affordances*, which relate to the physical attributes of objects [47, 20]. Usage and purpose can also be nested within the design to implicitly signal interactions. Later, Norman suggested that the term *signifiers* would be a more appropriate formulation for perceived affordances [98].

For instance, Affordance++ is a useful concept that builds on affordance by suggesting interactions with real objects [164]. The authors argue about how, instead of changing objects, we can change bodies to communicate action possibilities. Affordance++ uses electrical muscle stimulation to give users suggestions about how they can act. The simulations are triggered when the user's hands are in certain positions relative to objects. This technique expands the classical notion of affordance by *communicating dynamic use* that can be seen during multistep actions and behaviors involving motion or dependent on time. The authors tested the technique for objects with poor natural affordances and found that it helps users interact with them.

These notions underlie our design of feedforward for virtual reality and are carefully considered during the development of the feedforward design space. We emphasize the teleological aspects of interaction by first breaking down interactions in VR in terms of Norman's interaction model and, further, in terms of actions and outcomes. Such a teleological approach allows us to focus on the user goal by designing a system to achieve this goal as a means to an end.

4.5 RELATED WORK

To use any interactive system, you need to know what to do and how to do it. Our paper draws on the body of HCI concepts aimed at helping users interact, such as feedforward, feedback, and affordance. Next, we discuss those concepts and how they have been used in VR.

4.5.1 *Feedforward*

While feedback is widely used and remains a key design principle of HCI, its twin concept of feedforward has not gained the same attention. Feedforward informs the user what must be done to achieve a certain outcome, whereas feedback helps the user understand what has occurred or is occurring. One key difference between the two is user action. In feedforward, the action is hypothetical or is to be executed, whereas, in feedback, the action has already happened or is in progress. Simply put, feedforward tells the user "*what the result of their action will be*" [148], and feedback tells the user *what the result of their action is*. Both of these techniques are closely related to setting user expectations and goals for an interaction [144].

Vermeulen et al. discussed feedforward at great length, both conceptually — by relating it to affordance — and practically — by describing 2D applications [148]. First, they described Wensveen's framework on coupling actions and reactions on *time*, *location*, *direction*, *dynamics*, *modality*, and *expression* [77]. Wensveen et al. also named *inherent feedforward*, which shows what action is available and how to execute it; *functional feedforward*, which gives information about certain features or some purpose; and *augmented feedforward*, which refers to messages, labels, or pictograms that supplement existing action possibilities. Vermeulen et al. remarked that feedforward can be expressed in other modalities besides the visual, such as the tactile. Second, borrowing from Gaver's notion of affordances [27], Vermeulen et al. came up with *hidden feedforward*, *false feedforward*, and *nested feedforward*. Feedforward can also be shown at certain points in time or updated continuously in what they call *static* or *sequential feedforward*.

An application that uses feedforward is Bau and Mackay's dynamic guide called *OctoPocus* [91]. It applies feedback and feedforward techniques to gesture learning with cursors. Here, feedforward is described as being a "cheat sheet" accompanying gestural commands. Bau and Mackay presented two dimensions of feedforward: the *level of detail* and the *update rate*. The *level of detail* captures how much of the gesture is displayed to the user. *Update rate* refers to the frequency of this display. Building on dynamic guides, *Shadowguides* uses Microsoft Surface to help users learn gestures using their prints or *shadows* left on the tabletop [106]. These shadows are presented through feedback and annotated through feedforward. As annotations, they use arrows, highlighted keyframes of print deformations, and dynamic markers, which are text labels that surround the prints only when relevant. Shadowguides emphasizes how to prevent occlusion by showing the user's handprints at a different location, while dynamic guides introduce the concept of *level of detail* to capture the sequence of actions within a feedforward. Both concepts are useful for our discussion of feedforward in VR, where many objects often overlap and the user's range of motion is increasingly complex.

Using feedforward to improve existing GUI techniques has a positive impact on user experience. For example, in *Fortunettes* GUI widgets may show their future states to users [215]. As a proof of concept, the authors implement these functionalities in a commercial airline application, adding feedforward to a weather radar's widgets. In a subsequent online study of these enhanced widgets, researchers compared their use in a variety of demo applications with traditional non-feedforward widgets. While participants had fewer clicks with enhanced widgets, the results showed that the completion time was longer. Some participants mentioned that the annotations of the feedforward took longer to process, but most agreed that these were helpful to perform quicker despite the results. In the study, the participants reported using trial-and-error strategies whenever feedforward was absent, with most agreeing that *Fortunettes* would be useful when faced with unfamiliar interfaces.

While researchers agree that feedforward fits emerging technologies, no attempt at introducing a design space for this technique to a specific field has been made. We address this by developing a design space for feedforward in VR. The design space mostly deals with *inherent feedforward* because we aim to show people what they can do in VR and how to do it. Some concepts just described are part of the design space, such as the *level of detail*. Other concepts are new, describing what it means to preview actions and outcomes within VR. Using feedforward in VR contributes to setting the user's expectations and goals for the interaction, complementing existing feedback techniques [144]. For possible applications in VR where interfaces are unfamiliar, feedforward has the potential to similarly increase perceived performance and aid in navigating the UI.

Of course, existing VR systems use some way of showing users what they can do. Next, we analyze such systems to determine how they differ from the concepts discussed so far.

4.5.2 *Showing What to Do in VR*

Some emerging interaction techniques within VR display feedforward-like qualities, especially in the field of enhancing motor control. Here, researchers show people what to do through guidance cues or other types of feedback. For example, ghosted hands are an additional pair of hands rendered differently from the user's own, commonly translucent. Researchers have often used ghosted hands when demonstrating to users how to move their tracked hands [271]. In the work of Lilija et al., the objective was to teach users how to move their hands to perform mid-air gestures. While this work focused on hand guidance with ghosted hands, there is also research on how to teach motor skills with objects. For instance, other researchers have used a ghosted brush to teach calligraphy [62]. Apart from ghosting, external renderings and visualizations on the body, such as arrows on the hands, can also guide movement [135].

Recently, Fennedy et al. introduced a VR version of OctoPocus for mid-air gestures that builds on Bau and Mackay's design of feedforward [267]. This implementation uses colored pathways and text labels to help users complete mid-air gestures. Some degree of ghosting is used by making less likely gesture pathways transparent. We have also seen variations of feedforward used in commercially available apps. For example, the game *Elixir* shows an additional floating pair of ghosted hands coupled with audio instructions to teach the user how to teleport using hand gestures. For VR exergames, Barathi et al. used the term "feedforward" as a training technique in which users see previous versions of their avatar performing [197]. The authors implemented a VR system wherein the participants see themselves cycling under different conditions to study performance and motivation during workouts. In this case, feedforward describes a form of playback. While feedforward is used across fields with different meanings, in this paper, we present a VR design space for the HCI concept of feedforward. The work of Barathi et al. is similar in that it also involves recording and playing back the avatar. The difference is that the VR feedforward system captures interactions and introduces levels of abstraction, such as ghosting. The purpose of the VR feedforward system is to show users what to do in VR and how to do it, not to motivate the user to continue their current action.

Dillman et al. presented a framework consisting of interaction cues after analyzing 49 video games, some of which are immersive [199]. The framework contained three dimensions, *purpose*, *markedness*, and *trigger*. The *purpose* dimension categorizes cues by how they are intended to help players. *Markedness* describes the visual characteristics of the cues (such as *subtle* or *emphasized*). The *trigger* dimension describes how these cues are revealed to users. While Dillman et al. showed how their framework can be used to generate visualizations for augmented reality, we suggest that some of these cues are also suitable for VR and warrant further research. For example, Hu et al. compared delayed and immediate interaction cues in their research and found that immediate cues are more efficient in VR applications [239]. In this paper, we adopt some of these notions within the feedforward design space and show how they can help users in VR.

Feedforward also relates to the emerging AR field of *situated visualization*, where the key concept is that interactions are "visualized in situ, where it is relevant to people [259]." In a review, the researchers identify five components of situatedness: space, time, activity, community, and place. Feedforward may instantiate situated visualizations because it connects outcomes to actions within the environment. Considering its physical and temporal aspects, feedforward may be particularly suited to augmented reality. Apart from embodiment, hand-tracking is a unique feature of immersive mediums. For this reason, we have chosen to instantiate feedforward with hand-tracking in practice.

Because of the situatedness of VR feedforward, proxemic interaction principles may be leveraged for feedforward in practice. Within ubiquitous computing, Ballendat et al. propose exploiting the knowledge between devices and people to inform the design of *proxemic* interactions [113]. Later, Marquardt and Greenberg further describe location, movement, orientation, and distance as dimensions that define proxemic interactions, and identify six challenges concerning seamless and embodied design [132]. Vogel and Balakrishnan provide an interaction framework to describe interactive public displays, from close—explicit to distant—implicit [76]. The researchers present *self-revealing help techniques* to help users interact, similar to the feedforward concept. In their implementation, when users do not perform actions in a particular state for some time, they receive a video that loops with the actions available for the display.

4.5.3 *Challenges for VR Interactivity*

Even with the interaction techniques discussed so far, fundamental challenges remain for VR. In particular, VR may share similar challenges with ubiquitous computing, such as revealing interaction possibilities and directing actions, as described by Marquardt and Greenberg who propose proxemic interactions as a way to solve such challenges [132]. This section briefly reviews challenges to VR interactivity.

In general, designers may draw on two types of knowledge when bringing users' expectations of interactivity from the real world. This type of knowledge is related to real-world objects and processes that convey the everyday know-how of how the world operates. First, this knowledge includes people's experience of their bodies interacting with the world and with others. The reality-based computing framework further exemplifies how to use this knowledge to build emerging interfaces [96]. For VR, an example of this is the virtual hand metaphor, which allows users to manipulate nearby objects from a first-person perspective, just like in their day-to-day lives [49].

Second, designers also build on people's experiences with technology. This means that we can find established ways to interact with the technology nested within VR. For example, VR applications can provide users with elements from traditional graphical user interfaces: virtual keyboards, cursors, pointers, scroll bars, labels, drop-down menus, and many other familiar interface elements. These elements are sometimes adjusted to the VR experience, resulting in novel interaction techniques such as the flexible pointer [66].

A type of knowledge specific to VR users is the awareness of the outside world while immersed in the virtual world. McGill et al. looked at usability challenges faced by immersive VR and found that interacting with the real world posed a significant challenge [166]. While bringing real-world objects in virtual environments could negatively impact presence, the same was not true for renderings of non-immersed users. The researchers suggested that this difference emerged from people's expectations. The participants expected other users in the virtual environment, whereas objects like keyboards were "unnatural" for that scene.

In VR, both real-world and GUI knowledge is largely conveyed to users through affordances, relying on their previous knowledge of the world [e.g., 193]. However, nesting real-world affordances within VR and mimicking classical UIs can prove problematic. Some of these downsides relate to VR hardware, such as the resolution and other graphics capabilities [253]. Other downsides concern the way users interact within VR, namely how embodied interaction may hide discoverability [175], similar to Norman's criticism of gestural interfaces [118]. While most common in AR, situated visualizations may overcome the ambiguity of affordances by making interactions explicit, tying usage to context [259]. However, guidelines on how to represent these visualizations lag behind, particularly concerning layout design, cognitive limitations, and complex interactions [264].

Moreover, VR does not provide the same haptic feedback as the real world. This can cause problems for users when they encounter *false signifiers*, which means that the virtual objects are treated as real objects. For example, users can lean on a virtual pool table and fall because they interpret the object as providing this type of interaction as in the real world. Furthermore, users may have different expectations of what they can see in the virtual scene based on the context with which they interact [166]. In VR, where the lines between the virtual and the real are especially blurred, setting expectations about interaction is essential.

4.5.4 *Summary*

A key challenge in designing human-computer interactions is to help users understand what they can do and how they can do it. Existing concepts for doing so (e.g., suggested interactivity and feedforward), have been explored mainly for 2D graphical user interfaces. We propose using feedforward to address these challenges and improve discoverability and understanding in VR systems.

We identify several applications for feedforward within VR whose common thread is that they cause confusion for the user:

1. Interacting with objects that have low perceived interactivity [164];
2. Performing interactions that require a specific sequence of actions, possibly involving multiple controls [259, 264];
3. Setting expectations for interactivity, e.g., for unfamiliar or different types of objects [166, 164, 262];
4. Discovering action possibilities within a VR [175, 118, 132].

In VR, the strategies used to help people act have different strengths and drawbacks. For possible applications in VR where interfaces are unfamiliar or work differently from the real world, feedforward can be the answer. This is because feedforward can increase perceived performance by helping users navigate the UI. Once the users understand the available interactions, they can be more likely to engage with them. While feedforward is a concept that informs design, its usefulness for VR has yet to be spelled out.

4.6 FEEDFORWARD IN THEORY

In this section, we define and illustrate the concept of feedforward in VR. VR allows many actions on a single object, and even more for the entire environment. Moreover, actions in VR may vary greatly in the level of detail, from a single movement to a long sequence of actions. We discuss these aspects by spelling out the options and considerations associated with each of the stages and parameters of feedforward within the design space. We illustrate these concepts with figures showing how to apply feedforward parameters to specific VR interactions and the design changes it enables.

4.6.1 *Methodology*

The concept of feedforward was initially created to help generate novel interactions for 2D interfaces. In this paper, we have specifically tied feedforward

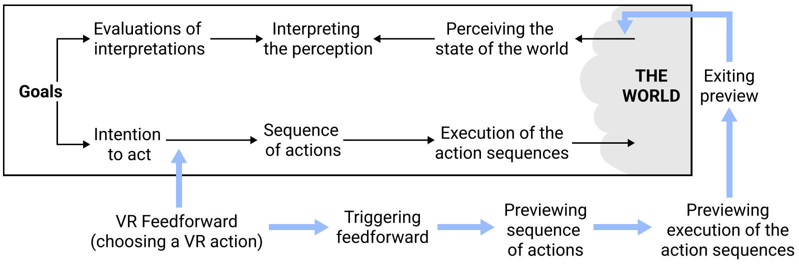

Figure 14: This figure shows Norman's "stages of action model" and the position of VR feedforward within the model. This figure is based on *Figure 1* from [148] and [20].

to VR and imagined a context for its use. For that reason, we call VR feedforward a bridging concept [152]. Although feedforward is a broadly applicable concept, by constraining it to VR, it gets what Beaudouin-Lafon [70] calls generative power, "the ability to help designers create new designs". We developed the design space for feedforward in virtual reality in several phases, which often informed each other:

1. In the **first phase**, we performed a *morphological analysis* starting from the phases of feedforward (shown in Figure 14) and three different VR interactions;
2. In the **second phase**, we aligned the new parameters with the existing literature on feedforward and interaction cueing.
3. In the **third phase**, we implemented feedforward for virtual reality and developed three VR demo applications with varying feedforward designs. New considerations with respect to practice emerged.
4. In the **fourth phase**, we performed an expert evaluation of the design space by running 14 one-on-one workshops with HCI researchers with a background in VR. The experts also reacted to the demo applications and helped formalize design guidelines for using feedforward in VR.

4.6.1.1 *Morphological analysis*

We started with a breadth-first exploration of feedforward parameters in VR using Zwicky's "morphological approach" [2] applied to Norman's stages of action model [144] for the case of VR. We show the first part of this process in Figure 14, where we adjust Vermeulen et al.'s figure to include VR feedforward and the stages (*Figure 1* from [148]). The figure illustrates how feedforward can cross the gulf of execution and reveal the purpose of a system to the user. From Norman's model, we derive a model of a feedforward interaction, which denotes the dimensions of the design space: *triggering*, *previewing*, and *exiting* the preview. We begin the morphological analysis by identifying parameters within these three dimensions.

While the starting point of the morphological analysis is grounded in 2D feedforward and interaction design, VR differs from 2D interfaces because actions may happen at a distance (e.g., through raycasting) and because the scale of the environment is larger than the typical user interface. To account for these VR-specific interaction features, all authors met during three workshops to perform morphological analysis [159]. The design space represents a mapping from an interaction to a feedforward design for that interaction. Therefore, the triggering, previewing, and exiting lie on one side of the Zwicky box. And on the other side, we represented an interaction split by level of detail and targets (action, outcomes, or the avatar).

During the morphological analysis, we designed three Zwicky boxes for interacting with a spray can, operating a forklift, and using the Go-go technique [38] in VR. Combining morphological analysis with real-world interactions helped us ground the process in "concrete reality" [147]. We chose these interactions because they each pose unique problems for the immersed VR user:

1. the spray can has low perceived affordance and requires manipulating before interaction with an occluded object part[1];
2. operating the forklift requires expert knowledge and performing complex operations in a specific order[2];
3. the Go-Go long-arm technique [38, 39] is a VR-specific interaction that combines immersion and embodiment and therefore may reveal specific VR dimensions. Furthermore, generating it does not depend on any existing objects in the environment.

We iterated over each interaction, generating different types of feedforward within the Zwicky boxes, collapsing similar parameters, adding new values for parameters, and adding new parameters if necessary until all authors reached a consensus on the final parameter space.

4.6.1.2 *Theory alignment*

In the second phase of the design space development, we aligned emerging parameters with related work on 2D feedforward research and 3D design spaces. In this phase, we added new parameters from related work or renamed parameters we obtained from the morphological analysis that had a counterpart in the existing literature. For example, at this phase, we included signifiers in the design space to allow users to explicitly engage with feedforward. Signifiers are directly derived from Norman's work on perceived affordances [144]. On the other hand, we identified "granularity," to mean the scope of interaction and renamed it to "level of detail", which appeared with a similar meaning in Bau and Mackay's work on OctoPocus [91]. We aimed to maintain a consistent methodology across the feedforward literature with this process. In addition, we also formalized the definition of VR feedforward, using concepts that underlie feedforward from the literature, such as actions and outcomes [148]. We continue discussing the theoretical background of the feedforward parameters in section 4.6.5. In Figure 15 we show an overview of the background for each parameter and value of the feedforward design.

4.6.1.3 *Feedforward in practice*

In the third phase, we implemented a feedforward authoring system and designed three simple demos to illustrate the concept in practice. The implementation process helped us ground the research in practice [157] and

---

[1] Kingspray Graffiti VR

[2] Forklift Simulator 2019, Best Forklift Operator

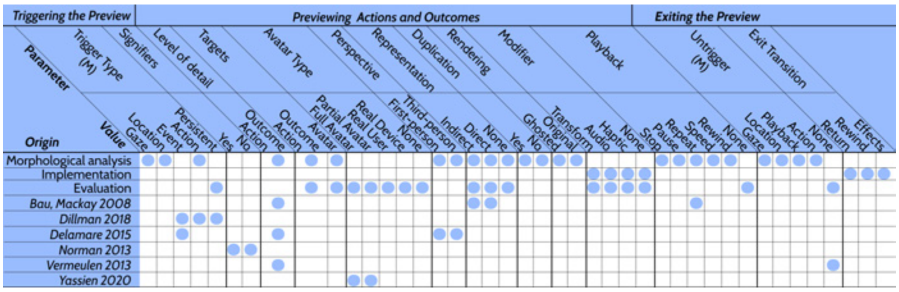

Figure 15: This figure shows the background for the design space of feedforward in VR. The X-axis contains the stages or dimensions of feedforward, the corresponding parameters, and the parameter values. The Y axis contains the development phases of the feedforward design space and the literature that informed it. Parameters marked with *(M)* denote multi-triggers, which may nest conditions.

move from a *general concept* towards a *bridging concept* [152]. The feedforward demos and implementation serve to "demonstrate the scope and potential of" the feedforward concept in VR. Dalsgaard and Dindler introduce the notion of *bridging concept* that captures intermediate knowledge, grounded in theory but "reflecting the span from theory and practice" to reveal "untried design opportunities and potential theoretical advancements". They differ from *strong concepts* [**hook1012strong**] that derive only from theoretical grounding and observation to aid design practices and relate more closely with "concept-driven approaches" to advance theory [119]. The design space of feedforward aims to support designers in generating ideas to show people what to do and how to do it in VR. Thus, the design space can be used as an artifact to aid designers in brainstorming feedforward interaction techniques.

The implementation process yielded new design considerations and revealed potential issues with occlusion and visual clutter. Thus, a gap emerged between the theoretical parameters of feedforward and its practical applications. At this stage, we added the *Exit Transition* parameter to the design space. We also identified ways of dealing with clutter, such as offsetting the targets or designing a feedforward lens.

4.6.1.4 *Expert evaluation*

We performed an expert evaluation of the model proposed above. The aim was to (1) validate the generative power [70] of the feedforward in theory, (2) identify potential issues of the feedforward in practice, and (3) adjust the design space based on the interplay of theory and practice. We ran 14 one-on-one workshops in which an HCI researcher with VR experience used the model to generate ideas, identify potential implementation issues, and specify how to improve the design space. Based on the workshops, we made final adjustments to the design space, identified design guidelines, and future work to successfully integrate feedforward in practical VR applications. We discuss the expert evaluation methodology and analysis at length and present the results in Section 4.9.

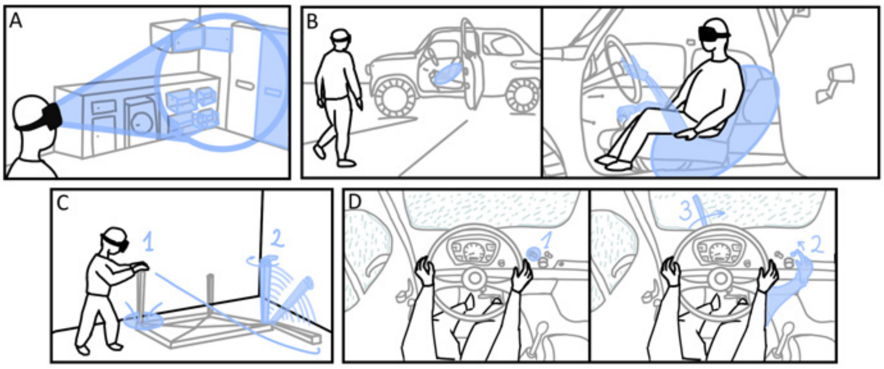

Figure 16: This figure shows different ways to design VR feedforward triggers. In panel A, the user triggers a preview with their gaze. They look around the room to preview possible interactions with objects in their sight. Panel B shows a location-based trigger. When the user is close to the front seat, they trigger a preview that reveals how to start the car. In panel C, the user triggers a preview by doing an action. Once the user bolts in the leg of a table in moment 1, they see the next step in the assembly in moment 2. Panel D shows an example of an event trigger. Once the rain starts in moment 1, an event triggers the preview in moment 2. A ghost hand shows the user how to flip the switch to activate the windshield wipers. The wipers preview after the switch turns in moment 3.

Since this work is informed by theory and practice, to make it easier for the reader, we first provide the theoretical definitions of the feedforward design space in its final iteration. Then, we describe feedforward in practice, the analysis, and the result of the expert evaluation, which yielded changes to the design space. Throughout the paper, we exemplify the concept using two types of demonstrators — the examples of feedforward within the comic strip figures and the screenshots of the implemented feedforward examples.

4.6.2 *Triggering Feedforward*

The trigger represents the mechanism of starting the preview of actions and outcomes in the virtual environment. These triggers represent conditions dependent on the user, on certain objects, or both.

4.6.2.1 *Trigger types*

We identified different types of triggers that vary according to the conditions they must fulfill. Triggers can have more than one type and may require multiple conditions to start feedforward. In such a case, we refer to the trigger as *multi-trigger* and provide a particular graphical representation for it in the design space. Thus, triggers may have a *and/or* relationship. Next, we describe the trigger types we identified, some of which are illustrated in Figure 16:

- *Gaze triggers* – Users may trigger previews by looking at objects or locations. Triggering by gaze is useful when the user's gaze (or estimated gaze, i.e., headset) is enough to signal the user's intent to interact with objects. For example, a user can simply gaze around the room and preview ways of interacting with the objects they have in focus. In Figure 16, panel A, the user sees a preview of how to open cupboards and doors upon gazing at them.
- *Location triggers* – Users may trigger previews by their proximity to objects or by being in certain locations. Location-determined triggers show what actions users can do once they are close to an object (i.e., through proximity). For example, in Figure 16, panel B, the user is shown how to start a car. First, the user is heading toward the car, which has a trigger near the chair. Once the user sits on the chair, the preview starts. The teaser (Figure 12) also shows an example of a location trigger.
- *Actions triggers* – Users may trigger previews after performing particular actions. For instance, if users have to follow instructions that contain many sequences of steps, completing a step can trigger a preview of the next step. We exemplify this in Figure 16, panel C, for the case of assembling a VR table. Once the user bolts one leg of the table, they see how to do the next step (i.e., assembling the next leg). Action triggers represent a group of object conditions that need to be fulfilled in some order by the user explicitly.
- *Event triggers* – Events or states within the virtual environment that do not depend directly on the user may trigger previews. The system fulfills the conditions necessary for events and not the user. For example, the VR environment's weather variables are considered events. We can use an event trigger to show users how to activate windshield wipers during rain, illustrated in Figure 16, panel D. The previews are triggered only when there is rain in the environment. The user has no control over the weather in the system. Event triggers may also nest temporal requirements. For example, an event may trigger a preview ten seconds after starting an application. Thus, context and timing are nested within this type of trigger.
- *Persistent triggers* – Feedforward may already exist in the environment and may not need triggering. Persistent triggers refer to a preexisting feedforward in the environment. The iPhone's "slide to unlock" function is an example of this.

Event triggers combined with other trigger types may create multi-triggers. Multi-triggers allow the designer to add conditions that do not explicitly depend on the user. For example, a cup must be empty to preview itself being filled. An event trigger can add the condition of the cup being empty to a location trigger, resulting in a multi-trigger with an *and* relationship.

4.6.2.2 *Signifiers*

In situations where triggers are optional or need to be made explicit, designers may leave *signifiers*. Signifiers are visual cues in the environment that signal the presence of a feedforward and make triggering explicit. These cues may be text labels or images with additional details about the feedforward,

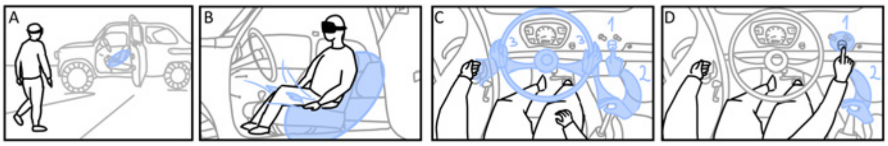

Figure 17: This figure shows how to design different types of feedforwards with varying levels of detail. In panel A, the user approaches a car containing a location trigger on the front seat. In panel B, the user triggers a preview showing how to start the car. In panel C, the feedforward continuously previews all the actions involved in starting the car one after another. In moment 1, the ghost hand starts the engine by pressing the button. In moment 2, the hand puts the car in gear, and in moment 3, both ghost hands steer the wheel. However, each action in panel C can constitute its own feedforward. By varying the level of detail, we could alternatively design a feedforward that shows the user how to start the car in moment 1 and then exits. In panel D, the user triggers the next step with an action trigger after pressing the button in moment 1. Moment 2 shows the new preview, which contains a ghost hand setting the car in gear. Using multiple feedforwards in this way means the users can follow instructions step by step.

such as the type, status of the previewing, or options to pause or interrupt it. The signifiers allow users to choose whether to engage with the feedforward. In the figures from this section, the blue circles are an example of a trigger signifier for the reader. The blue circles convey that a trigger exists relative to specific objects. This rendition is purely for the reader.

4.6.3 *Previewing Actions and Outcomes*

Once the preview is triggered, the next step is to show the users a preview of some actions and their outcomes. The design space for previewing actions overlaps with previewing outcomes, so in our discussion, whenever we mention actions, the same considerations apply to outcomes.

4.6.3.1 *Level of detail*

In VR, interactions may be highly complex, stretching over time and space, and involving many objects and actions. The *level of detail* refers to the order, grouping, and granularity of these actions in one feedforward instance. This parameter allows the designer to select the scope of the interaction to be previewed. The designer may preview the interaction fully, select only the necessary actions to preview, or iterate over specific key static steps. For example, showing users how to operate a virtual car involves actions such as starting it, steering it, stopping it, speeding up, and slowing down. Panels A to D from Figure 17 contain different types of previews involving different steps.

From an interaction design perspective, the levels of detail refer to showing interactions *step by step* or *continuously*. Instead of having one long feedforward previewing all actions, we may split it into smaller feedforward instances, each with a separate trigger, previewing a distinct action. In this

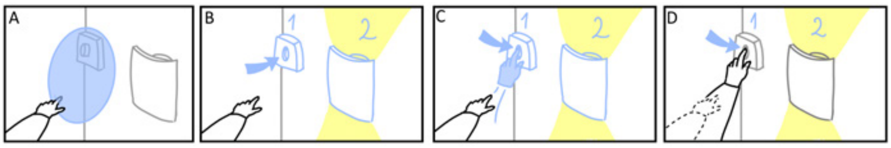

Figure 18: This figure shows how to design different feedforwards using *target* and *representation*. We illustrate this using the interaction of *turning on a light* in VR. Panel A shows a location trigger relative to a light switch. In panel B, the feedforward has non-duplicated direct action and outcome targets. Moment 1 shows the action, namely pressing the switch, whereas moment 2 shows the outcome, the light turned on. In panel C, the feedforward targets all visual elements, the avatar, the outcome, and the action. In panel C, moment 1, a ghosted copy of the avatar presses the button and turns on the light, as illustrated in moment 2. We illustrate embodied feedforward in panel D, where the feedforward controls the original avatar from a first-person perspective. Here, the user's virtual hand (i.e., their un-ghosted avatar) moves toward the light switch and presses it. The intermittent lines denote the user's real-world position outside of the preview. This example also targets all visual elements involved, but in addition, the preview includes the original avatar, which is not duplicated.

way, instead of a sequence of previews within a feedforward, we have a sequence of individual feedforwards.

4.6.3.2 *Targets*

Interactions contain certain visual elements. These can be the avatar that performs the action, the visual elements involved in the action, and the visual elements involved in the outcomes. These visual elements can be interactable 3D objects, text, labels, or other 2D images. When the feedforward includes a preview of the user's avatar, we say that the avatar is a *target* of the feedforward. This avatar may be targeted as a copy or as the original user-driven avatar. The feedforward can target outcomes and actions in the same way, by either previewing them on the original or duplicating the corresponding visual elements. The *duplication* parameter captures this distinction. We illustrate how to design feedforward by changing the targets in Figure 18, panels A to D. *Embodied feedforward* is a particular instance of feedforward when the previews show the user's own avatar acting out the interaction. We illustrate this in Figure 18, panel D. The feedforward may have any combination of the targets mentioned previously (actions, outcomes, and avatars) at any level of detail.

4.6.3.3 *Duplication*

When duplication is selected in the design space, the feedforward shows previews with copies, instead of changing the original objects to preview interactions. This parameter helps distinguish which virtual entities change during previewing: the original visual elements or the copies. The copies can either be ghosted or retain their original material, a distinction captured with the *rendering* parameter discussed below. In Figure 18, panel C, the

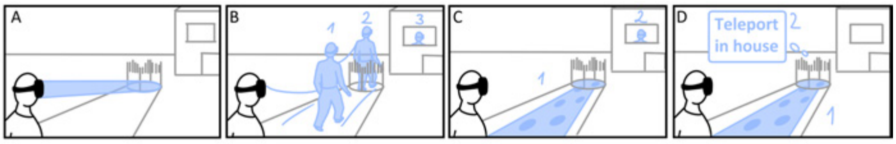

Figure 19: This figure shows how to design different feedforwards by changing the representations of targets. In panel A, the user triggers a preview by gazing at a portal, which would teleport the user inside the house. In panel B, the feedforward previews the interaction using direct targets. In panel C, the feedforward previews the interaction using both direct and indirect targets. In panel C, moment 1, an abstract rendition of steps reveals the action *indirectly*. In panel C, moment 2, the user is teleported inside the house, which is a *direct* representation. In panel D, both the actions and the outcomes are indirect. In moment 1, the user sees the abstract footsteps meant to convey walking to the portal. In moment 2, a text label with the portal's function appears upon reaching the portal.

user triggers a preview of the action required to turn the light on. This preview shows a ghosted copy of the avatar pressing the switch, which means duplication was selected for the avatar. The visual elements of the action and outcome change to a ghosted appearance, but do not duplicate. The *duplication* parameter may be applied to any target. This means a feedforward may contain different levels of duplication, for example by showing previews with copies of the avatar and original objects. When the targets are not duplicated, the original objects and the user's avatar preview these interactions.

4.6.3.4 *Perspective*

So far, we have assumed that previews are shown from the user's tracked perspective; again, VR offers other design options. Rather than using ghosts, designers can let users *embody* the feedforward and experience it from a *first-person perspective*. In these cases, the user's own hands perform the actions. Here, the avatar also moves to a place suited to perform these actions. For example, in Figure 18, panel D, the user may trigger a preview only near the light switch. To design this feedforward with an *embodied perspective*, we could trigger previews from a distance with gaze. The avatar would then move near the light switch during the embodied feedforward as if the user were walking there. Once in proximity, the user's virtual hands would start performing the previewing actions. Conversely, the *third-person* perspective does not change the user's position during the preview.

4.6.3.5 *Representation*

We say the previews contain *directly* represented targets when they simulate the actions and outcomes involved in the previewing interaction. In this case, the visual elements involved move as if the user is interacting with them. *Indirect* representations refer to visual or audio cues that communicate what to do instead of showing what to do during a preview. *Indirect* representations involve some level of abstraction. For example, the iPhone unlock screen contains indirect outcomes and actions in the form of a text label ("slide to unlock"). Upon seeing this message, the user knows that performing the "sliding" action will unlock the phone. A *direct* action would involve the slider moving. Figure 19 shows some examples of direct and indirect representations in VR.

4.6.3.6 *Rendering*

As part of previewing VR interactions, feedforward targets may be *rendered* in different ways with different consequences. For example, Figure 19, panel B, shows a blue ghostly avatar performing the appropriate actions within a preview. The user's original avatar remains unchanged. The different renderings serve to distinguish the previewing avatar from the user's tracked avatar. We refer to this type of target rendering as *ghosted*, which is traditionally translucent and blue. The design space, however, does not specify rendering particularities. As such, the designer may opt for coloring and opacity suited to their context of use. Alternatively, the targets may be rendered in their *original* form. This rendering applies to the previewing targets and may change the original targets if *duplication* is not selected.

4.6.3.7 *Playback*

Since feedforwards show previews of actions and outcomes, they share some properties of video playbacks. For example, previews can be played at different speeds, repeated, paused, or interrupted (stopped), as we discussed previously. Thus, we can repeat the same previews several times until the user understands what movements to perform. To communicate that the previews are static instead of animations, the designer may select the value *none* for the *playback* parameter.

4.6.3.8 *Avatar type*

The *avatar type* parameter describes the avatar that performs the previews. It refers to the duplicated avatar if duplication is selected, or the user's avatar if duplication is not selected. This parameter serves to accommodate interactions without a direct mapping between the user's movements and their avatar. Hand-tracked avatars represent direct mappings, whereas controller-based interactions rely on animations that do not reflect the user's tracked movements. For example, pressing a button on the controller may generate an animation of the avatar that reloads a sling and shoots it. Here, the movements to generate the actions differ from the ones visually performing the action. Therefore, it is important to reveal the correct real-world action to support the intended virtual action to have the desired outcome. In such cases, the designer may choose to visualize both the avatar and the *real user*'s hands with the *real handheld device*. If *none* is selected for the *avatar type* parameter, the previews do not contain the avatar and only show action and outcome visual elements. This parameter also supports the visualization of a *partial avatar*, for example, to allow showing only the user's hands during previews instead of the full avatar.

4.6.3.9 *Modifier*

During previews, targets may change more than their rendering. The *modifier* parameter reflects which target modalities change during previews. For example, a target may change its *transform*, which describes its location, rotation, and position parameters. For example, a feedforward may preview an action at a different place than where the objects are. A target may also change its *audio* during a preview in terms of pitch or speed. Also, in cases where VR is accompanied by haptic feedback, such as controllers vibrating, the design space enables *haptic* modifications during previews. For example, the designer may wish to decrease the intensity of the vibrations or remove them completely to signal to users the action did not take place. If *none* is selected for the modifier, no such changes are applied to targets.

4.6.4 *Exiting the Preview*

Seeing previews aims to help users discover what they can do and how. After learning from the feedforward, users may want to exit the preview to explore other feedforwards or interact with the environment. By definition, previews are transient. As such, to maintain a consistent user experience, the virtual world must return to its initial state before feedforward. When designing the *moment of exit* from previews, we identify two key aspects: when to stop the previews and how to stop them. Figure 21 illustrates a feedforward from triggering to rewinding.

4.6.4.1 *Untrigger*

Starting and stopping a preview are similar concepts. We can stop previewing if the user acts, gazes elsewhere, or moves someplace else. These are the same parameters for triggering the feedforward. However, since we use these conditions to stop the previews, we call them *untriggers*. However, instead of using events, we can stop previewing through *playback*, such as ending a preview after repeating it a few times.

4.6.4.2 *Exit transition*

After stopping the preview, the virtual world must return to its initial state. We identified several ways of transitioning the virtual space before feedforward. First, we can simply play back the previews. This means the interaction is *rewinded* like a video as the user sees the feedforward clock turning back. This sends a clear signal that outcomes do not persist and that the feedforward ends. Second, objects can *return* to their initial states without any transition. Third, we can signal exiting with *visual effects* at the previewing locations, like blacking out the user's view.

4.6.5 *Theoretical background*

The VR feedforward design space integrates knowledge from 2D feedforward and various other frameworks/models and design spaces. In what follows, we describe the theoretical notions that relate to the feedforward design space.

The concept of feedforward is related to Boy et al.'s suggested interactivity and Dillman et al.'s interaction cues [170, 199]. Targets are similar to Boy's attractors because they serve to attract the user to the location of the feedforward. Previewing is similar to animations as it describes the targets over time. Yassien et al. in their design space for social presence in VR [254] discuss a parameter similar to feedforward's *partial avatar type*. Yassien et al.identify the parameter of *completeness* within the *self-embodiment* dimension to describe how much of the avatar is represented in social VR.

Dillman et al. provide a framework describing interaction cues from 49 video games and showcase the usefulness of their framework to generate cues for interacting in augmented reality [199]. They also identify triggering as a dimension and identify four potential values. Thus, in their framework cues may be triggered by a *player*, by *context*, or by some *other agent*, or simply *persists* in the environment. These triggers map to the feedforward design space, namely to the *trigger type* parameter. Here, *action* triggers are determined by the player, while the *event* triggers are determined by the system through context and timing. *Persistent* triggers map to Dillman et al's persistent cues. In addition, during the morphological analysis, we identified *location* and *gaze* and included them in the design space as immersive parameters particular to VR. Dillman et al.'s *markedness* relates to rendering and duplication aspects of feedforward duplication. Ghosted copies can be *integrated* and *emphasized* within the VR scene. In addition, both feedforward and interaction cues encode the distinction between user-driven and system-driven activation.

The models are similar because 2D feedforward techniques are nested within the VR feedforward design space. While Dillman et al. do not mention feedforward explicitly, they discuss labels that tell users what to do. The feedforward design space further builds on these dimensions with details for immersive feedforward such as avatar type, trigger type, rendering, and perspective. It also encodes VR-specific concepts like agency and body ownership.

Delmare et al. develop a design space for designing guiding systems for 3D and 2D gestures that involves feedforward and feedback-specific parameters [162]. Despite being focused on gestures, their feedforward model includes similar parameters like *trigger initiative*, which maps to *trigger type*, and *execution*, which is similar to the *level of detail* parameter. The similarities between the design spaces emerge from the hand-tracking modality of interaction in VR. However, we do not use gesture recognition in the examples discussed in this work.

In this work, the *level of detail* is a combination of Bau and Mackay's work on OctoPocus [91] and Vermeulen et al.'s *static* or *sequential* feedforward. If the *level of detail* is sufficiently small to capture a frame, then feedforward can be represented as a static visualization (e.g., a 3D gesture in VR space). Increasing the *level of detail* to actions, or even multiple actions can describe a sequential type of feedforward that may span whole interactions. The *playback* parameter together with the *level of detail* makes it explicit whether the preview is static or dynamic, like an animation. When *none* is selected for *playback*, the design space describes a static feedforward.

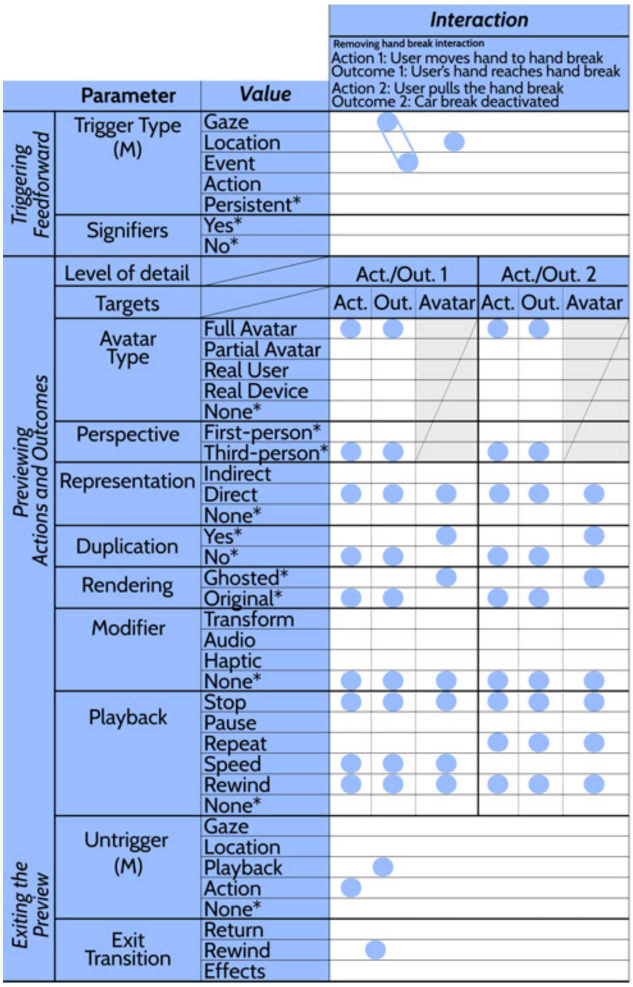

	Parameter	Value	Interaction Removing hand break interaction Action 1: User moves hand to hand break Outcome 1: User's hand reaches hand break Action 2: User pulls the hand break Outcome 2: Car break deactivated					
Triggering Feedforward	Trigger Type (M)	Gaze		●				
		Location			●			
		Event		●				
		Action						
		Persistent*						
	Signifiers	Yes*						
		No*						
Previewing Actions and Outcomes	Level of detail		Act./Out. 1			Act./Out. 2		
	Targets		Act.	Out.	Avatar	Act.	Out.	Avatar
	Avatar Type	Full Avatar	●	●		●	●	
		Partial Avatar						
		Real User						
		Real Device						
		None*						
	Perspective	First-person*						
		Third-person*	●	●		●	●	
	Representation	Indirect						
		Direct	●	●	●	●	●	●
		None*						
	Duplication	Yes*			●			●
		No*	●	●		●	●	
	Rendering	Ghosted*			●			●
		Original*	●	●		●	●	
	Modifier	Transform						
		Audio						
		Haptic						
		None*	●	●	●	●	●	●
	Playback	Stop	●	●	●	●	●	●
		Pause						
		Repeat				●	●	●
		Speed	●	●	●			
		Rewind	●	●	●	●	●	●
		None*						
Exiting the Preview	Untrigger (M)	Gaze						
		Location						
		Playback		●				
		Action	●					
		None*						
	Exit Transition	Return						
		Rewind		●				
		Effects						

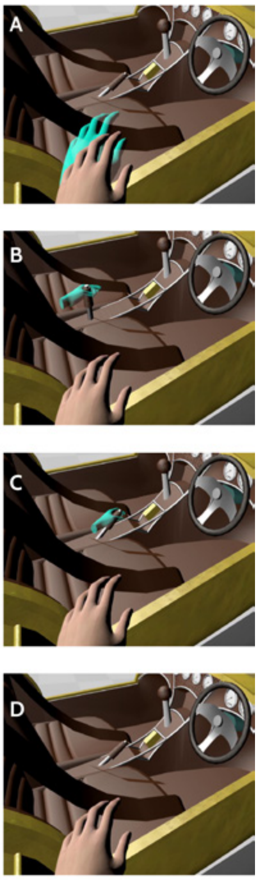

Figure 20: This figure shows an example of a filled-in VR feedforward design space and the corresponding interaction that it generates. The X-axis contains the stages or dimensions of feedforward, the corresponding parameters, and the values each parameter can take. The Y axis contains an interaction split by *levels of detail* and split by *targets* actions, outcomes, and the previewing avatar. Parameters marked with *(M)* denote multi-triggers, which may nest conditions. Values marked with an asterisk may not be combined with other parameter values. Panels A to D capture an implementation of the feedforward designed by this example. Panels A and B show the actions and outcomes 1 and 2, whereas panels C and D showcase the rewinding and exiting moments.

4.6.6 *Using the design space*

The examples of feedforward that we discussed here are non-exhaustive. In practice, values can be combined to generate feedforwards with different degrees of duplication, ghosted rendering, and embodiment. Figure 20 shows an example of a feedforward design space layout that enables this multitude of combinations. For example, most of the previewing parameters may be selected per target, thus enabling unique combinations, such as showing targets ghosted during one action and showing them enlarged during the outcome. Within this design space layout, the fields that are highlighted gray and crossed out are invalid and cannot be filled in. On the right side of Figure 20, panels A to D, we show an example from the feedforward implementation that corresponds to the particular feedforward design space instance. The layout is an artifact resulting from the expert evaluation and serves to accommodate the various design choices the experts wanted to make.

The values annotated with an asterisk (*) may be the sole selection if chosen for that parameter, for a particular target. For example, *duplication* can either be selected or not, and when *none* is selected for *avatar type*, it cannot be combined with any other values. The rest of the parameters can be combined in two different ways. For example, the designer may select to preview several avatar types per action or outcome. In such a case, they would place a dot for whichever avatar type they desired. The fields can either describe an *or* relationship, or an *and* relationship, as shown in Figure 20. We adopt a notation similar to that in the original Zwicky box [159], namely a connected graph notation to denote an instance where multiple conditions need to be met within the same instance. In this particular example, the user may trigger previews when they are close to the car or when they look at the car for a certain amount of time. The gaze and event trigger connect to denote the multi-trigger condition. The disconnected graph notation refers to a single instance.

We make the design space available as a tool for designers using Figma[3]. The Figma template contains several ways of filling in the design space, from bullets that denote *yes* and *no*, to expanding on numerical and qualitative values of dimensions (e.g., speed and action descriptions).

While various types of feedforwards can emerge from the design space, successful implementation requires application-specific considerations. Next, we further develop the notion of feedforward in VR by prototyping it within a VR system and implementing three real-world demo applications.

4.7 FEEDFORWARD IN PRACTICE

In this section, we describe the feedforward system and particularities related to its implementation, such as the parameters it instantiates from the design space and the underlying logic behind the parameters. A video of the feedforward concept and the implementation is shown on YouTube here[4].

---

[3] The Figma design space may be viewed here.

[4] This video is for the benefit of the thesis reader.

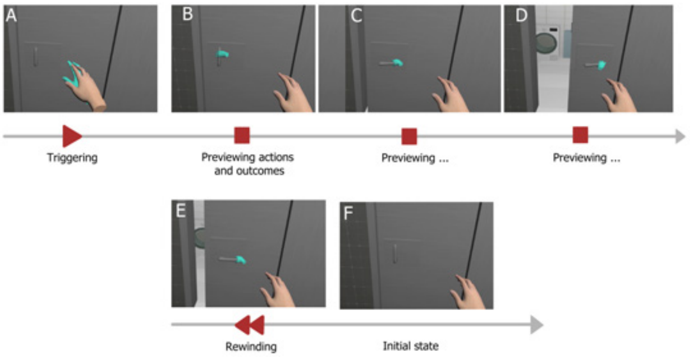

Figure 21: This figure contains an example of feedforward from the implementation. We show the user how to open a virtual door. Its handle suggests pulling, whereas the door opens by rotating the handle and pulling right. In panel A, the user triggers the preview when their hand is close to the handle. Next, the user sees a ghost hand previewing the movements to open the door — the actions in panels B, C, and D. Since feedforward is a preview, the movements are rewinded in panel E. After, the door returns to its initial state in panel F. After the feedforward ends, the user may interact with the door.

4.7.1 *Setup*

The prototype is implemented using the Unity game engine and uses the Oculus Quest 1 headset for tracking. Instead of controllers, we used hand tracking for three reasons. First, it removes the tool between the user and their avatar. This allowed us to capture a broader range of hand movements and use the virtual hand metaphor. Second, we used hand-tracking to better visualize feedforward in-game. When using controllers, the device can occlude the user's hands, and the finger movements on the controller are generally too minute to clearly display in a video or screenshot. Third, using hand-tracking generates an increased feeling of body ownership, allowing us to implement feedforward that is truly embodied [177].

4.7.2 *Pilot*

We arrived at our current feedforward system through an iterative process. First, we implemented the recording of interactions, interpolating the hands to the interaction spot and showing the recorded previews relative to objects. The first implementation only had the user's avatar as a feedforward target, with duplicated and embodied hands. We had a short informal evaluation of our prototype with four members of the Human-Centred Computing section (1 female). The researchers tried to interact with objects of different shapes on a table and triggered previews under two conditions: embodied and duplicated hands. The researchers were encouraged to talk aloud. We found that seeing multiple pairs of similar-looking hands confused the researchers during the informal study. Some thought embodied feedforward was "weird", but they mentioned it was easier to focus without the extra hands. Some researchers found the hand movements ambiguous without the corresponding object movements. We changed the feedforward prototype according to this feedback arriving at the current iteration of the feedforward system. We detail the current feedforward system and other aspects of the implementation below.

4.7.3 *Feedforward implementation*

The implementation contains a few key decisions that we will now describe. For hand-tracking, we used the HPTK library [5] to allow the virtual hands to have realistic, physics-based interactions.

4.7.3.1 *Trigger*

We implemented distance and gaze triggering. The triggers can be placed on any object and reference interactions they can preview. In practice, any object can start the feedforward. Design-wise, the trigger objects should be related to the feedforward. Objects can also have multiple triggers, allowing users to start different types of feedforwards. Triggers also deal with playback management by repeating the previews, either continuously or on trigger. The system keeps track of the state of all triggers, interactions, and hands: if they are active, during feedforward, finished, and so on.

4.7.3.2 *Previewing actions and outcomes*

We implemented interactions by manually recording the user's actions within the virtual space. Interactions are recordings of the user's virtual hands and the objects they interact with. During previews, we simply play back these recordings and adjust the speed as desired. We made the recordings relative to a chosen object in the environment during runtime. This means the system calculates the location and rotation of the hands and objects through a transformation from local space to world space relative to a reference object that does not move. This transformation allows offsetting and rotating the interaction playbacks. If interrupts are disabled, by moving an object during previewing, the playback moves relative to the object. We needed this feature to play interactions accurately if the boundary of the headset is redrawn or if there are changes in the location or orientation of the play area. Scripts that alter targets are disabled when previewing. This prevents any alteration in the movements during feedforward.

4.7.3.3 *Ghosts and targets*

The implementation contains ghosted hands, ghosted objects, target duplication, and embodied feedforward. We can change the material of the non-duplicated targets to make them ghosted, use the non-duplicated non-ghosted targets, or use ghosted copies of these targets. The ghosts are copies of the original objects in the scene, with ghosted materials and no associated scripts. When feedforward is triggered, the ghosted hands appear at the location of the user's tracked hands and then move to the first recorded gesture in the preview (as seen in Figure 21). During embodied feedforward, the tracked hands move to the location of the previews instead. To allow ghosts to appear from the avatar, we keep an invisible ghost hand pool

---

[5] `https://github.com/jorgejgnz/HPTK`

that continuously tracks the user's hands. After the preview ends, the feedforward hand, whether ghosted or embodied, returns to the pool through linear interpolation to the tracked hands.

4.7.3.4 *Exiting previewing*

In the implementation, we can exit the feedforward after some repetitions or through an interruption. If the interrupt feature is enabled, users can interrupt feedforward by touching the targets. During the preview, users cannot interact with the previewing targets. If duplicated, users can only interact with the original objects. If interrupted, the objects return to their initial states without any transition. The exit transition is implemented either by rewinding or by simply returning objects to their initial state. We rewind the interactions by reversing the recorded previews. Also, the rewinding is done at a higher speed than the previewing.

4.7.3.5 *Line of sight and perspective*

We implemented a few strategies to improve the user's feedforward line of sight, by adjusting *modifier*, *rendering*, and *perspective* parameters. We can offset targets during feedforward to prevent overlap with the original objects. We can trigger feedforward from a distance and bring the user to the feedforward location. In addition, we also prototype a technique called *feedforward lens*, which overlays a sphere over the feedforward location. This represents a particular case of the *ghosted rendering* value. The lens only shows the ghosts previewing within the sphere, while the original objects remain unchanged outside of the sphere. The feedforward lens is implemented using the *Amazing World Fading* asset, which can be found on the Unity Asset Store website[6]. We cannot make this part of our code freely available as the asset is not open-source. Instead, the GitHub VR feedforward repository contains an APK file with a lens demo, which developers may install on the Oculus Quest headset.

Having described the implementation and design space, we move on to the practical applications of the feedforward concept. We implemented some interactions that would benefit from having feedforward. Then we designed and implemented feedforward for those scenarios.

4.8 EXAMPLE APPLICATIONS OF FEEDFORWARD

This section contains lessons from applying feedforward in real-world contexts. First, we implemented some interactions that mimicked usability problems presented in Section 4.5.4. Then, we used feedforward to solve them. During this process, we iteratively generated various previews of interactions. After each iteration, we tried the application, noted its effect on user experience, and improved where possible. The aim was to arrive at a feedforward that would benefit the user and minimize any additional cognitive load from the virtual environment. The practical applications here serve as "exemplars" to ground the feedforward concept in practice and reveal new design considerations [152].

---

[6] `https://assetstore.unity.com/packages/vfx/shaders/amazing-world-fading-51172`

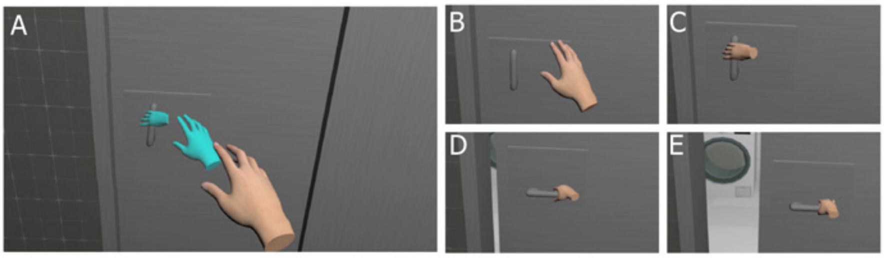

Figure 22: Panel A shows a preview after triggering a feedforward as the ghosted hand starts moving to the handle. This panel does not capture the entire feedforward. The leftmost panels capture the moment after the feedforward ends when the user interacts with the door. In panel B, the door is back to its initial state. The user rotates the handle in panels C and D, and opens the door in panel E by sliding it right.

4.8.1 *Improving perceived interactivity*

Sometimes virtual objects within VR are new to the user. Other times, virtual objects mimic real-world things without operating similarly. When objects do not clearly communicate their use, poorly perceived affordances or wrong signifiers may lead to a breakdown of the user experience. The HCI community often uses Norman doors as an example of bad design (described in *The Design of Everyday Things* [144]). These doors have misleading signifiers that lead onlookers to push instead of pull[7].

We implemented a VR Norman door to show how feedforward can solve issues of incorrectly perceived interactivity. Feedforward can let users know how the door works and prevent possible struggles. We note that Norman doors are merely a placeholder for objects with these usability problems. Often, simply changing the design of these objects is not feasible. For example, objects in VR training applications must be faithful replicas of those in the real world. These are the cases that can benefit from feedforward the most.

The implemented Norman door signals pushing; however, it is opened by turning the handle to the right (Figure 22, C, D) and sliding it (Figure 22, E). We showcase the main phases of the feedforward for the Norman door in Figure 21. For this example, the user triggers the previews when they are close to the handle (Figure 22, A). After triggering, a ghosted hand reaches for the handle and opens the door.

Figure 23 shows selected feedforward variations with different levels of duplication, targets, and ghosting: A — non-duplicated action/outcome targets; B — non-duplicated avatar and action/outcome targets, i.e., embodied feedforward; C — duplicated ghosted avatar target, non-duplicated action/outcome targets; D — ghosted duplicated avatar target; E — non-duplicated ghosted avatar and action/outcome targets; F — duplicated ghosted avatar and action/outcome targets; G - duplicated ghosted action/outcome targets; H — duplicated ghosted avatar and action/outcome targets, with transform modifier offsetting location; I — the user interrupts a preview with ghosted lens rendering; J — the user grabs the handle after interrupt in I; K — the user interrupts a preview with non-duplicated ghosted

---

[7] An example of a Norman door can be found here: https://99percentinvisible.org/app/uploads/2016/02/pulldoors.jpg.

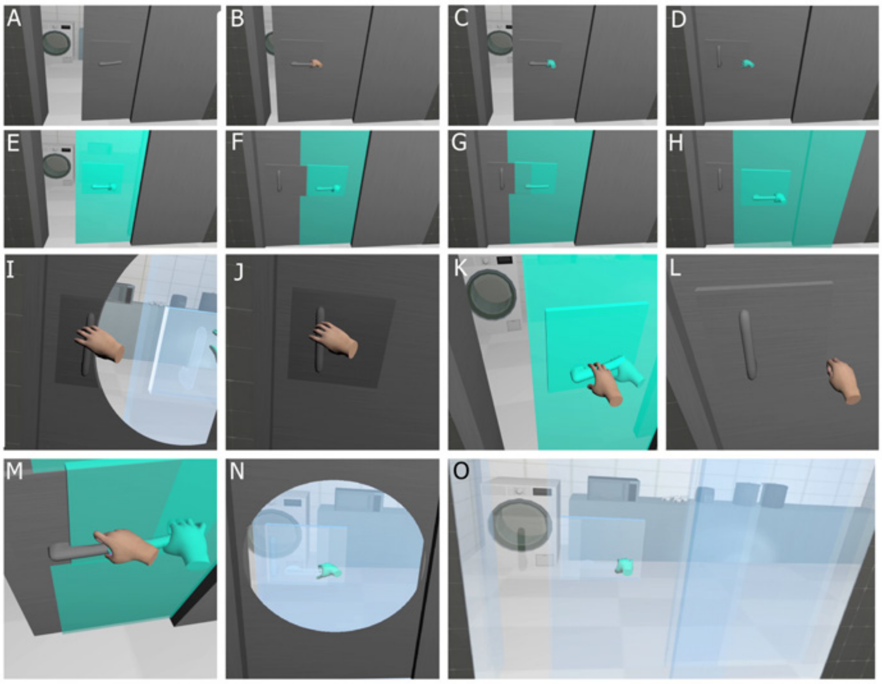

Figure 23: This figure shows variations of feedforward for the door demo, explained below.

action/outcome and duplicated ghosted avatar targets; L — the interrupt shifts the door to its original position; M — the user performs movements at the same time with the preview when interrupts are disabled; and N and O — two different types of feedforward lens. For example, Figure 23, panel A, shows a feedforward that targets the original unghosted objects, while Figure 23, panel G, shows an example in which the targets are ghosted and duplicated.

In practice, we have found varying trade-offs related to ghosted rendering. Mainly, ghosting makes feedforward more salient to the user. In contrast, previews with non-ghosted objects may generate confusion. First, the user may miss subtle movements during feedforward. We observed this when previewing the door. Since the handle is small and fades into the door color, the user may miss the preview if they gaze elsewhere. Second, the user may be unaware a preview is taking place, believing instead that an interaction is occurring. Except for embodied feedforward, we suggest ghosting objects during feedforward whenever possible. Ghosting can help distinguish between interactions and previews and bring the user's attention to the preview.

During embodied feedforward, the user's avatar draws attention to the previews instead of the ghosted rendering. The momentary loss of control over the avatar lets the user know they did not start the interaction. Users can interpret this as the system communicating through the avatar. We show an example of embodied feedforward in Figure 23, panel B. The users cannot act in the environment during embodied feedforward. As a drawback, it can be very jarring for the user to lose control of their avatar. This sudden loss of agency could potentially lower presence and the feeling of body ownership. Despite this, embodied feedforward has one main benefit: focusing the user toward the preview. Keeping track of only one pair of hands and preventing any other interactions could lessen the cognitive load of the feedforward. Thus, using embodied feedforward when the preview requires a high degree of focus could be beneficial.

While embodied feedforward highly focuses the user's attention, duplicating targets enables the user to multitask by repeating the movements during the preview. We exemplify this in Figure 23, panel M, where the targets are duplicated and ghosted, and the user moves the original door during the preview. Especially for long and complex feedforwards, allowing the user to perform the actions simultaneously with the preview could lessen the cognitive load. The number of objects involved and the precision of the avatar's movements represent other factors that could increase cognitive load and add complexity. The feedforward model enables handling such complexity by changing the level of detail or splitting the previews into smaller action-triggered feedforward.

One of the main drawbacks of duplication relates to visual clutter. While transparency is a VR standard for ghosts, we noticed that overlapping more than three transparent objects made them hard to distinguish. Thus as the number of overlapped objects grows, rendering becomes a technical and perceptual challenge. To overcome this, designers may consider adjusting opacity, transparency, and material properties. In Figure 23, panels N and O, we showcase the feedforward lens, which resulted from an adjustment of the ghosted material properties. In this example, the lens follows the hand and differs in color and opacity from the ghosted avatar. The different materials prevent the ghosted targets from causing visibility issues and allow the user to understand the preview.

Ghosts overlapping generate only one kind of visual clutter. Another issue arose from the original objects overlapping with the ghost copies. We solved this by offsetting the duplicates to a different position. Thus, for duplicated targets, the designer may consider moving ghosts someplace the user can easily see the preview. We exemplify this in Figure 23, panel H, where the ghosted door is offset slightly in front of the original door. However, when previews contain complex objects with a high degree of overlap, we suggest removing duplication and ghosting the original targets to avoid this issue altogether.

During the implementation process, we noticed that interrupts could also be jarring when non-duplicated targets are involved. We exemplify this in Figure 23, panel K, where the interrupt is triggered, and panel L, where the door snaps back to the original position. The effect is jarring because, as the user approaches the door handle in panel K to interact, they trigger an interrupt instead. The location of the door then changes after the feedforward ends, leaving the user to grab thin air. To prevent this jarring effect, we recommend avoiding interrupts when the user may easily interact with non-duplicated previewing targets. Instead, the feedforward may continue the preview to completion without repeating it further.

As discussed previously, we designed and implemented the feedforward lens to prevent visual clutter. In addition, the lens may also prevent jarring interrupt effects. With this technique, the original objects overlap with the ghosts but disappear at the lens area, as seen in Figure 23, panel I. Programmatically, the lens is a sphere centered on the ghost hand, which follows the hand, "disappearing" all objects except the ghosted copies. In Figure 23, panel I, the user notices the handle's original location and approaches it correctly. In Figure 23, panel J, the user grabs the original handle after the preview ends instead of the previewing handle that disappears. We show another type of lens in Figure 23, panel O, where intersecting objects are affected. The lens may also be desirable when previews contain large objects to prevent redundant ghosting and visual clutter. As a downside, the lens may hide the original objects preventing users from performing actions during feedforward.

4.8.2 *Guiding users through multistep interactions*

Feedforward has the potential to aid in VR training- and tutorial-type applications. These apps often have interactions that involve many steps in a specific sequence and involve unfamiliar objects. Examples of such apps are assembling 3D printers, putting together furniture, or learning how to drive a forklift. We showcase how feedforward can help users navigate multistep interactions by implementing a virtual car demo. Figure 24, panel C, shows the car's controls: users can change the gear, start the car by pressing a button under a cover and rotate the wheel. Figure 24, panel A, shows a feedforward previewing these controls. The leftmost panels show the user interacting with the car after seeing the feedforward. The user lifts the hand brake in panel B and the cover of the ignition button in panel C. Then the user pushes the ignition button in panel D and changes the gear in panel E while keeping their hand on the steering wheel.Figure 25 shows the rest of the feedforward and other variations.

This example contains two different feedforwards, each triggered by proximity to the steering wheel: a right-hand trigger shows how to use the steering wheel, and a left-hand trigger shows how to start the car. We can see a compilation of the ghosted hands after the moment of trigger in Figure 24, panel A. Figure 25 shows different types of feedforward with varying targets, ghosting, and duplication for the car demo: A — non-duplicated ghosted action/outcome targets; B — ghosted duplicated avatar target; C — non-duplicated ghosted action/outcome and duplicated avatar targets; D — ghosted duplicated action/outcome targets; E — ghosted duplicated action/outcome and avatar targets, with transform modifier by offsetting location; F — triggering feedforward from a distance; G — the user experiences feedforward from the first-person perspective (i.e., embodied feedforward); and H — returning the user when the preview ends; I, J, K, L — showcasing how the user may follow the previews during playback for duplicated

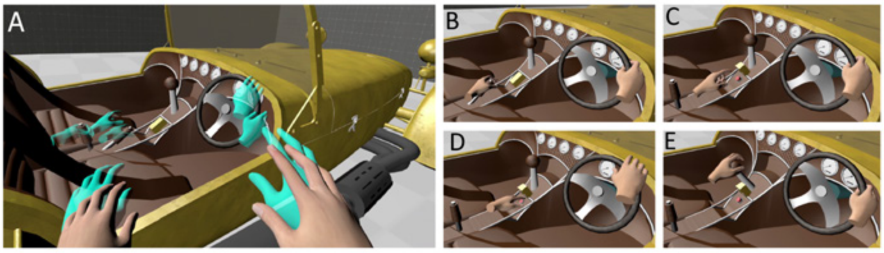

Figure 24: Panel A contains the beginning of a feedforward, which shows the user how to start and steer a virtual car. After triggering, the ghost hands move to the car controls.

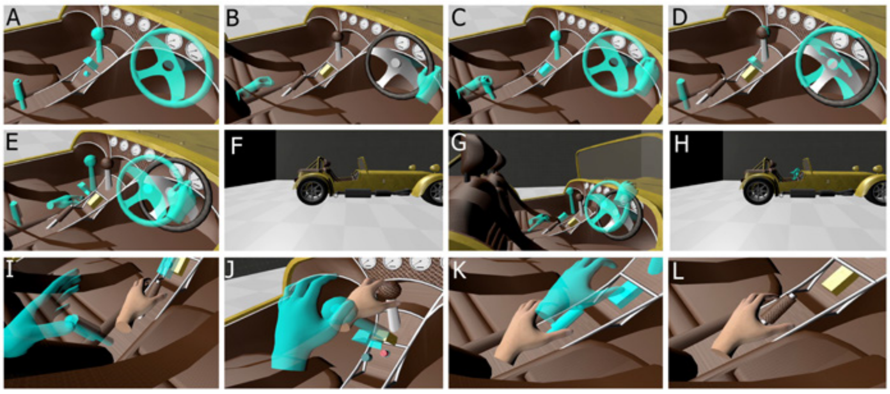

Figure 25: This figure shows feedforward variations for the car demo, described in the text below.

targets and disabled interrupts; in L the previews disappear after playing once.

For target duplication, we observed some immediate problems with occlusion, as shown in Figure 25, panels B and D. In panel B, the button cover occludes the ghost avatar during feedforward. Thus, the user cannot see what the ghost is showing them. In panel D, the original objects hide the ghosted wheel and the button. When the button moves to convey pressing, the user cannot see it because of the cover. In both cases, feedforward would fail to communicate the appropriate actions to the user. Thus, we recommend using ghosting and offsetting the ghost copies to prevent occlusion issues. By offsetting the targets in Figure 25, panel E, the user can see the button under the cover and the wheel.

Multistep interactions are particularly suited for feedforward because users can repeat the actions simultaneously. We showcase this in Figure 25, panels I and J. Here, the user pulls the hand brake up as soon as the ghost hand does the action. In Figure 25, panel I, the ghost hand is slightly ahead of the user. The same appears in Figure 25, panel J, where the user sets the car in gear. The previews must be slow enough to allow the users to follow the feedforward at a comfortable pace. The speed can increase as complexity decreases. However, we note that even simple interactions may cause problems. We recommend approaching the level of detail with caution. If the actions previewed are too complex, the designer can increase comprehension by repeating them, using rewinding, decreasing the level of detail, or decreasing playback speed.

In the demo, feedforward may end by rewinding after one repetition or through an interrupt. We see this in Figure 25, panel K, where an interrupt occurs after the user touches the original car handle. Figure 25, panel L, captures the moment after the exit when objects return to their initial states. While rewinding is good for comprehension and continuity, it has some downsides. First, the user might want to interact with objects after one previewing ends if they understood the actions. In these cases, rewinding could become frustrating. Thus, for simple interactions, we suggest allowing users to interrupt rewinding to interact with the objects sooner. Second, the rewind could signal that the "opposite" interaction has the opposite effect. While pushing the button again does stop the car for this demo, not all controls respond this way. Thus, we suggest speeding up the rewinding to signal that the movements are not part of the preview.

We can also trigger previews at a distance by adjusting the user's field of view and changing their perspective. We showcase this in Figure 25, panel F, where the user is not close to the car when triggering. Since the controls are too far to see clearly, we bring the user to the preview location in Figure 25, panel C. The user remains with this view until the untrigger, after one preview. Then, in Figure 25, panel G, the user returns to their initial tracked perspective. This strategy of *embodying* the feedforward perspective helps whenever the actions are away from the trigger. In this implementation, the user moves smoothly to the feedforward location and back. Instead of linearly interpolating the movement, the designer may use portals or teleportation to change perspectives and prevent motion sickness.

4.8.3 *Feedforward as a tool for discoverability*

Often, designers mimic the real world by bringing familiar objects and interactions into the VR space. This can lead users to have unmet expectations of interactivity, for example, expecting all real-world-looking objects to be interactable. Feedforward can aid users in discovering which virtual objects are interactable and what they do. We showcase the discoverability application of feedforward through a virtual kitchen demo. We implemented a VR kitchen with various objects and interactions available to the user, shown in Figure 26. The user can interact with the following objects in the kitchen: a washing machine, a microwave, cups, an air fryer, a pot, an oven, a fridge, and a light switch. The figure description continues below. We can see a preview of the light switching on and off from panels A to B. In panel F, we show an example of a GUI feedforward technique brought into VR — the iPhone lock screen. In panel G, we show how the ghost hands move the knobs of the oven. In panel H, the ghost hands cannot move the leftmost knobs because they are not interactable. The feedforward here has the objects and the avatar as targets, with the avatar being ghosted and duplicated.

The idea of the virtual kitchen demo is to give the user *an overview* of the available interactions at a glance. For this, we implemented gaze triggering, whereby the user triggers previews when they look around the kitchen. The previews consist of pre-recorded simple interactions with various objects in

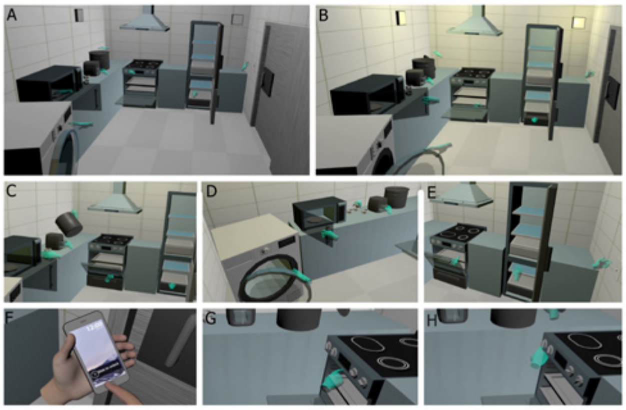

Figure 26: This figure shows the virtual kitchen demo, continued below.

the room, like opening doors or picking up objects. During the previews, ghosted copies of the user's hands fly to the targets and move the original objects. During the implementation, we noticed even simple previews could be overwhelming. Thus, we recommend previewing simple interactions and repeating them multiple times when playing out many previews. Therefore, increasing feedforward comprehension is the goal for dynamic and visually complex environments.

Most demo applications discussed in this section contain co-located actions and outcomes. However, we implemented a light switch to showcase dislocated actions and outcomes for the kitchen demo. We show this in Figure 26, panel A. Here, the preview is triggered when the user looks at the light switch; thus, it is an action-relative trigger. After that, a ghost hand switches on two lamps across the room (Figure 26, panel B). The outcome is not subtle since the lights are noticeable wherever the user looks. However, this may not always be the case for VR interactions. We recommend directing the user's attention to the outcome after triggering when the actions are far from the outcome. Designers may achieve this by changing the perspective, using teleportation, or leveraging other VR techniques like portals. While ghost hands draw the user's attention to the actions, this does not stand true for the outcomes.

We also introduced an example of abstract action and outcome representation in the kitchen demo. While the focus of this work is not 2D feedforward, we implemented the classical "slide to unlock" iPhone screen as an example of indirect representation (shown in Figure 26, panel F). Previous demos mostly showcase direct representations, which previewed actions and outcomes as they would occur when performed by the user. The iPhone feedforward has an indirect action and outcome represented by a text label and an arrow symbol. In contrast with the interactions throughout the kitchen, the sliding control is simple, which lends itself to an indirect representation. When the interactions are complicated and require multiple avatar movements and virtual objects, thus we suggest using direct representations. The light switch control could also have an indirect representation. Instead of a ghost hand, there could be a text label such as "press the button to turn on the light."

4.9 EXPERT EVALUATION

Typically, design spaces are validated in two ways: *theoretically*, in terms of their usefulness in generating new ideas by identifying knowledge gaps and suggesting applications that fill them [261, 220, 196]; *practically*, by implementing novel applications and evaluating their performance [146, 39] and usability [237, 150]. In some instances, design space analysis concepts overlap with how designers talk among each other and may provide additional structure and reasoning for the discussions [29]. So far, we have demonstrated the feasibility of feedforward by showing its usefulness for three VR application scenarios. Through the scenarios, we also uncovered a gap between theory and practice. We conducted an expert evaluation to further explore this gap and assess the generative power of the feedforward design space.

4.9.1 *Methodology*

The feedforward design space is intended as a tool to help VR researchers and practitioners generate ideas about previewing the results of a particular VR interaction. The expert evaluation, therefore, aims to investigate the general capability of the design space with the following goals:

- **Goal 1:** Evaluate the generative power of the design space based on its perceived completeness and clarity.
- **Goal 2:** Evaluate the usefulness of the design space for real-world VR problems.
- **Goal 3:** Improve the design space and develop guidelines for its use based on expert feedback.

We designed the expert evaluation as a series of one-on-one workshops. Each expert was invited to bring one or more VR challenges to the workshop. Here, a challenge represents a VR design question where the expert did not know how to explain to immersed users how to perform an interaction. In addition to these challenges, we prepared two backup tasks for the experts, a *sequential task* and a *system gesture* task. During the workshops, the experts were asked to apply feedforward to a challenge and a backup task chosen by the experimenter.

The *sequential task* involved designing a feedforward that revealed to users how to open a hidden door masked as a bookshelf. The opening mechanism consisted of taking out two particular books in a particular order. This is a common game mechanic in puzzle games, appearing even in VR applications[8]. For this task, the previewed interactions involved objects that already exist in the environment.

The *system gesture task* involved using feedforward to reveal how to use a gesture-based teleportation mechanism. Teleportation is one of the most common navigation techniques in VR and has no direct mapping in the real world [276], which makes it a unique experience not bound to physical reality. For this example, the previewed interactions represented key knowledge the users must have before using the application.

4.9.2 *Experts*

Fourteen researchers and developers with backgrounds in VR and AR participated in the workshop (8 male, 6 female). We recruited participants from the VR research labs at the University of Copenhagen, Aarhus University, and the University of Barcelona. The experts had developed and researched VR/AR between one and six years and were aged between 21 and 38 (M = 28, SD = 4). Participants were rewarded with a 40-euro Amazon gift card for their participation.

4.10 PROCEDURE AND MATERIALS

On average, each workshop lasted 52 minutes (SD = 12 min) and carried the experts through four phases.

---

[8] In The Walking Dead Saints Sinners users must use a secret key to open a hidden bookshelf door.

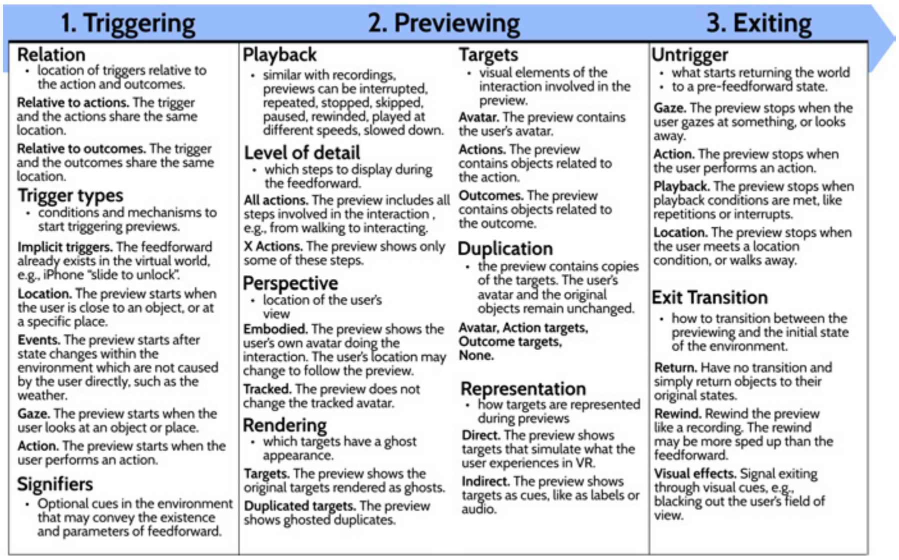

1. Triggering

Relation
- location of triggers relative to the action and outcomes.

Relative to actions. The trigger and the actions share the same location.

Relative to outcomes. The trigger and the outcomes share the same location.

Trigger types
- conditions and mechanisms to start triggering previews.

Implicit triggers. The feedforward already exists in the virtual world, e.g., iPhone "slide to unlock".

Location. The preview starts when the user is close to an object, or at a specific place.

Events. The preview starts after state changes within the environment which are not caused by the user directly, such as the weather.

Gaze. The preview starts when the user looks at an object or place.

Action. The preview starts when the user performs an action.

Signifiers
- Optional cues in the environment that may convey the existence and parameters of feedforward.

2. Previewing

Playback
- similar with recordings, previews can be interrupted, repeated, stopped, skipped, paused, rewinded, played at different speeds, slowed down.

Level of detail
- which steps to display during the feedforward.

All actions. The preview includes all steps involved in the interaction , e.g., from walking to interacting.

X Actions. The preview shows only some of these steps.

Perspective
- location of the user's view

Embodied. The preview shows the user's own avatar doing the interaction. The user's location may change to follow the preview.

Tracked. The preview does not change the tracked avatar.

Rendering
- which targets have a ghost appearance.

Targets. The preview shows the original targets rendered as ghosts.

Duplicated targets. The preview shows ghosted duplicates.

Targets
- visual elements of the interaction involved in the preview.

Avatar. The preview contains the user's avatar.

Actions. The preview contains objects related to the action.

Outcomes. The preview contains objects related to the outcome.

Duplication
- the preview contains copies of the targets. The user's avatar and the original objects remain unchanged.

Avatar, Action targets, Outcome targets, None.

Representation
- how targets are represented during previews

Direct. The preview shows targets that simulate what the user experiences in VR.

Indirect. The preview shows targets as cues, like as labels or audio.

3. Exiting

Untrigger
- what starts returning the world
- to a pre-feedforward state.

Gaze. The preview stops when the user gazes at something, or looks away.

Action. The preview stops when the user performs an action.

Playback. The preview stops when playback conditions are met, like repetitions or interrupts.

Location. The preview stops when the user meets a location condition, or walks away.

Exit Transition
- how to transition between the previewing and the initial state of the environment.

Return. Have no transition and simply return objects to their original states.

Rewind. Rewind the preview like a recording. The rewind may be more sped up than the feedforward.

Visual effects. Signal exiting through visual cues, e.g., blacking out the user's field of view.

Figure 27: This figure describes the design space at the moment of the expert evaluation. A similar figure was given to the experts as a cheat sheet to describe the parameters of the design space.

4.10.0.1 *Introductory phase*

First, each expert discussed the VR challenges they brought with the experimenter (or if needed, the experimenter introduced the backup tasks). Then, the experimenter gave the expert two information sheets describing VR feedforward. One sheet described the feedforward concept, and the second sheet briefly described the parameters of the design space (similar to Figure 27). The expert could ask for any follow-up clarifications about the concept and the design space.

4.10.0.2 *Co-design phase*

Next, each expert designed feedforward ideas for the two challenges or backup tasks based on the design space. The experimenter selected the second task based on the experts' problem. If this problem involved interacting with a sequence of objects, they received the system-gesture task. If it involved performing system gestures, they received the sequential task. During this phase, the expert selected values for each parameter of the design space and generated feedforward instances. The experimenter mediated the session by guiding the expert to their desired solution, answering questions, and clarifying misunderstandings. The experimenter made notes to extend the design space whenever it could not map to the expert's desired solution.

4.10.0.3 *Demo phase*

Similar to related work evaluating taxonomies and design-space driven interactions [266, 237], we performed a short walkthrough evaluation of six feedforward implementations. Each feedforward instance revealed a particular clash between theoretical and practical aspects of the design space identified during development.

More precisely, we were interested in feedback on the ghosting versus duplication designs of feedforward, the embodied aspect of the feedforward, the usefulness of the lens, and the usability of the gaze trigger.

4.10.0.4 *Interview phase*

Last, we conducted a short interview, requesting the experts' direct feedback on the usefulness of the design space, clashes between theory and practice, and theoretical improvements.

4.10.1 *Data collection and analysis*

We recorded the audio from the workshops and collected demographic data about the gender, age, and professional experience of the experts. After transcribing the workshops, we performed a thematic analysis inspired by Braun and Clarke [80] aligning with the goals of the expert evaluation mentioned in Section 4.9.1.

4.10.1.1 *Co-design phase*

First, we coded instances of confusion related to the clarity and completeness of the design space. We formalized *confusion* as instances when the experimenter clarified concepts more than once; when the expert misunderstood parameters or values; and when the experts could not map their desired feedforward instance to the design space. Second, we coded direct feedback about the design space (e.g., comments about the order or graphical representations).

4.10.1.2 *Demo phase*

We coded the experts' positive observations about the implementation and instances of confusion or critical feedback, often voiced directly. The feedback from this phase allowed identify practical design recommendations and matters of further investigation.

4.10.1.3 *Interview phase*

Finally, we coded the experts' answers during the interview as being positive or negative about the model. Further, the experts reflected on their expectations of feedforward after designing it versus experiencing it in practice, revealing gaps between theory and practice.

Last, we identified themes across these codes as opportunities to extend the design space, address *completeness* and *clarity* related to **Goal 1**, and provide design recommendations grounded in practice for **Goal 3**. By having the experts design feedforward based on problems from their own experience and work on predefined VR tasks, we address the usefulness of the design space set by **Goal 2**. Next, we discuss the themes that emerged during the study and the subsequent changes to the design space.

4.10.2 *Results and discussion*

Six of the 14 experts brought challenges to discuss during the feedforward workshop. P1 wanted to show users how to operate a catapult while lying in bed in VR. In their example, the users would have to teleport back while holding the catapult and then release it. P2 needed users to assume a specific position and posture during a VR pointing and selecting the task. P4 wanted to show users how to define particular surfaces for interaction in AR, which required them to perform a series of actions. P7 wanted to explain to their users how to use gaze simultaneously with hand gestures as a keyboard input in VR. P8 found it difficult to explain to participants how to perform specific gestures in a VR study. P10 had difficulties explaining to users how to interact with a control surface, especially without getting outside the play area. For the first task, the remaining eight experts received either of the two pre-designed tasks at random.

4.10.2.1 *Collapsing and removing parameters*

We found that *relation* and *perspective* were the main theoretical sources of confusion within the design space. *Relation* was a previous parameter of the design space that described the trigger locations relative to the action and/or outcomes. The previous iteration of the design space is shown in Figure 27. This parameter was similar to Delamare et al.'s *spacial locality*, which describes the location of the guidance cues relative to the inputs [172]. In total, nine experts asked for further clarifications on the *relation* parameter, and some mentioned being confused (P9: "I'm a bit confused about the relation"). The problem often arose when considering this parameter during the workshop tasks. For most, if not all cases, triggers were relative to both actions and outcomes. We removed this parameter from the design space since it served more as a descriptor of the trigger placement rather than a design variable.

The *perspective* parameter was confusing for ten experts (P6: "I'm not sure what that means"). While most understood the meaning of changing perspective during a preview, their names caused misunderstandings. Initially, instead of *first*- and *second*-person, the values were named *embodied* and *tracked*. P13 mentioned that naming the perspective *embodied* was confusing since it has a different meaning in embodiment literature. To address this, we changed the names of the values from the *perspective* parameter to *first-person* and *third-person*.

The *action* value of the *triggering* parameter posed problems on occasion. For example, P11 could not map the controller input or voice commands to *action*, which involves explicit user input. *Action* may take on many meanings, and, as P9 mentioned, that may be too broad to help generate ideas in some cases. P14 was confused about whether walking toward an object was an *action* or a *location* trigger. The action is important because it captures specific user-avatar requirements for interaction. For example, if the experts must trigger a feedforward only when walking, and not teleporting, *action* is more appropriate. We chose not to expand the *triggering* parameter with types of actions since there are other existing taxonomies about manipulation [49], input [26], and interactions [150]. Instead, we expanded on the types of actions discussed within the paper.

4.10.2.2 *Expanding parameters*

A few experts wanted to make some design choices that did not fit in the design space. Some experts desired different representations for outcomes and actions (P11, P12, P8, P7, P5). To support this, we changed the design space by placing the interaction and targets on the X-axis. Thus, before starting the design process, designers should model their interaction in terms of the level of detail, and identify the outcome and action-related objects or interface items in the interaction. While we discussed multi-triggers (or untriggers), we did not have a way to represent them in the design space. We encode this by marking the bullet point selections in the design space with a connected graph representation, shown in Figure 15.

Some of the experts desired particular avatar representations within the previews. For example, P4 wanted to preview a part of the user's avatar, and P1 had an incongruent mapping between the avatar and the user[9]. In P1's case, the movements performed by the user and by their avatar did not coincide and, therefore, required hand movements to convey the appropriate action. To support this, we added another parameter called *avatar representation* and allowed different types of renditions of the avatar within the preview.

In the study, there were few cases in which the experts wanted to highlight the books as a cue to suggest interactivity (P4, P7, P11). It was not clear to them how to convey a static cue. Then, we expanded the *playback* parameter and the *untrigger* to include *none*. This would allow persistent feedforwards, such as text instructions or signs, without an option to exit. We also change the name of *implicit* feedforward to *persistent* to describe this value more accurately. We added the *modifier* parameter to reflect visual changes to targets, such as changing the location, or size (P4) of targets, since they might play a "bigger role" than initially thought (P2).

Timing aspects also eluded some experts (P7, P4, P8) who could not find a mapping when adding timers to triggers, like gaze or with application start. We collapsed the timing aspect into the *event* value. Initially, there was no timer description in the event, but the *action* value included instances when users would not take any action for some time.

4.10.2.3 *Practical feedback*

In general, the experts agreed that the types of feedforward were clear. Despite having relatively simplistic demos, the most surprising aspect was the added visual clutter from the duplicates, which "clogs the environment too much" (P12). The experts gave the most positive feedback for the car demo in which the actions and outcomes targets were non-duplicated and ghosted, and the avatar was duplicated. The main reason for this preference was the lack of clutter, compared to the other demos. Several experts kept their hands still while waiting for the ghosted hands to finish a preview (P1, P2, P5, P9: "If I don't need to keep my hand there it's better"). The car had two triggers on the wheel, one for the car controls and another for the wheel itself. Triggering both the wheel and car controls was overwhelming for some of the experts (P2, P4, P6, P12). Therefore it seems even with a limited number of objects, the visual complexity and mental load appear to

---

[9] P1 says: "I think it should be the physical controller and not the in-game hands with a virtual hand on the controller that shows you to press the button."

grow exponentially. To solve occlusion and avoid duplication from clutter, P6 suggests reducing the level of detail.

4.10.2.4 *Car demo*

Duplicating the car feedforward targets caused different types of confusion related to clutter, mental load, and sense of space. Most of the experts mentioned that duplication caused clutter and posed challenges to apply in every scenario (P2, P6), being mediated by the interaction space size (P14). Some disliked moving through the ghosts to reach the objects (P14), and instead "automatically" grabbed "the ghost version of the wheel" (P7). For the non-offset targets, P3 mentioned that there were overlap issues when the ghosts reached through the original objects. By dealing with overlap using rendering, P3 mentions, the in-place ghosts would be "ideal" feedforward. The experts were also confused by the offsetting for the duplicated version. It was seen as a bug (P2, P4, P12) rather than a design decision. Offset interactions affected experts' "sense of space" (P5), causing them to adjust to the preview (P8). In contrast, the non-offsetting version made movements easier to relate to (P4), especially when watching and learning from a preview (P13). On the other hand, following complex tasks could only be possible with duplication (P3). Complex actions may need to be previewed more (P10), while the user might be annoyed to wait for simple actions before interacting (P9). With duplicates, users need not wait for a preview to end and may "just start immediately" (P9).

Another surprising aspect of the implementation was the triggering and rewinding. After trying the demo, P10 noted that gaze and location could be too ambiguous as triggers and would include signifiers in their design after all (P10). Other times, location untriggers made the experts interrupt the feedforward too quickly by touching the targets (P2). Moreover, gaze could lead to triggering unwanted previews (P13). Some experts were also confused by the rewinding (P3, P12, P13), thinking they had received instructions to reverse the car, when in fact it had been an exciting transition. Considering this feedback, we may expand signifiers across all phases of feedforward, to also reveal how to stop or pause the playback.

4.10.2.5 *Embodied feedforward*

Most experts disliked first-person perspective feedforward because of the lack of agency that initially confused, (P1, P2, P6, P13), caused discomfort (P8), was "jarring" (P3), and alienating (P12). This type of embodied feedforward could blur the line between what the user wants to happen, and what is happening, perhaps convincing the user they had triggered an animation instead (P4). The experts experienced this first-hand with the door demo, where the experts' movements to open it coincided with the previewing actions (P4, P12, P14: "I thought that I was doing the actions").

4.10.2.6 *Feedforward lens*

The lens received mixed reviews. Some experts disliked it completely (P1, P11). Others liked it (P6, P8, P13), and enjoyed the preview of what lies behind the door (P5). Others thought it could be suited to some scenarios (P9) to help prevent overlap issues, but could add clutter (P3). Some preferred it to be smaller or less colorful (P2, P14), reveal less information about what lies behind the door, and follow the feedforward hand more smoothly (P12). Some experts also reported that feedforward could cause Midas problems [41] (P6) and add world inconsistencies (P5[10]). Indeed, if the user wanted to look behind a door, feedforward would reveal enough information to achieve this goal (P10: "I don't need the preview after that because I already did it").

4.10.2.7 *Kitchen demo*

The experts generally liked the kitchen demo visualizations (P1, P3), but some would have preferred all targets to be ghosted (P3). Some mentioned the movement on the original and unghosted objects could lead people to believe the action was happening (P5: "like something is happening in my world and not something is being previewed in my world"). Indeed, previewing with original objects could also affect how users plan and execute their actions (P9: "I'm not sure how to prepare any actions [...] they are already displaced from their initial starting point."). In addition, without signifiers, the location triggers caused uncertainty (P6: "I didn't know how to get it started again"). The many previews in this demo also posed issues with mental load and caused accidental previews (P6). To lessen the mental load, some experts mentioned untriggering the previews when users would look away (P11, P14). As a limitation of the demo, P7 mentioned that head position had been used to mimic eye-tracking, which was not an accurate eye-tracked interaction.

4.10.2.8 *Usefulness of the design space*

Except for P7, all experts said the design space was useful for generating new ideas. In their view, the design space was "well thought out" because it allowed many possible combinations (P5). Some found the designing experience "educative" (P13) and mentioned they could even map their previous ideas to the design space (P11), which seemed "kind of silly now given a new idea" (P1). All experts found the workflow of the design space useful to model VR interactions by levels of detail and actions and outcomes. Some regarded the design space as particularly useful for training and tutorial applications (P10). It could prevent designers from skipping "details that actually play a good role of [sic] your embodiment" (P12). The design space was also regarded as a useful tool to explore possibilities (P3), arrive at a design that "is quickly feasible" (P5), and "understand the experience of participants" (P13). Others mentioned the design space made them consider aspects they had not thought about before (P8), like exiting the preview (P6).

Some experts had difficulties understanding the concepts "in the beginning" (P6), found some inconsistencies, like persistent triggers not working with the *relation* parameter (P7), and desired a clearer distinction between actions and outcomes (P1). In addition, P7 disliked the mental load from being constrained to the terminology of the design space ("I'm spending my cognitive resources to figure out and try to feed the solution I already have."). They found the design space helpful "as a framework of thinking",

---

[10] When talking about non-duplicate target previews, P5 said: "If I have something on that plate, it feels like it will take out the food I had there because it's kind of doing it."

but preferred describing parameters in their own words with less constricting techniques, like brainstorming.

We re-wrote the workshop design space and made it available as a cheat sheet, or "quick version" of the design space, seen in Figure 50 in Section A.6. The cheat sheet is available as a Figma project here. The Figma template and Section 4.6.6 provide further instructions on using the design space. For the updated feedforward design space, we made further clarifications related to cues and outcome/action visuals, changed the order of the parameters to support the design process (P1, P2, P5, P6), and provided "illustrations" (P6) to aid the decision process (P2, P3). Together, the cheat sheet, the newly added parameters, and the video accompanying the design space describe feedforward design choices grounded in theory and practice.

4.11 DISCUSSION

In the introduction, we asked whether feedforward could be applied to VR with benefits. The answer is positive. We have shown that feedforward can be instantiated within VR. We have also demonstrated how feedforward eases some issues in planning and executing actions for particular applications. For users, the key benefit is that they have learned how to act after exiting the feedforward. More precisely, they should be able to generate the previewed outcome themselves. In practice, feedforward helps generate interesting design ideas and, in theory, provides a useful framework to think about virtual reality feedforward design. Based on the theoretical grounding of feedforward, the experience developing feedforward exemplars, and expert feedback, we provide a few guidelines for using feedforward in practice. By extending the feedforward concept to VR, we aim to "bridge and span the gap between theory and practice" [152] and transform feedforward into a bridging concept. Next, we discuss the theory of feedforward further and raise some questions about using it in practice.

4.11.1 *When and how to apply feedforward*

How to trigger previews? In most cases, triggers may be designed so that the users activate them without thinking, by being in a location, or by performing a sequence of actions. These triggers contribute to the seamless integration of feedforward in applications and maintain the feeling of presence, similar to Vogel et al.'s *self-revealing help mechanism* [76]. The drawback of these implicit (e.g., gaze) triggers is that users may accidentally activate them.

The design space in its latest iteration represents an attempt to accommodate any combination of the outcome, action, and avatar previewing relative to any interaction at any level of detail. Still, during the study, some experts designed guidance cues that do not fulfill all requirements of a feedforward in VR, as we have defined it: to reveal how to perform an interaction and its outcome. Inadvertently, they designed a type of feedforward conceptualized by Vermeulen et al. based on false affordances, namely *hidden feedforward*. These types of feedforward hide outcomes from the user, or, in the case of the design space for VR, may hide the actions and the user representation as well.

Hidden feedforwards may be desirable due to their subtlety and may depend on context. For example, the purpose of a puzzle game is for the user to solve it. In such a context, revealing the solution would defeat the purpose. Instead, the design space can be used to brainstorm hidden feedforward or guidance cues. Frameworks for guidance cues exploring feedforward and feedback parameters have been described before [162]. Thus, we recommend designing hidden feedforward by previewing only actions in contexts where subtlety is desired by the user (**G1**). Pohl et al. provide an overview of applicable subtle contexts based on HCI literature [228]. Some examples include nudging users to look at certain locations or providing additional information without distracting them from their primary task.

4.11.1.1 *Implicit vs. explicit triggers*

Immersion is a key aspect of VR systems. Hand and headset tracking present an opportunity to create contextually-aware interactions which respond to the user in 3D space. For context-aware computing, feedforward is a useful concept to support the *intelligibility* of the system, which researchers consider a "key requirement of context-aware systems [55]."

However, inferring intent is a difficult task. Belotti and Edwards propose design principles that support *intelligibility* for systems that inform users about "what they know, how they know it, and what they are doing about it" [55]. These design principles map well to the concept of feedforward. The researchers provide some recommendations for guiding users to desired outcomes by allowing the user to *correct*, *confirm*, or *reveal* the actions of a system when in doubt. More recently, Fender and Holz introduced a mixed-reality system to capture and preview actions and outcomes in what they call *Causality-preserving Asynchronous Reality*. In this system, a generated causality graph between events and actions represents actions and outcomes. To showcase the system's usefulness, the researchers present a workplace scenario where immersed users may preview past interactions around them while in deep work. The authors discuss similar aspects related to VR feedforward, such as *triggers* and *playback*.

The design space parameters enable triggering feedforward implicitly or explicitly. In the examples of implemented feedforward, the triggers are implicit and do not reveal the existence of feedforward to users beforehand. We assumed that users would be aware of such mechanisms. We have kept this assumption during the expert study, where participants were informed about the types of feedforward before experiencing them. Despite being told where the triggers were, triggering the feedforward was sometimes confusing since we did not inform the experts about specific triggering parameters such as distance. Implicit triggers may share similar challenges with ubiquitous computing. In particular, users may trigger accidental previews when simply acting out in the environment [132], causing "Midas Touch" problems [273].

The *signifier* parameter makes the presence of feedforward explicit to users. Thus, to overcome ambiguities related to triggering, we recommend that signifiers include additional information about the triggers, such as distance or gaze timers (**G2**). We echo Jacob and colleagues in that a useful design does not fully mimic reality, going beyond instead and augmenting it with artificial elements to supplement drawbacks [96]. Outside of tutorials, when users are confused about their next step in VR applications, signifiers can also reveal hidden previews and allow opting into the feedforward.

4.11.1.2 *Level of detail vs. representation*

Level of detail is the broadest parameter and may be unbounded. This parameter allows designers to map interactions by actions and outcomes and with respect to granularity. Thus actions can be modeled continuously or in a static manner. Together with *playback*, these parameter allows the generation of *static* and *sequential* feedforward [148]. The design space allows changes in representation for any level of detail. It enables combinations of indirect or direct actions and outcomes at any point of the interaction and with any avatar representation. However, *none* values must be the sole selection within a parameter. Together, these parameters characterize the preview of feedforward. And except for *perspective* and *avatar type*, they can be applied to the action and outcome elements and the avatar. This enables interesting design spaces where, for example, by selecting an indirect representation for a non-duplicated avatar target, the user's own hand may become an abstract trajectory line telling them where to go. However, we recommend keeping the same style of representations across interactions to maintain consistency whenever possible (**G3**).

Actions appear in the *level of detail* as well as in *trigger*. Indeed there are many broad ways to describe actions, from manipulation techniques [49] to particular sets of gestures [125]. Blom and Beckhaus develop an interaction taxonomy solely for what can be considered actions in VR, and attempt to categorize them further by *dynamic components* [150]. We leave the designer to describe the actions and outcomes for their particular interaction as this can vary greatly in the VR space and may depend on the types of tasks.

4.11.2 *Parameter inconsistencies*

By not constraining the generative power of the design space, we may have allowed some inconsistencies. For example, to generate static indirect representations like images, the user must select *none* for *playback*. This is not immediately obvious and requires the designer to familiarize themselves more in-depth with the concept. By selecting a *playback* value, the designer states the indirect representation is an animation (**G4**). In the same way, *original* rendering may not apply to indirect representations that are very abstract, such as lines or shapes. Since we did not expand the types of indirect representations, there is no way to convey to the reader which specific visualization it implies. Indirect representations incorporate countless ways of designing cues and visual elements, which researchers have attempted to categorize already [199, 170]. The feedforward design space is meant to provide a framework for thinking about how to preview actions, how to trigger previews, and how to end them, specifically for immersive environments. Describing all ways to visualize 2D interactions is not within the scope of this work. For example, the lens does not appear directly in the design space. Instead, the lens is a particular type of *ghost rendering*.

Due to their broadness, indirect representations may add some inconsistencies to the rendering and perspective parameters. For example, combining *first-person perspective* with indirect representations or static direct representations may not make sense. The *perspective* parameter only applies to actions that can be embodied and not static or abstracted actions (**G5**). For example, it would not make sense to change the user's perspective to read a label from a different point of view.

4.11.3 *Usefulness of the design space*

Beaudouin-Lafon [70] suggested that researchers may use interaction models in several ways. First, feedforward in VR may be used *generatively* to design VR. One way to do so is by going through all stages and parameters of the design space and choosing a value for each category. The designer can use the examples and the overview in Figure 15 as inspiration and as a way of going systematically through options.

More concretely, since feedforward involves showing users' actions and the outcome of these actions, the designer must identify (1) which VR interactions they would like to preview and (2) how to preview them in terms of outcomes and actions. This involves identifying the actions and outcomes within the interactions. It also involves identifying the virtual elements which show the interactions (i.e., previewing targets). During the evaluation, most participants chose direct representations and used ghosted hands to preview. Ghosted hands frequently appear in VR, which might have biased them to select this type of design.

Feedforward in VR may also be used *descriptively*. For example, OctoPocus in VR (and similarly 3D OctoPocus) uses indirect outcomes that are hints instead of replicating them [267, 172]. In this application, the user highlights the most likely path as they move their marker avatar. Simultaneously, the less likely paths are more and more transparent. The feedforward targets the actions as possible paths the user may take. The outcomes are also a target, represented by text labels at the end of each path, informing the user of the result. The feedforward design space is useful because it provides a common descriptive ground to compare these types of techniques.

As mentioned in the introduction, feedforward can give benefits similar to those discussed for classic user interfaces [e.g., 148]. In particular, it helps bridge Norman's gulf of execution. Techniques such as Affordance++ [164] help guide users to appropriate actions, whereas feedforward more directly shows actions to users. The concept of previewing actions is similar to Lopes and colleagues' idea of communicating action possibilities through bodies [164]. However, feedforward also allows replicating interactions and previewing them as simulations to the user. Furthermore, a unique possibility of VR is showing how to perform actions using the avatar. The *representation* parameter captures the difference between simulated and abstracted actions. Applying the same concept to 2D interfaces, we can show a mouse cursor moving from A to B. However, we cannot show the user's hand moving the mouse to generate this action. In VR, feedforward may reach outside the system to pull the user's real-world embodied experience within the application.

4.11.4 *Challenges of feedforward in practice*

Implementing feedforward is challenging because VR toolkits do not support it. We intend to solve this by making the feedforward implementation available to designers. LeapMotion and Pavlov VR[11] were two examples of commercial applications with feedforward-like tutorials mentioned by the experts. Indeed, many tutorials or instructions could be explained in terms of Norman's feedforward concept, both for 2D and 3D games and applica-

---

[11] You may find an example of feedforward used in the game Pavlov VR as a tutorial in this video.

tions. Despite this, the design space of feedforward has not been mapped out in VR, and design and implementation guidelines are lacking.

4.11.4.1 *Maintaining coherence*

Having the virtual world fall back to its initial state maintains the integrity and logical flow of events. Within the implementation, we uphold this consistency by preventing the experts from changing feedforward during previewing. The point of the feedforward is to let users know what objects do and how to interact with them. If users can alter the state of objects during feedforward, the message becomes corrupted. This means that feedforward loses the power of the expected result of the action. Therefore, we recommend disallowing changes in previewing (**G6**).

Rewinding is another strategy to maintain the coherence between feedforward and the virtual space. In practice, rewinding returns the state of the virtual world pre-feedforward. It does not surprise the user by suddenly shifting objects. However, it can signal inaccurate actions by implying that the rewinding action does the opposite. This has caused some degree of confusion during the study. We suggest making the rewinding salient, by changing the speed and rendering of actions and outcomes (**G7**).

Another type of coherence relates to the continuity of objects' states. Outside of tutorials, users may interact and change the objects before previews. Thus, objects should fulfill certain conditions before being able to preview actions. For example, an empty container cannot preview how to pour water. To address this, objects could be duplicated at the adequate state without changing the original objects and interfering with continuity. The designer must ensure that the continuity of objects is preserved throughout the application (**G8**).

4.11.4.2 *Dealing with clutter*

Why use previews instead of text instructions? This question captures a trade-off given by feedforward in practice —between clarity and comprehension, mental load, and visual clutter. The implementation should strike a balance between useful information and enough information for the user to be able to generate expected outcomes [55]. As some experts mentioned in the study, this trade-off may not be evident in the conceptual design phase.

If the previewed actions are simple enough, writing some instructional text might be preferable to embodied feedforward previews. However, if the range of movement is broad and includes many actions, users might find it easier to follow feedforward. In such cases, using large amounts of text, labels, and other symbols could make it harder for the user to follow instructions. The benefit of previewing with indirect representations is simplicity. Very familiar actions lend themselves to indirect representations, such as using a switch to turn on a light. When actions are complex, we recommend using direct representations to leave no room for interpretation (**G9**).

In practice, clutter may appear whenever there is duplication. As reported by the experts, most found offsetting ghosts unpleasant, especially when they had to go through the objects. Therefore, minimizing duplication is advised. Moreover, designers should minimize offsetting duplicates, especially toward the user (**G10**). Instead, designers may deal with occlusion by decreasing the level of detail (**G11**). For example, a feedforward may wait until the user has reached a certain action to preview the next step. The feedforward lens could also help in dealing with clutter. However, further work is required to establish an appropriate rendering of the lens before we can recommend it.

Augmenting signifiers or indirect representations with audio or haptic cues presents another opportunity to design unique feedforward and deal with visual clutter. In Norman's view, signifiers need not be only visual representations but include "any perceptible signals of what can be done", like audio or haptic feedback. However, conveying the preview quality of haptic or audio cues remains to be investigated.

4.11.4.3 *Addressing mental load*

The experts from the study remarked on the limits of their perceptual abilities, mentioning that previewing multiple objects was confusing, especially in small spaces. Leveraging human perception to aid in designing interactions has been a common practice among VR developers and researchers [176]. Because of limited multitasking abilities [32], especially when tracking multiple moving targets [99, 92], users could struggle to track more than five distinct objects [21]. By grouping the ghosts in the same material, we leverage human perceptual features like pre-attentive processing, e.g., shape and color [242], to instantly convey that an interaction is a preview. Therefore, we suggest grouping targets by features to leverage pre-attentive processes and lessen mental load (**G12**). Apart from size, location, and color, these features can include playback speed and ghost material adjustments.

For big spaces, clutter might not be a salient issue, though designers should draw the user's attention toward the preview locations. During the study, spawning the ghost hands from the user's avatar seemed useful to draw attention to the feedforward location. This is especially important for cases where the actions, outcomes, or triggers are not in the same place or may lie outside the user's field of view (**G13**).

Apart from clutter, transparency and playback speed may also increase the mental load. Sodhi et al. present an AR system called *LightGuide* that projects on-body feedforward and feedback visualizations to guide the user's hand movements [135]. During an evaluation of *LightGuude*, experts performed movements more accurately with the system compared to watching a video of 3D hand positions. The researchers identified *system-imposed timing* and *self-guidance* as two possible approaches to on-body hint design. Further evaluating several guidance types, the researchers found that a *3D Self-Guided Arrow* had the highest accuracy, albeit the slowest performance. Similar to this type of hint, the playback speed from VR feedforward can alter the timing of the previews. When selecting the playback speed, designers may consider the complexity of the previewing interaction, the users' prior knowledge, and their learning style. Thus, we recommend allowing users some degree of control over the playback, at least replaying and attaching signifiers to reveal the playback control mechanisms (**G14**).

By revealing unnecessary information, the additional mental load may be caused by the ghosts' degree of transparency. Considering this, the feedforward should strike a balance between the provided information and the information needed to perform an interaction (**G15**). While we introduced the feedforward lens to address clutter, some experts suggested it added mental load. As such, further research should investigate how to adequately deal with clutter without increasing mental load. Disregarding transparency entirely when adjusting ghosted materials may improve comprehension. Concerning the intelligibility of a system, Lim and Day further describe ten ways of what can be done [108], later enacting some in a toolkit [117]. When explaining the context in applications, the researchers recommend revealing information related to *input*, *output*, *certainty why*, *why not*, *how to*, and *what if*. The feedforward design space provides the variables to generate visualizations that reveal such information to the user.

So far, we have made some recommendations for the design and implementation of feedforward based on implementation practicalities and the expert evaluation. We show an overview of these recommendations here:

1. **G1:** When users desire subtlety, use hidden feedforward and only preview actions.
2. **G2:** To make feedforward explicit, use signifiers with additional information about the triggers, such as distance or gaze timers.
3. **G3:** To maintain consistency, aim to have similar visual target representations across feedforward interactions (i.e., rendering).
4. **G4:** Use the *playback* parameter to communicate whether the target representations are animations or static previews
5. **G5:** Use *perspective* to state whether the user's viewpoint changes, particularly useful for animations.

6. **G6:** Do not allow changes to targets during previewing to maintain continuity.
7. **G7:** To prevent misunderstandings, make the rewinding salient by changing the speed and rendering of actions and outcomes.
8. **G8:** To maintain continuity, duplicate the targets if objects need to fulfill certain conditions before previewing.
9. **G9:** When actions are complex, use direct representations to leave no room for interpretation.
10. **G10:** When using duplication, refrain from offsetting, especially towards the user.
11. **G11:** To deal with occlusion, decrease the level of detail to preview interactions step-by-step.
12. **G12:** Pre-attentively group targets to reduce mental load, for example, by adjusting playback speed, ghosted material, shape, etc.
13. **G13:** For dislocated actions and outcomes, direct the user's attention to the preview using ghosted hands animations.
14. **G14:** For complex interactions, allow users some degree of control over the playback using signifiers.
15. **G15:** Feedforward interactions should not reveal more information than needed to the user, e.g., through

4.11.5 *Feedforward for learning*

Learning by demonstration is a key principle of human learning, particularly from the perspective of *social cognitive theory* [54]. In VR, *vicarious reinforcement* has shown promise in promoting healthy behavior. For example, when people saw their avatar change weight, they exercised more on the treadmill [103]. Another example of learning by demonstration is *TutoriVR*, which implements a video-based teaching system to aid 3D painting in virtual reality [232]. The immersive system allowed users to better relate to actions in 3D space and perform more critical tasks in the teaching process. Feedforward is a concept rooted in learning by demonstration and, therefore, may benefit learning or teaching applications. In *MirrorFugue*, participants took piano learning sessions with three different representations of the instructor's piano [121]. The researchers found a representation that showed the instructor's piano at 90 degrees in front of their own piano to be the most useful. A representation in which dots indicated the keys to be played had similar learning outcomes. Albeit in AR, designers may use the feedforward design space to describe both types of piano visualizations.

Further work is needed to specify the effects of different feedforward designs on memory and learning. Most of the experts disliked embodied feedforward. However, there may well be appropriate use cases for it. Embodied feedforward was confusing because it prevented the experts from distinguishing between the movements of their hands and the preview. Users of *LightGuide* had similar issues with projected visualizations [135]. Despite this, embodied feedforward creates the least amount of clutter, which was an issue often mentioned by the experts. Furthermore, removing the extra ghost hands might allow people to retain information better. Or perhaps by previewing first and then replicating the preview, users may learn better.

4.11.6 *Feedforward in Augmented Reality*

Feedforward could perform well in augmented reality learning experiences, enabling applications within fields such as situated visualizations and proxemics. Moreover, feedforward may address interactivity problems across mixed reality mediums since some of the challenges we identified at the beginning in Section 4.5.4 also occur in AR, XR, and ubiquitous computing. With a focus on hand-tracking, the feedforward design space also lends itself to real-world use cases for AR and ubiquitous computing. Feix et al. provide a detailed taxonomy of hand grasps, which may be used in related design spaces as immersive visualizations [174]. However, the hand-tracked interactions described within the paper result from the physics engine, not pre-programmed animations. We added the *avatar type* parameter to support integration with various tools like controllers or styluses, which can extend to AR and VR [227].

For piano learning, AR has shown positive outcomes using similar techniques discussed in the feedforward design space [280]. Indeed, researchers highlight the importance of *Action-Concept Congruencies* for embodied learning in immersive environments [142], which underpins the feedforward concept: to link actions to new concepts, or outcomes. In a review of challenges in immersive analytics, 24 experts from various fields create 17 key challenges for future research [264]. The researchers mention the lack of controlled studies concerning the efficacy of visualization methods. This research addresses a few challenges related to spacially situated visualizations. Namely, the tradeoff between the degree of understanding and "information bandwidth". We uncovered potential issues with clutter, transparency, playback speed, and spatial reference, which are particularities that go beyond placement, and relate to immersive qualities of the environment, e.g., rendering. Indeed, in VR, the designer may control the layouts in the environment and minimize clutter. However, when overlaying visualization over real-world objects, we expect issues related to clutter and mental load to the surface to a higher degree. In addition to feedback-related mental load, we provide design recommendations for visualizing and designing *gestural input* that can reveal precise information related to complex interactions involving the environment. Thus, the feedforward design space may describe cross-platform visualizations and establish a common vocabulary to aid in comparing and evaluating such visualizations [264].

Guided visual hints, learning by demonstration, and digital twins are concepts often applied to AR, for example, in manufacturing and design [133, 279]. *Rapido* [270] and *Pronto* [244] are related AR applications that allow developers to prototype and design using *programming by demonstration* techniques by sketching on a tablet. Such systems can be envisioned as practical applications of feedforward in virtual and augmented reality. Jetter et al. implement a VR world editor for simulating *natural* interactions using avatars and 3D objects and report on various complexities of design and implementation [241]. Introducing an interface to record feedforward interactions could enable the implementation of tutorial interactions in 3D spaces through enactment. The goal of such a system would be to allow VR developers to test interaction designs without users.

By revealing interaction possibilities, feedforward extends VR-mediated communication beyond non-verbal means [281], towards immersive learning experiences [271] that are not passive [245] and allow embodied control. Developing a usable, non-technical interface to support recording feedforward interaction remains to be investigated.

4.11.7 *Privacy and ethics*

As a medium, virtual reality sets itself apart by bringing a series of unique ethical and privacy issues related to immersion and presence. For example, Slater et al. discuss these various challenges at length [250]. Regarding feedforward in virtual reality, we identified key challenges in data collection that could lead to misuse. Data management to generate immersive analytics is a growing concern among experts in the field [264]. Mock-ups of environments with different degrees of virtuality are becoming increasingly feasible [256] and together with the "digital twin" concepts [279], AR and VR applications create more and more opportunities to record movement data. Motion data can be used as a playback tool to preview experiences in VR [189] and beyond virtual reality to map movements to robots and play ping-pong in VR with oneself [293]. While these are not feedforward instances, they exemplify the various applications of user motion data and the importance of privacy and data protection. Recording and storing interactions may expose private data to developers. If compromised, such motion databases could lead to instances of identity hacking. We recommend that VR developers and designers follow stringent data protection and sharing practices.

4.12 LIMITATIONS AND FUTURE WORK

The primary limitation of our work is the lack of controlled empirical analysis of interactions enabled by the design space. Instead, we conducted an expert evaluation to address the generative power of the model as a tool for designers. Despite developing a set of practical feedforward recommendations, we echo Ens et al.'s call to address immersive analytic challenges, specifically to investigate how to balance visualization complexity with cognitive load [264]. Further works should investigate how ghosting, transparency, and duplication affect learning outcomes and usability.

We have given a few guidelines on handling complexity; however, none have been precise. The design space enables so many variants for the same interaction that an adequate formal comparison may not be possible. Moreover, we expect these metrics to be application and user-specific. For example, if the target audience of a VR training application is skilled workers, previews may contain more complexity since the users are familiar with the concepts. If the VR application targets novices, even short instructional previews could be problematic. The preview complexity may depend on the user's movement, the number of objects involved (including their parts), the domain knowledge of the user, the VR and technical knowledge of the user, and the end goal of the application. Whether an interaction is better suited for indirect or direct representation may also depend on its complexity. Further work should address the impact of such types of knowledge on user experience.

The lack of novice VR end-users represents another limitation of the evaluation. Momentarily, it is unclear how non-expert and novice VR users may understand and discover feedforward techniques using the design space. A future study with novice VR users may reveal how to enhance the design space for a non-expert audience. Moreover, we acknowledge that the expert evaluation could have suffered from interviewer bias as the experts could be more likely to agree that the design space is useful under the assumption the interviewer participated in the research [127].

The limited scope and simplicity of the scenarios we used for exemplifying practical feedforward represent another limitation. With this work, we focused on the more immersive aspects of the design space, such as direct representations, targets, and hand-tracking. However, indirect representations may yield new and interesting design considerations. In future research, we could focus on establishing the benefits and trade-offs of indirect and direct feedforward, the particularities of the feedforward lens, embodied feedforward, integrating controller-based interactions, and investigating context-aware signifiers and their trade-offs between presence and user experience.

We have discussed potential applications for feedforward design, such as tutorials and programming by demonstration. In future work, the feedforward design space can be extended to other mediums, such as augmented reality. We could also experiment with previewing different types of stimuli like haptic sensations. While we make the code required to develop feedforward freely available, the interface to record interactions is cumbersome and highly technical. The feedforward system could be upgraded with a non-technical interface for recording and previewing interactions.

Lastly, the system may suffer scalability issues when previewing more than ten interactions. As the number of stored interactions grows, it becomes more difficult to record them. To address this technical limitation, we could upgrade the system to use a more efficient way of storing motion data, for example, by storing fewer gestures and interpolating between them to show the interactions.

4.13 CONCLUSION

With the widespread adoption of VR, the complexity of virtual worlds will inevitably increase. Thus, efficiently communicating the action possibilities in VR becomes more pressing. Users need to know about the newly formed multitude of interactions. We have proposed feedforward in VR as a solution to this problem, and we have surveyed the design space of feedforward in virtual reality. We have discussed how the design space can generate ideas for feedforward interactions in VR that apply across a variety of application scenarios and domains.

We have illustrated how to use the design space by implementing a feedforward system in VR, comprising three demos of feedforward in rich settings. We have run an expert study to evaluate the generative power of the design space and compile design recommendations for practical use. We have focused on a breadth-first account of the possibilities of designing for VR using the concept of feedforward. The evaluation also provided expert feedback concerning the implementation and yielded new design changes and practical recommendations. Based on this feedback, we have changed the feedforward design space and compiled 15 guidelines for applying feedforward in practice. Finally, we have made available a cheat sheet version of the design space alongside a Figma template as tools to aid designers in applying feedforward in practice.

Throughout this work, we have introduced feedforward as a bridging concept. While the implementations serve as examplars that "embody the properties of the concept" of feedforward in virtual reality, "reflecting the span from theory and practice", whereby the expert evaluation revealed new theoretical considerations [152]. Together, these contributions illustrate a way of helping users understand what they can do in VR and how to do it, grounded in exchanges between theory and practice.

5

ACTING THROUGH AVATARS

This chapter is based on my latest work, MultipleAvatars, entitled *Why and How to Act Through Multiple Avatars in Virtual Reality*, currently under review: *Andreea Muresan, Teresa Hirzle, and Kasper Hornbæk. "How and Why To Act Through Multiple Avatars in Virtual Reality." In: Pending Peer-Review Process*.

This work also evolved from a simple concept that undertook a complete rewrite. This research highlights the versatility of concepts and how they may evolve and change.

Figure 28: This figure illustrates three key scenarios for acting through multiple avatars in virtual reality brainstormed by experts during the formative workshops. These avatars share the user's appearance and act in sync or out of sync with the user. A illustrates P2's idea (A13) where avatars serve as a *"music looping machines"* to generate performances from past recordings. B shows P4's idea (A41) where a user records a concert by themselves. C shows P4's idea (A43) where a trainer and a trainee embody another avatar to help the user learn movements from the third-person perspective. User 1 (yellow) controls the torso of avatar A, whereas User 2 (blue) controls its hands.

5.1 ABSTRACT

In virtual reality (VR), users typically control one virtual body — their avatar. Previous works enable users to act through multiple avatars simultaneously. These works, though, lack a systematic account of *why* users want to interact with multiple avatars and do not explain *how* users can manipulate and generate avatars. To address this, we run six workshops with 12 VR experts and develop a design space that captures four fundamental dimensions for acting through multiple avatars in VR: *Appearance*, *Context*, *Input/Output*, and *Control*. Researchers can use the design space to generate novel interaction opportunities involving multiple avatars or analyze existing work. We then run a usability study with 17 participants to understand the practicalities of an interface that integrates parts of the design space, which reveals conceptual and technical challenges that we address through design recommendations

5.2 INTRODUCTION

Have you ever considered cloning yourself? It sounds like a scene from a Sci-Fi movie, but, in virtual reality (VR), everything is possible. In VR, you can race against yourself to learn from your performances or generate a whole band performance by yourself. You can jointly control your clone with a professional dance trainer to reveal mistakes in your movement (use cases shown in Figure 28). You can switch between your own avatar and your clone to better understand your movements and the instructor's guidance. You can clone yourself with a friend to brainstorm dancing choreographies for dance performances.

Researchers have explored different types of interactions where the user has multiple avatars for particular use cases, like interacting at a distance to improve selection [277], enabling conversations from different perspectives for therapy [168], visualizing information [289], and capturing and replaying past events [282]. However, in most of this work, the user has no control over how many avatars they possess or how to generate movements for these avatars. the user is put into a scenario with a set number of avatars, leaving questions like creating or controlling more than a few parameters within that interaction unanswered. In works that allow some degree of control over multiple avatars, such as *puppeteering*, the avatars' ranges of motion are limited to head-tracking and preset gestures [222]. This limited control does not fully convey the expressiveness of human motion required for simulating human behaviors through avatars.

In this paper, we explain *why* users desire to have multiple avatars in VR. In contrast with single use cases, we aim for a full mapping of the space of acting through multiple avatars and set out to reveal *how* users may control these avatars. First, we developed the concept of acting through multiple avatars by running six workshops with 12 VR experts to identify *(1) application scenarios*, *(2) system requirements*, and *(3) potential challenges*. Based on the workshops, we generated a design space that maps out different application scenarios of users controlling multiple avatars in VR into key dimensions. Second, based on the identified requirements and the design space, we implemented a VR prototype that allows users full control over multiple avatars, combining motion and 3D controls in a single interface. Finally, we ran a usability study with 17 users to explore the feasibility of the prototype interface and the concept. During the study, users controlled multiple copies of their avatars in four application scenarios derived from the workshops — to play games with themselves, practice arguing, make choreographies, and perform tasks in a pipeline.

We found that avatars can serve as *stand-ins* not only for the user themselves but also for *other people* and *things*. Additionally, we identified *multiplex* stand-ins as particular avatar types that allow users to perceive and act through multiple contexts simultaneously. In the first subsubsection of this introduction, we gave other compelling use cases of multiple avatars that resulted from this research. We discuss how these application scenarios map into four dimensions that define the design space we put forward: *appearance*, *context*, *input/output*, and *control*. The development and evaluation of the interface revealed conceptual and technical challenges related to generating and synchronizing multiple avatars. We discuss how users may overcome these challenges by incorporating mechanisms to reduce clutter within the system and by executing complex operations like switching and implicit syncing to achieve better crowd control. Together, this work aims to help researchers and practitioners organize, define, and implement interactions with multiple avatars in VR. The design space serves as a design tool that captures the most fundamental properties one should be aware of when realizing multiple avatar interactions in VR.

5.3 RELATED WORK

Abtahni et al. reviewed motion-based interactions from the past 30 years and provide a framework splitting VR interactions into *illusory*, *beyond-real*, and *reality-based*. They highlight the need to research interactions that have no direct mapping in the real world and may lead to better learning outcomes of motor skills. While the specifics of a system that supports interaction through multiple avatars, especially for hand-tracking, have not been investigated, themes involving duplication, *parallel avatars*, and *digital doppelgangers* have emerged in the VR interaction landscape. Below, we discuss how different kinds of mixed reality (MR) interaction techniques that involve replication and how this research sets itself apart.

5.4 MULTIPLYING OBJECTS ACROSS REALITIES

The *digital twin* is a related concept that refers to replicating real objects in VR, mimicking their appearance. This concept is used in domains like situated visualization [259] and by guidance and control [287]. *MIRIA* [260] and *Corsican Twin* [248] are both augmented reality (AR) that which allow users to replicate real objects into VR and overlay them on top of real objects, similar with[214, 270, 244], and [167] which leverages digital twins for VR and AR collaborative tasks. *CorsicanTwin* aids in debugging, playback of remote processes, and simulations [248]. Digital twins are generally used for training in fields like manufacturing or in healthcare and education [279]. With respect to interactions, *Voodoo Dolls* is a VR interaction technique that allows people to manipulate objects at a distance by creating a hand-held "doll" [48]. The dolls are temporary copies of distant objects. With this technique, whatever the user does to the dolls is reflected in the original object.

Similarly, the *through-the-lens* metaphor involves replicating viewports and two environments [61]. In this VR technique, the user resides in the primary world having the primary viewport, while they can see a *secondary world* through a *magic lens*. The second world is a referenced copy of the main world and allows the user to navigate or manipulate remote objects within the primary world. *Poros* [275] and *Photoportals* [156] are more recent examples of reference-based techniques which use the portal and photograph metaphor instead.

In these reference-based VR interactions, the environment contains copies that act as a reference to the original object. Because these objects share the same state, when the copies are manipulated, the changes reflect in the original instances. In contrast, Jetter et al. [241]'s VR world editor allows replicating objects without propagating changes and behaving independently, sharing the same context. This enables use cases like simulation and prototyping. This would allow designers to simulate interactions with various technologies, like tablets or touch-based displays, to spur design ideas and identify flaws [241]. *Parallel objects* from Xia et al. [211]'s *Spacetime* combines shared and independent states, allowing users to propagate changes at different stages of object creation. We aim to create an interface to the interaction cloning side of VR-world builders, which has not been achieved with respect to hand-tracking and moves beyond parallel interactions, toward the asynchronous independent object and avatar states.

5.4.1 *Multiplying Places Across Realities*

Besides manipulating objects, the concept of replication can be leveraged in VR for navigation. Stoakley, Conway, and Pausch [37] were the first to introduce the World-In-Miniature (WIM) metaphor in an attempt to provide immersed users with multiple perspectives for navigation [37]. This reference-based technique has been classified as an *exocentric metaphor* whereby it allows users to adopt a "God's eye viewport", interacting outside-in [43].

Indeed, some level of duplication appears frequently in interactions for changing viewports. *Holoportation* is an AR/VR system that reconstructs 3D spaces in real-time for telepresence [179] and is envisioned to keep memories and allow users to receive feedback on dances. Similarly, in *Remixed Reality*, users can cause temporal and viewpoint changes to jump and pause time and inspect themselves [203]. Exploring the concept of being in multiple places at the same time, in *OVRLap* ghosted objects are overlaid on top of the user's primary, opaque environment, allowing the user to change the active viewpoint with their controller [291]. *ShadowSclones* leverages a similar concept in 3D, where the user perceives their cursor in four different contexts simultaneously by splitting the screen into four parts [294].

Another application scenario to include levels of duplication refers to time control. In *Temporal Links*, the authors overlay previously recorded environments in real-time as *ghostly flashbacks* [59], which can be used to leave messages or respond to users asynchronously. Trajectories and timeline scrubbers can alternatively be used in VR to query spacial recordings by manipulating objects to identify causal relationships between events [245]. In *AsyncReality* [282], the authors present future MR work scenarios in which users can enable a *focus mode* to capture and playback real-time recordings of events, such as a colleague leaving a message. Recordings of past events create duplicates of either real or virtual objects, however, they create their own contexts outside of the viewer's (the past), and are limited to viewing, similar to video recordings.

5.4.2 *Interactions With Multiple Avatars*

When we refer to the user's avatar, we refer to the *self-avatar*, which is co-located and synchronous with the user [200]. We refer to the avatars that are not co-located with the user but under their control as simply *avatars*. Multiple avatars are common throughout immersive analytics and serve as a representation of users' movements, usually displaying snapshots of movement across a timeline. For example, *AvatAR*, combines AR and tablet to visualize motion data in-situ [289] by overlaying generic humanoids over trajectories and scatter plots. *GhostAR* is a similar AR temporal and spacial authoring tool where the user creates ghosted snapshots which can then be mapped to different robots and used for programming [214]. With *MoSculp*, users can create printable 3D *motion scultures* from people's movements [212]. The sculptures display the movements as trajectories in time and generate snapshots of the user along that timeline. These techniques separate movements from the user's context and serve only to visualize the interactions, similar to how videos are independent of environments. This means the visualized interactions maintain their own context without propagating changes in the user's environment. The duplicated user movements, represented by avatar copies or snapshots in time, cannot be altered by the user to perform interactions. They may be moved or scaled for the benefit of the visualization, similar to watching a video.

Most interfaces used to control multiple avatars involve motion-capture technologies. To generate tutorials, the programmer must record movements programmatically, which results in a *"cumbersonme"* process [305]. Known as *"digital puppeteering"*, non-continuous gesture-based interaction methods have enabled control of multiple concurrent avatars in VR [258], albeit with lesser degrees of embodiment. In such cases, a human *"interactor"* wearing an exoskeleton triggers animations and voiceovers of virtual avatars mediated by AI [201]. Ingraham et al. [222] discuss how an interaction may control multiple avatars through puppeteering to simulate patients when training therapists, though they do not reveal their interface. Similarly, [230, 277] do not have an in-game interface, and the avatars simply follow the user synchronously without any other affordances, in the latter switching between users with a controller. Indeed, multiple avatars are used to record one's dance moves and preview them in 3D [114] or train on dances from a customizable instructor avatar in VR [297]. Liu et al. [246] develop a Tai-Chi VR learning application that incorporates mocap technology and voice commands, allowing a learner to duplicate their instructor to observe movements from more perspectives. For example, in *Ninja Hands* the user concurrently manipulates four and eight hands with a controller during a selection task [277]. In *Spacetime*, the user may create *Parallel Avatar* objects that can be *dynamic* — copying the user's movements continuously, or *static* — functioning as teleportation pads and switching views using the controller. *Parallel Avatars*. In these cases, however, users may not asynchronously perform changes outside of the *parallel objects* since the contexts are not completely independent. For mixed reality, users may interact with their mirrored and annotated self [100], creating the illusion of looking inside one's body to learn anatomy [126].

We can find asynchronous instances of multiple avatar interactions in body ownership literature. For example, users can alternatively embody two avatars (a lookalike and Freud) and *record* themselves having a conversation by switching between the two using a controller [168]. Having conversations with a Freud-lookalike seemed to improve participants' happiness and mood more than with a self-lookalike. Exploring *vicarious agency*, Gorisse et al. [268] concluded that seeing a self-lookalike avatar correctly interacting with crowds may reduce anxiety by aiding in task planning. The asynchronous embodiment of *digital doppelgangers* may yield different benefits for learning and psychological therapy [112].

5.4.3 *Multiple Avatars in Commercial Applications*

We can find examples of using multiple avatars to interact in WIMP GUIs [1], where the user's avatar (the mouse cursor) is duplicated and interacts with the system. Some examples are to clone interface elements [72], extracting

---

[1] WIMP GUI: Graphical user interfaces based on windows, icons, menus, and a pointing device, typically a mouse.

user-system interactions to generate system use cases [58], to re-create errors [145], or for UI-automated testing [154]. More recently, VR games have emerged where users act through multiple avatars. For example, the VR game The Last Clockwinder allows users to spawn robots that clone their interactions to complete puzzles by pressing and releasing a controller button. The game We Are One has a similar interaction cloning mechanism in the context of a shooter game. In other games such as Rick and Morty: Virtual Rick-ality, users play as clones of fiction characters and may create and interact with concurrent clones themselves. Other times, cloning is used as a story-telling tool or respawn mechanic when the user dies, such as in The Persistence. Outside of VR, some 3D and 2D games use multiple avatar interactions, especially for puzzle-solving. In the game The Swapper, users may control up to four synchronous clones of themselves and swap between them to solve puzzles. In Hourglass, users solve puzzles by cloning themselves, performing interactions for a short time, and then observing a ghosted clone perform them. SOMA is a sci-fi horror game that uses cloning as a storytelling mechanism that confronts its users with existentialist themes about cloning and human consciousness. In the game Elden Ring, players can use an object called the *Mimic Tear Ashes* to create an AI-controlled clone of themselves that assists them in fights. Popular games, like Minecraft and Roblox allow users to make crowds of clones which are asynchronous from the user and may even be scaled[2].

Applications that integrate multiple avatars are emerging within VR, indicating that this concept is useful. Most of these works assume a spatial lens over the interactions, where the spatial relationship of the user within scenarios is emphasized. The user may interact at a distance [277, 294], from different places [291], through portals [275, 156], etc. In contrast with this, we assume the perspective of the user's avatar to investigate the interfaces and affordances nested in these avatars. That is, we investigate how different types of avatars span various interactive scenarios and how an interface enabling such scenarios may look like.

5.5 A DESIGN SPACE FOR ACTING THROUGH MULTIPLE AVATARS

To identify *why* users may desire to have multiple avatars in VR, we involved pairs of VR experts as potential users and conducted six formative brainstorming workshops to identify application scenarios, system requirements, and potential challenges of using multiple avatars to interact in VR. Based on the workshop results, we developed a design space spanning four dimensions defining *how* to manipulate and generate multiple avatars in terms of *(D1)* appearance, *(D2)* context, *(D3)* input/output, and *(D4)* control. In the following, we describe the theoretical foundations and the results of the design space development. Lastly, we explain how to use the design space to *describe and compare* existing work and to *generate* novel designs for acting through multiple avatars.

---

[2] Some examples of users showing demos of the cloning mechanisms can be found on YouTube, for Minecraft and Roblox.

5.5.1 *Methodology and theoretical grounding*

We followed Beaudouin-Lafon, Bødker, and Mackay [257]'s model for *generative theories of interaction*, which describes a *"path from theory to artifact and a principled method for exploring the research design space."* This framework includes designers, developers, and even users in the creation of interactive systems in the earlier design stages [70]. Since our goal is to generate a design space that practitioners can use as a tool when *describing, comparing, and generating* interactions involving multiple avatars [70], we adopted a co-design approach inspired by the requirements elicitation and analysis process [147, 44, 68, 73]. The generative and descriptive power of the design space derives from incorporating the experts' creativity and their field-specific knowledge. Thus, the resulting design space nests *a conceptual space* [149]. Halskov and Lundqvist remarked that design spaces often *"conceptualize design alternatives"* such as *"accumulated knowledge"* or collections of design materials like ideas or sketches [269]. For this research, the workshop materials represent the *primordial soup of design materials, activities, and ideas* used for generating the design space. Similar to related work generating design spaces for VR, we analyze data from formative studies involving experts [272, 305] to derive a parameterized design space through a morphological analysis process [220, 237, 26] combined with thematic analysis [80].

5.5.2 *Participants and procedure*

We recruited 12 experts (3 identified as female, 9 as male) with different degrees of VR experience, of which one self-described as one VR user, four as VR researchers, and one as a VR developer. The remaining four described themselves as a combination of the above. Participants were recruited through the university network and extended social networks and were 28 years on average (SD= 3.9). Each workshop lasted on average 55 minutes and could accommodate remote participation if requested. Thus we had two remote workshops using Zoom and four in-person workshops.

Before the workshop, the experts completed a consent form describing the procedure and data collection and then filled out a demographics sur-

Table 4: This table shows VR the scenarios and applications used during brainwriting — we associated each scenario with a commercial application appearing on the Oculus Store and the Steam Store in July 2022, ordered by popularity and number of sales.

Category	Application
Science, education and training	Gravity Sketch
Sports and improvement	Beat Saber
Social and cultural experiences	Blade and Sorcery
Moral Behavior	VRChat
Travel, meetings and industry	Job Simulator
News and entertainment	Virtual Desktop

vey. During the workshops, the experts discussed acting through multiple avatars in VR in the following three phases:

5.5.2.1 *Phase 1: Introduction*

First, the experts saw a *concept-demonstrating* [44] video, which explained that making an avatar involves recording movements and generating a look-alike avatar that performs those movements.

5.5.2.2 *Phase 2: Interview*

Second, the experts discussed the application scenarios for multiple avatars, requirements or operations suitable for a multi-avatar system, and possible challenges brought on by using multiple avatars in VR.

5.5.2.3 *Phase 3: Brainwriting*

In the third and last phase, the experts were guided through a brainwriting session [16], where each expert wrote down ideas about using avatars for three minutes and then expanded on their colleague's ideas for two minutes. We used VR fields and applications randomly chosen from a list as prompts to *ground* the interaction techniques *"in concrete reality"* [147]. We adapted the scenarios from Slater and Sanchez-Vives [181]'s survey of VR applications. We had six scenarios and six associated applications, as seen in Table 4.

5.5.3 *Workshop Results*

To analyze the brainwriting data, we printed the ideas and used the KJ method [165] to generate a preliminary hierarchical affinity diagram where ideas are categorized in terms of codes and/or subcodes. The KJ method was developed by Kawakita [28] and is an affinity diagramming method consisting of four basic steps: bottom-up label-making, label grouping, chart-making, and explanation. Following this method, all authors met during four sessions lasting one to four hours to iteratively code each idea and merge and collapse codes as needed. Two authors analyzed the interview transcripts separately following the open coding process [107]. Then, the lead author merged the interview and brainwriting codes into the final affinity diagram. Analyzing all coded ideas, we identified the *Stand-In* code, which describes four broad ways *why* a user would like to have more than one avatar in VR and covers 18% of the ideas. The most frequent code here was *Operation* (28%), which categorized ideas that captured system requirements within this code. Other codes reflected *Tasks* (16%), *Training* (15%), *Moral* (6%), *Crowds* (6%), *Avatar = You ?* (3%), *User-to-avatar mapping* (2%), and with a prevalence lower than 1%, we coded ideas as *Memories*, *E-Commerce*, *Privacy*, *Non-avatar*, *User to non-VR Mapping* [3]. We identified 131 ideas from the interview and 328 from the workshop, totaling 459 ideas.

---

[3] We discarded 30 notes from the brainwriting process for being unclear or unrelated.

5.5.3.1 *Application scenarios for multiple avatars*

The subcode *Stand-In For Me* captures ideas where avatars serve as replacements for the user and, therefore, must copy the user's self-avatar appearance. Within all ideas from the *Stand-In* code, 2̃1% captured *Stand-In For Me* use cases. Another 1̃6% of ideas in this category refer to use cases where the user embodies multiple self-avatars, coded as *Multiplex Stand-In for Me*. This particular kind of avatar would enable users to perceive and act through multiple contexts or VR sessions simultaneously. Coded as *Stand-in For Thing*, 9̃% of ideas described using avatars that practically replaced objects, for example, to serve as marks in the environment for mnemonics or progress (P7). However, most *Stand-In* use cases involved using the avatars as *Stand-Ins For Others* (5̃2%) — here, what the avatar does is more important than how it looks. During the workshop interviews, all experts considered that multiple avatars would be suited for application scenarios that involve movement. Therefore, scenarios lacking physical movement (P4, P3, P7, P8) or involving social interactions (P3, P9) were not considered appealing for multi-avatar interactions. Multi-avatar interactions seemed useful for entertainment application scenarios as a game mechanic in "Blade and Sorcery" (P11, P5), especially for multiplayer contexts (P1). "Gravity Sketch"'s repetitive actions could also be useful to multiply (P12) or model as tools for design (P2).

Apart from the *Stand-in* code, we identified other codes relating to application scenarios. Next, we briefly describe them in natural language. Most use cases involved using avatars to perform *repetitive* tasks, to free the user from tedious tasks, and *increase productivity*. Apart from this, many ideas described using avatars for training by learning from recordings or simply as memories of events. Many ideas integrated aspects of switching perspectives to enable better learning for social and psychological training or as an artistic experience. For example, the avatars could teach users how to perform specialized tasks by serving as their practice partners, for example, to practice throws. This also enables users to simulate spacially-aware processes, like generating tutorials that involve particular motions or many objects. Conversely, avatars could simulate social immersion by mimicking crowds. Each idea received a unique code, which we refer to in parenthesis when discussing it below.

5.5.3.2 *System requirements*

The experts discussed basic operations to create, delete (EG40), and hide avatars (EG83), for example, based on event triggers (G63), such as holding doors when opened (A65), motion, gestures, or voice (IA24). Additionally, the experts brainstormed operations to change the appearance of avatars (G23, G23), like clothes and hair (EA36, OP5), material visibility (EG25), deforming its mesh (OP15), scaling (OP3), moving or rotating them (OP46, OP44). To perform repetitive tasks, the experts discussed playback operations like setting the number of replays (EG71), adjusting timing (OP35, OP7) by controlling speed (OP38, OP39), by delaying playback start (EA109), or pausing, for example, through gaze (G40). P6 described navigating by simultaneously controlling two clones at different scales and switching "between them" (OP22), similar to the "world-in-miniature" metaphor (OP26). Others proposed *blending* different users' avatars into one (EG70, EA43, EA76) or partially controlling limbs (EA49, EG8, G4), or changing avatar output parameters, like adjusting gravity (OP13).

5.5.3.3 *Challenges*

The workshops also surfaced potential technical and ethical challenges regarding the use of multiple avatars. Some experts pondered how to sync avatars together (PR3), to keep track of and manage crowds of avatars (G92, EG92), considering that increasing their number would make them harder to manage (PR17), identify (PR1), and increase clutter (PR14), or overload the system (PR20). Realizing interactions where multiple users control parts of the same avatar could also be difficult (PR10), and multi-user interaction could enable users to harass people en masse very easily (G25). The experts also discussed challenges related to movement data privacy and proposed establishing consent or disallowing users from having access to each others' online representations (EA31, A94, EA29).

The experts found realizing multiplex avatars to be challenging conceptionally since users would have to perceive multiple audio (PR9) and visual channels (PR16, PR18) at the same time, which may lead to motion sickness (P5). In contrast, the experts thought purely recorded avatars would require validation, either by the system (EA8) or by the user, "to ensure that the clone does what it's supposed to do" (EG58, PR4). These avatars could appear "creepy" (PR12) due to their limited abilities for social interactions (PR21, PR5). On the other hand, if avatars would become too advanced, cases of mistaken or stolen identity could arise (PR7).

5.5.4 *The Design Space*

To generate the design space, two authors performed *morphological analysis* by creating a pair-wise combination of all parameters [159] based on the affinity diagram and the application scenarios for multiple avatars (seen in Figure 29). The parameters were derived from the affinity diagram by following Braun and Clarke [80] steps for thematic analysis. The final design space is defined parametrically similar to related work [26, 237, 2] in *dimensions (D)*, *parameters (P)*, and *values (V)*. It spans eight parameters nested in four dimensions: *(D1) Appearance*, *(D2) Context (D3) Input/Output*, and *(D4) Control*. Next, we describe the dimensions, parameters, and values:

D1 *Appearance* captures whether the avatar looks like the user's self-avatar or not. It thus has two values: the avatars *Self-Avatar (V1)* and *Non-Self Avatar (V2)*.

D2 *Context* conveys whether the avatars share the user's context or have independent contexts. Thus, it has two values: *Shared (V1)* and *Independent (V2)*. When we say contexts, we refer to the virtual environment and whether it allows the avatar to interact with objects within it. With this definition, objects in different contexts cannot interact with each other.

D3 *Input/Output* captures multi-avatar users and multi-user avatars. It has three parameters: *Input (P1)*, *Output (P2)*, and *Input Type (P3)*. The values of *Input (P1)* describe the number of users controlling avatars from *1* to *N*. For *Output (P2)* this is the number of avatars a user controls from *1* to *M*. *Type (P3)* refers to which type of entities control the

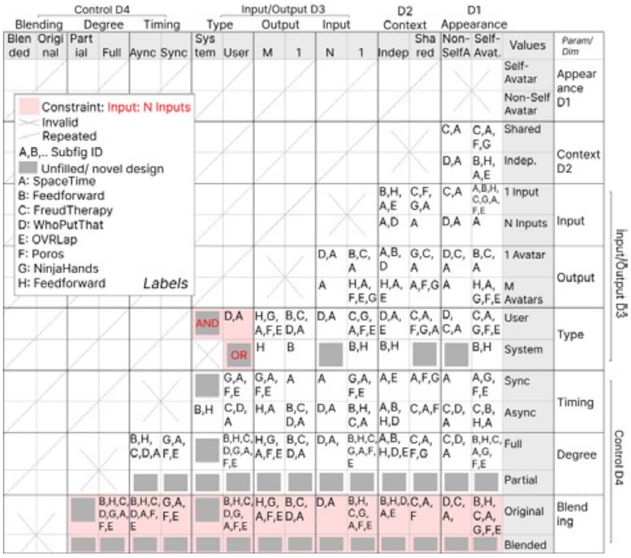

Figure 29: This figure captures the cross-consistency matrix of the design space. Impossible and repeated combinations of the parameters are crossed out. Within cells, we mapped out related work involving multiple avatars into the design space. Solid grey squares mark gaps in existing works that provide novel opportunities for design.

avatars. This can be the *User(s) (V1)* or the *System (V2)* itself through artificial intelligence (AI) or various adaptive heuristics.

D4 *Control* describes how the movement inputs are combined and the outputs are changed. It has three parameters: *Timing (P1)*, *Degree of Control (P2)*, and *Blending (P3)*. *Timing (P1)* refers to whether the avatars move *Synchronously (V1)* or *Asynchronously (V2)* with the user. *Degree of Control (P2)* refers to the part of the avatar a user has control over. This can be *Full (V1)* or *Partial (V2)*. *Blending (P3)* refers to how these parts are combined on the avatar. If the input is not changed at all, this parameter takes the value *Original (V1)*. Otherwise, if two movement sources are combined, we say the control is *Blended (V2)*. Various heuristics may be used, like distance, to generate a final avatar that moves coherently, blending the inputs from two users or a user and the system.

5.5.5 *Using the design space descriptively and evaluatively*

In the following, we discuss the *descriptive*, *evaluative* and, in the next subsection, *generative power* of the design space, based on workshop examples [4] and on related works involving multiple avatars. To use the design space descriptively, the practitioner may look at an interaction involving multiple avatars and select the parameters that describe it. In this way, we mapped out selected works from Figure 30 in the cross-consistency matrix shown in Figure 29 by putting them as single points into the space. The pairwise intersection of the design space allowed us to identify gaps in related

---

[4] Continuing, we refer to the ID of any scenario from the affinity diagram in parenthesis.

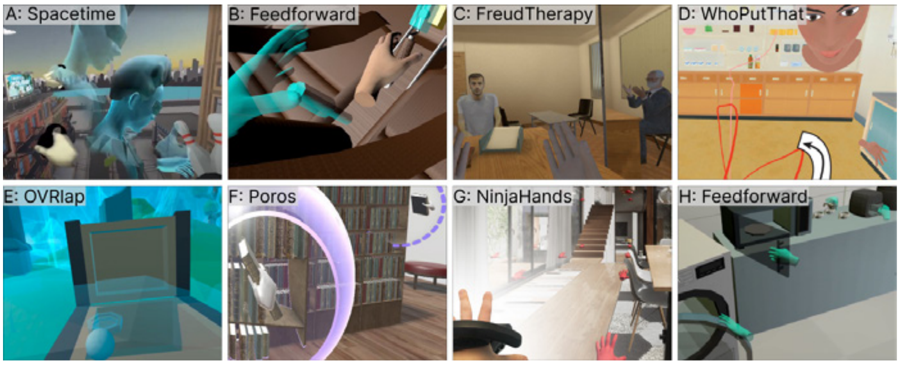

Figure 30: This figure shows selected works that use multiple avatars to: collaborate during world-editing in A [211], follow instructions to start a car in B [305], perform therapy on oneself in C [230], view an immersive recording of someone else in D [245], to interact from multiple places simultaneously in E [291], interact at a distance by creating portals in F [275], and by making many copies of one's hand in G [277], and lastly to preview interaction possibilities in a VR in H [305].

work, which are revealed by unfilled squares. To see the applications fully mapped into the design space, please refer to Section A.1. We provide high-resolution versions of all figures in the supplementary material.

5.5.5.1 *Appearance (D1) and Context (D2)*

Most applications represent the other avatars as copies of the user's self-avatar. In *FreudTherapy*, Figure 30 C the user also embodies Freud, which is not their initial look [230] and acts through the avatar in a *Shared* context, since they see their reflection in the mirror. Shared contexts capture how objects respond to the avatars. In contrast, interactions that are previewed, like recording in *WhoPutThat*(Figure 30 D [245]) or feedforward tutorials in B and H [305] are *Independent*. They may not change states where the user is, or the user may not change states within the recordings. These types of recorded avatars may not capture the likeness of the user, like in D, or may capture it, like in B or H. *Appearance* has a strong evaluative power since it allows the viewer to instantly identify the avatars and possibly their use case. In contrast, *Context* requires a closer inspection of the interaction and may not easily contrast work.

5.5.5.2 *Input/Output (D3)*

Different *Types* of actors facilitate inputs for moving the avatars. In these examples, most users control the avatars, but sometimes the *System* may trigger previews, for example for feedforward [305]. In this case, the user of the system may have no control over the movements generated by the recording. However, these movements were recorded by someone through motion-capture (mocap) technology. Before the user trains on the tutorial, a trainer loads a different interface of the application and records a tutorial for driving a car, similar to Figure 30 B. This use case changes the *Type* to *User* since now the user generates the asynchronous movements.

Concerning *Input* and *Output*, in some applications capture *1* user controls *M* avatars, like in *NinjaHands*, *OVRLap*, *Poros*, solely from their own input. In *Spacetime* [211], for example, a user may duplicate another user's avatar and move it around while in sync with the other user, which would be mapped as *N* users controlling *M* avatars in the design space Figure 30.

5.5.5.3 *Control (D4)*

Timing enables different types of interactions. In applications like *NinjaHands* [277], *OVRLap* [291], and *Poros* [275] the user controls multiple avatars *Synchronously*, which allows them to act from different places. *Asynchronous* control describes recordings that are not real-time like in WhoPutThat Figure 30. *Spacetime* offers synchronous control in *Shared* and *Independent* contexts through *Parallel Objects* and *Dynamic Parallel Avatars* [211]. These types of avatars are duplicated from the user upon teleport, leaving a clone that follows the user's movement.

5.5.6 *Using the design space generatively*

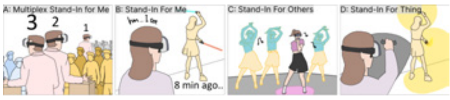

Figure 31: Subfigure A illustrates the same user having a talk in multiple meetings at the same time (A80). Each meeting happens in a different VR environment with different contexts (1, 2 and 3) so the avatars are different. In B, the user learns from the past by watching a recording of their movement in Beat Saber (G8). In C, the user generates backup dances for their performance (A18). Here, the user partially controls the avatars, specifically their legs (yellow), whereas their upper body is generated by the system (blue). In D, the use generates an avatar to light the way while they hold a shield to defend themselves (IA3).

In analyzing the previous works, we identified several gaps, namely related to *Control (D4) Degree* and *Blending*. Continuing with Beaudouin-Lafon [70]'s model for generative interactions, we discuss how to use the design space to generate *novel designs* and *evaluate* or compare design possibilities based on scenarios discussed by the experts during the workshops[5].

5.5.6.1 *Appearance*

Stand-ins reflect a broader theme of interaction since they may span all the workshop data, which we captured in the *Appearance* dimension. When the avatars are used as a stand-in for the user, the observer must believe they are interacting with the user themselves. Figure 31 A shows a user participating in several VR meetings at once, where the avatars must have the user's self-avatar appearance (A32). Conversely, multiple identical avatars can purposefully confuse onlookers. As decoys, multiple avatars allow the user to escape

---

[5] As previously, we refer to the ID of any scenario from the affinity diagram in parenthesis.

enemies (G55) or evade unwanted online conversations (G27, Figure 31 B). For movement or psychological training use cases, the user may observe their posture or facial expressions by syncing their behavior into another avatar (EA111, A111) to see improvements (A70) or mistakes (EA33). When preserving memories (G29) or capturing *"cool winning moments"* (G10), the avatars previewing the recordings also reflect the self-avatar which, in some cases, may be similar to the user's real-world appearance.

Conversely, some scenarios do not require the avatars to resemble the user — they serve as a generic stand-in for VR humanoids. What the avatars do is more important than how they look. For example, in Figure 32 C, the user generates a crowd of avatars to practice giving a presentation in front of an embodied audience (EA34, A46, A80). Generic avatars can also serve as models to preview use cases and clothes (EG48) or for painting certain postures (G18). As a *stand-in for objects*, the avatars may hold objects, like a torch (G13), to free the user to defend themselves in case of an attack in action games (illustrated in Figure 32 D).

5.5.6.2 *Context*

The *Context* dimension marks the distinction between using avatars or previewing avatars, or in other words, performing tasks or viewing recordings. Typically, the user may not interfere with the movements of these avatar types once recorded. Avatars with *Independent* contexts enable the user to generate simulations, for example, for lab safety demos (A45), tutorials (A112), experiments with participants (A47), or stress tests (EA77).

Merging the user's context with the avatar's context enables a whole different set of design opportunities. In *Shared* contexts, avatars allow the user to act out through them and, for example, use torches to light the way. Here, the lighting has an effect on the user's environment, whereas, in an *Independent* context, the outcomes do not propagate in the user's environment. We capture this distinction in Figure 31, B which shows a recorded instance of an avatar, and D, which shows an avatar used to light the way. These types may describe a form of feedforward to preview action outcomes, like teleportation (OP28, OP29).

5.5.6.3 *Input/Output*

So far, most examples we have discussed involved one user controlling one or many avatars. In real life, trainers demo movements for the trainees to follow. However, VR provides new opportunities for learning by demonstration using multiple avatars. Instead, the trainer may assume control over a user's avatar copy to show them how to perform tennis movements correctly (Figure 33 F). The third avatar may coincide with the user to allow a movement perception from the first-person perspective. The third avatar may be located elsewhere to allow the user to perceive their movements from the third-person perspective. Collaborative creative tasks that involve movement may benefit from allowing *N* users to control *M* avatars. This type of *Input/Output* enables scenarios like allowing a director and a writer to control a set of avatars to recreate a particular scene in a movie (G41) or art shows (A113) like choreographies.

During the workshops, the experts discussed avatars that are not purely recorded by the user. These use cases uncovered a gap in the related works analyzed previously, highlighted in Figure 29. For example, to enable use

Figure 32: Subfigure E shows a user and two avatars (yellow) that help with moving objects on a pipeline (EA50); the avatar's arm (blue) is adapted by the system to adjust to the grip of the box. In F, we show a user S and a tennis trainer T controlling avatar A at the same time (A43), where the user has control over the torso (yellow) and the trainer has control over the arms (blue). In G, we show the user escaping a confrontation with an online bully by activating a crowd of clones (G27). Yellow avatars act in sync with the user while blue avatars are controlled by the system. In H, we show a user escaping enemies in a game by creating an avatar clone (G55).

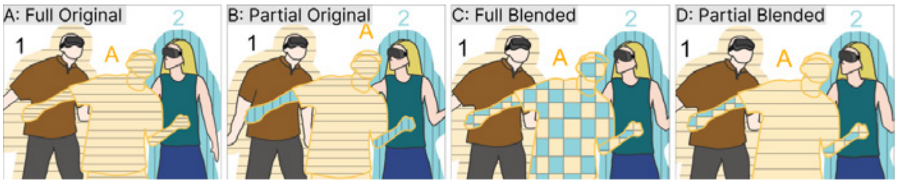

Figure 33: This figure illustrates how to control the motion of avatars through the *Degree* and *Blending* parameters. Subfigure A shows User 1 *fully* controlling Avatar A with their *original* movement (yellow). B shows User 1 (yellow) *partially* controlling Avatar A's torso, while User 2 (blue) *partially* controls Avatar A's arms, with *original*, unblended control. C shows User 1 (yellow) and User 2 (blue) *fully blending* the control over Avatar A. D shows a combination of control — User 1 (yellow) and User 2 (blue) *partially blend* control of Avatar A's arms, while User 1 maintains *full original* control of Avatar A's torso.

cases where avatars are enhanced with human reactions and reduce "awkwardness" (EA80). These avatars go beyond purely user-controlled recordings and towards complex human-system-controlled avatars. Such avatars may eventually perform highly autonomous tasks like learning from the user and cloning them (A29, G62). For now, these avatars may perform simple adaptations like adjusting their walking to different terrain, like stairs (OP12), adjusting grips when picking up objects (OP9), catching throws automatically (OP32), or randomizing actions like throwing tennis balls (EG41, EG46) to prevent the user from being bored in one task. The experts also discussed allowing the system to completely take over avatars for training, for example, to reveal mistakes in movement (A30) (similar use cases shown in Figure 33. The gap annotated with *and* and *or* from Figure 29 refers to these scenarios. The design space enables a *User* and the *System* to control avatars either at the same time (*and* highlighted gap), or take control of them separately at different times (*or* highlighted gap). The *Control* dimension describes in more detail how these types of inputs may be combined.

5.5.6.4 *Control*

The *Timing* parameter is a useful feature of the design space, as it enables the practitioner to think about tasks in two ways — those that can be parallelized and those that cannot. *Asynchronous* avatars enable parallelization, to do repetitive (IA59), tedious tasks, whereas *Synchronous* avatars enable perception —- seeing oneself perform to improve or learn. Still, during the workshops, the experts discussed substituting dance partners with avatars (A89) or using the avatar as player two (A63). For more adaptive and dynamic use cases, the avatar's movements may be combined and blended into a desired solution to a problem. For this, we may use the *Degree* and *Blending* to describe how to merge the motion controls. For example, users may control avatars to a *Full* or *Partial* degree to practice only certain parts of movements (EG8) or mirror movements on the same avatar (G89). When the control is *Original*, it means the user maps their own movement to the avatars perfectly (Figure 33 A). However, during the workshop, the experts mentioned adaptable avatars catch throws (OP32) or change timings (OP35). In this case, the avatar's movements are *Blended* with the system to maintain the user's intent but adapt to performing the task (Figure 33 C, D). These two parameters, *Degree* and *Blending* capture gaps in selected related works, highlighted in Figure 29. Most works discussed fully original avatars (Figure 33 A), whereas the design space enables partially original control (Figure 33 B), fully blended control (Figure 33 C), and partially blended control (Figure 33 D).

The experts also discussed switching to users to embody different avatars within their control to respond to actions instead of adapting to them automatically, implemented within the Freud scenario [230]. Multiplex avatars are particular types that emerged during the workshop when the experts discussed being in multiple conversations or meetings at the same time (A32). In these cases, contexts are independent, yet the user has synchronous full control of multiple self-avatars[6]. During the workshops, the experts discussed practical ways to leverage this concept, for example, to allow researchers to run multiple experiments simultaneously by being in multiple VR scenes (A47). *OVRlap* provides an implementation of this type of avatar, where the user perceives simultaneous contexts but may only act in them independently.

5.6 PROTOTYPING AN INTERFACE FOR MULTIPLE AVATARS IN VR

Since even good models may evoke "terrible interfaces", design spaces do "not guarantee" quality in the resulting design [70]. Therefore, a practical evaluation grounded in usefulness is key in revealing the quality of designs in practice [94, 89]. Continuing with the breadth-first approach of this research, we present qualitative results from a usability study with 17 participants who use an interface to generate and manipulate avatars in four different tasks. The interface represents an *artifact* contribution [184] and serves to filter parts of the design space [269] to allow us a better under-

---

[6] In optical communication systems, multiplexers allow transmitting multiple simultaneous channels through one optical fiber [14], hence the name for multiplex avatars. Because they allow combining multiple contexts into one and propagating the user's actions in multiple contexts

standing of how the design space works in practice and reveal insights into usability and implementation.

5.6.1 *Methodology*

We probe the design space by implementing avatars that share the user's self-avatar appearance and perform asynchronous operations in shared contexts. We followed a walkthrough procedure [36], which is common in breadth-first approaches to VR problems [266] and multi-functional XR systems like *Spacetime* [211], *Poros*, and *VRSketchIn* [237]. To aid in learning the interface, we introduced tasks of increasing complexity, similar to Jetter et al. [241]. The tasks represent common application scenarios from the workshop and determine participants to use the avatars as a *stand-in* for *themselves* and for *others*. To enable users to perform the tasks, we implemented 14 avatar operations based on the formative study operations, which we chose based on their frequency and prioritized CRUD[7] operations. We incorporated hand-tracking in the interface since it requires no prior training on the user's part and gives users a higher degree of control over motion than controller-based gesture generation from related work [222].

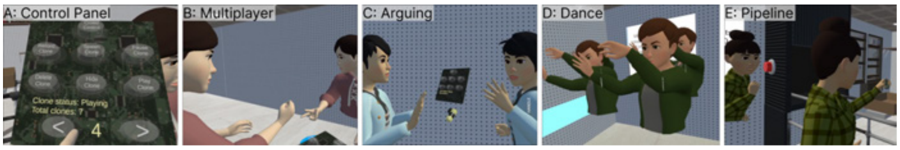

Figure 34: This figure captures screenshots from the participants of the usability study. Subfigures A shows the control panel from P2's perspective, while B shows P2's recorded avatars playing *rock-paper-scissors*. C shows P2's avatars playing *Patty-cake*. D shows P13's avatars performing a choreography during the dance phase. The avatar can be moved through a handle in the middle of its chest, shown as a white text box. E shows P16's avatars helping move objects during the pipeline task.

5.6.2 *Implementation and apparatus*

We implemented a control panel that allowed users to create, delete, and hide avatars (seen in Figure 34 A)[8]. We added a handle inside the avatars to allow users to move them, which was visible through all objects as seen in Figure 34 D. Most operations could be performed individually or for all avatars. Figure 34 A shows the crowd view of the control panel. Upon creation, each avatar received a unique ID, which appeared over their head. Users could also set the number of replays for the avatars' playback or pause and resume them. These playback operations could be used to sync the avatars with each other and the environment, all of which are described in more detail in Section A.4, and in a video shown on YouTube here[9].

---

[7] CRUD: Create, read, update, and delete.

[8] We attach in the supplementary material higher resolution images.

[9] The video was uploaded on YouTube for the purpose of the thesis reader.

We used Unity with Oculus SDK [10] and Meta Avatars[11] for the implementation. We used an Oculus Quest 2 running on a Lenovo ThinkPad T16G and NVIDIA GeForce RTX 2080 Super, with *Windows 11 Home*.

5.6.3 *Materials and procedure*

Upon arrival to the study, participants completed the consent form and a demographics survey. Then, they selected an avatar that resembled them most from a list of 31 Meta Avatars[12]. After being immersed in the VR application, the participants performed a mirroring task to aid with the feeling of body ownership [115] by inspecting themselves across a synchronous avatar (seen in Figure 34 D), while the experimenter introduced them to the procedure and the avatar generation system. To distinguish between the self-avatar and the other avatars, we told participants they may generate copies or clones of themselves that perform recorded movements using the VR system. Before each task, the experimenter followed the walkthrough procedure, explaining how to use the "*clones*" and describing the task, which involved the following:

- Task 1 (*Multiplayer*): Participants played two games with avatars (*Rock, Paper, Scissors* and *Patty-Cake*). After recording an avatar, participants played the games with it and then recorded a second avatar to play instead of them. Here, we probed the *Stand-In For Me* and *Others* capacity of the avatars for multiplayer games.
- Task 2 (*Arguing*): Participants simulated an argument with avatars. After recording an arguing avatar, participants practiced arguing with it and then created a second avatar to argue instead of them. Here, we probed the *Stand-In For Me* and *others* capacity of the avatars for psychological training.
- Task 3 (*Dance)*: Participants made choreographers using at least five avatars. With this task, we aimed to evaluate whether participants could use avatars to model real-world creative processes. Here, we probed the *Stand-In For others* capacity of the avatars for modeling within creative experiences.
- Task 4 (*Pipeline*): Participants used avatars to assemble boxes in a pipeline task. Here, we probed the *Stand-In For others* capacity of the avatars for increasing productivity.

After each task, the participants described how the "clones" helped them, talked about their experience concerning embodiment, perceived performance, and enjoyment by rating the following statements from validated questionnaires on a 7-point Likert scale (from -3 to 3): *Q1: I felt as if the clones were me* [200], *Q2: I felt as if the clones were someone else* [200], *Q3: I was successful in accomplishing what I was asked to do* [102], *Q4: I had a good time playing the game* [235]. In the end, we interviewed participants with open-ended questions about difficulties and improvements in using the system and the interface to support the application scenarios (*I*). We attached a video of the usability study tasks as supplementary material.

---

[10] https://developer.oculus.com/downloads/

[11] https://developer.oculus.com/blog/meta-avatars-sdk-now-available/

[12] We attached a picture containing all the available Meta avatars in the supplementary material.

5.6.4 *Participants*

We recruited 17 participants (9 female, 8 male) with a mean age of 29 (SD = 2) through convenience sampling and from the university. Two participants self-described as expert VR users (P13, P17), three as intermediate (P7, P8, P5), and ten as novice VR users. Two participants had not used VR before (P18, P9). The average duration of the study was 47 minutes (SD = 4). Participants received a gift for taking part in the study valued at 15 currency[13].

5.6.5 *Analysis and results*

To analyze the qualitative data recorded through the study, the principal researcher followed an open coding procedure [107] and grouped feedback related to application scenario usefulness and positive feedback in the tasks in Section 5.6.5.1 (Q4), avatar embodiment in Section 5.6.5.6 (*Q1*, *Q2*), usability challenges and improvements in Section 5.6.5.7 (*Q2*, *Q3*, *I1*, *I2*, *I3*).

5.6.5.1 *Were the avatars useful for application scenarios?*

5.6.5.2 *Multiplayer*

For the multiplayer games, some found it harder to synchronize the avatars (P13 P17, P2, P14, P16, P18, P12), and one participant mentioned the handle was difficult to manage (P16). They had mixed experiences, some finding it *"weird"* (P15), others *"interesting"* but not *"super fun"* (P7). Concerning the recorded part, P5 mentioned it felt unfair to control the outcomes of the games (P5), since *"you already know what you're gonna play"* (P4). P10 liked the *"embodied part of it"* and *"practicing it"* as they did it.

5.6.5.3 *Arguing*

The arguing task invited mixed responses, mostly due to the artificial nature of the task (P6: *"when you argue you have a reason"*) and lack of audio and facial expressions. Experiences using avatars varied, from being *"fun"* (P12), and *"funny"* (P7), to *"uncomfortable"* (P5), *"strange"* (P17), and *"interesting"* (P5, P10, P13). A few remarked that *"it's fun because they [do] not necessarily need to be synchronized"* (P2). Despite this, some participants found it *"more intense"* (P4, P11), since the avatars expressed a *"kind of aggression"* (P11).

5.6.5.4 *Dance*

The choreography-making tasks received the most positive feedback. Avatars were perceived as useful, particularly for seeing "how movements look like when you put them together" (P1), and were regarded as a "great visual tool" (P3) for this type of creative task. P13 mentioned having difficulties syncing the crowd with a singer avatar. Here, participants would have preferred having full-body tracking and modeling leg movement as well (P1, P4, P9, P7, P10).

---

[13] anonymized for review

5.6.5.5 *Pipeline*

On average, participants moved 30 objects through the pipeline (SD = 22, Min= 5, Max = 77). Some participants found avatars useful for pressing buttons to generate objects (P10, P15, P14), since they could perform the other complex tasks (P13: *"I could just handle the pieces into the boxes"*). However, this task generated synchronizing problems both technically and conceptually (P17: *"It was more [...] thinking about how they should sync in the world"*), which led to mixed experiences. For example, P18 realized that the avatars failed to grab objects since the objects had different initial positions from the recording. Syncing avatars among each other and with other objects in the environment added extra mental load (P18, P16, P13, P2). Others had different expectations of the avatars. For example, P14 expected the avatars to adapt *"to the movement of the boxes"*. Lastly, during this task, some faced some clutter issues (P9, P2: *"there were too many of them"*).

5.6.5.6 *How are avatars perceived?*

Despite mixed embodiment ratings, we identified several themes in how participants related to the avatars through the qualitative data. The avatars were seen as *"recordings"* (P10), *"representations"* (P11, P12, P13), *"things"* (P5), *"workers"* (P17) or *"programs"* (P13). Most often, participants mentioned movement, appearance, and the nature of the task as embodiment mediating factors. In addition, the avatars' perceived personality and quantity seemed to affect relatedness to some degree. Lastly, a few participants mentioned the avatar's agency, performance, and their own personal beliefs as mediating factors.

The avatar's lack of expressiveness put off some participants (P6: *"The face kind of ruined the experience"*). Participants would have preferred more *"personalization"* (P14) to fully identify with the avatars (P15: *"it's just not a picture of me [...] like the movements are me"*), in addition to leg tracking (P3, P6, P11, P16) and voice (P11). For the *arguing*, and sometimes *dancing*, participants role-played during the tasks (P4 [14], P18 [15]), *"exaggerating"* their hand movements (P7) to accomplish the task, however, this conflicted with their personality, thus they felt a *"disconnect"*(P10), since they did not see themselves *"as somebody who is like, kind of aggressive"* (P13). Quantity also influenced how participants related to avatars (P12: *"there were so many, it's like, it's not me, all of them"*). By not responding to the user or adapting to objects, the avatars failed to meet the users' expectations of performance, which may have decreased enjoyment and relatedness (P2: *"I feel that they're stupid"*, P4: *"They weren't doing what I was telling them to do"*). P5 mentioned how their own beliefs prevented them from relating to the avatars: *"it's giving a lot of value to the clone [avatar], which I don't agree with"*.

5.6.5.7 *What usability challenges did participants face?*

Participants were negatively affected by hand-tracking limitations during the study. When the hands lost tracking, the avatars' mesh was still connected to the hands, causing the avatars to look *"creepy"*, or *"broken"* (P12).

---

[14] P4: *"Maybe the second one would be me [...] I was like the guy, who wanted to defend himself""*

[15] P18: *"because I'm asked to like, take upon a role to feel a certain way, which I actually don't"*

Participants suggested creating an idle pose during the loss of tracking to improve the avatar's appearance. Besides this, there were other technical limitations, such as issues with grasping, which affected grabbing the menu (P9), and approximating depth, which affected pressing buttons (P8).

Synchronizing avatars was the most prominent usability issue during the study. To improve this, participants suggested different ways to increase playback control: by playing the recordings within a fixed time frame (P16: *"I want this action to take 20 seconds"*), adding a slider in the control panel to control the playback timeline (P17), displaying a *" numerical value as if it was a frame-based animation"* instead of the ID (P3), and triggering avatars based on each other (P3). Some syncing issues were also caused by recording unwanted actions (P5: *"I would also record my movement, like walking to the button"*). Participants further suggested selecting avatars by touch instead of the control menu (P12), setting a countdown before spawning avatars (P16, P18) to help with *"getting a position"* (P11), or using voice control for generating avatars (P16, P13).

In addition to syncing, clutter made identifying the avatars difficult during the pipeline task.

First, visually linking avatars to their IDs was problematic for some due to overlap and movement (P2: *"they're all moving, it's hard to understand which number is assigned to which'*). Second, recording too many avatars introduced problems with recall for others (P13: *"I don't even remember the number he has"*). Lastly, the avatars sometimes obstructed each other and the environment, causing issues in handling or identifying them. While moving the avatars was useful, participants sometimes struggled to find the handle (P16: *"because they [referring to the avatar] were big, and I was like - where is the handle?"*). To improve crowd visualization, participants suggested making avatars transparent (P10) or completely invisible while being moved (P14). Some other suggestions included adding legs (P2, P6, P17), facial expressions (P6, P17), appearance editing features (P2), and permanently enabling the copy-cat avatar (P6).

5.7 DISCUSSION

This work has generated a resource for describing and sharing ways to use multiple avatars in virtual reality. We have adopted an avatar-centered lens when investigating multi-avatar interactions. This has allowed us to discover novel exciting opportunities, like blending movements within the same avatars and varying degrees of control. Some related works adopt a spatial lens when discussing interactions with multiple avatars, referring to proxemics, like interactions at a distance [291, 277], a notable exception here being Xia et al. [211] and Slater et al. [230]. While this spatial lens helps identify use cases to perceive oneself or interact from different places, an avatar-centered lens reveals the interfaces embedded in the avatars that enable these application scenarios. Next, we discuss recommendations for developing an interface and a system to act through multiple avatars based on our findings.

5.7.1 *Recommendations for integrating multiple avatars in VR applications*

Based on the feedback received during the usability study, we recommend some interface improvements to practitioners wishing to implement interactions with multiple avatars. Participants mentioned syncing as the most difficult practical task with the avatars. While adding more operations to facilitate timeline scrubbing may result in easier syncing, we believe using the control menu will not remove syncing as a conceptual problem. Some participants seemed to have adaptive expectations of the avatars and regardless, having objects in the exact same state for pre-recorded interactions may not be possible in messy environments. Therefore, we recommend building an implicit and visual syncing mechanism into the system, similar to the concept of *causality* [158], which already considers system context in its model. *Asynchronous Reality* describes a mechanism to keep track of actions and outcomes in the virtual space [282]. Using this kind of mechanism, users may implicitly connect avatars with each other or with the environment before recording to generate a causal graph of events.

Recent work exploring how users can embody two robot arms and play pin-pong in VR changes viewpoints automatically for the correct player [293], similar to switching between avatars from *FreudTherapy* [230]. Takada et al. [293] found that switching can cause loss of ownership when the switch target does not reflect the user's pose after switching. When in therapy with Freud though, users do not interact with objects in the environment since they only visualize themselves in a mirror. When switching between avatars that perform interactions though, we recommend allowing the user to take explicit rather than automatic control of the avatar to prevent disturbing any ongoing task. Furthermore, measures to prevent motion sickness may be needed. Based on the workshop and the usability study feedback designers may consider integrating voice control in the interface, similar with Liu et al. [246]. This could prevent the users from recording unwanted actions, like walking to and from the control menu. Indeed, during development, we had to remove the avatar's access to the control menu to prevent Midas Touch problems [95]. Therefore, completely shared contexts may be impractical unless the designer wishes to enable avatars to recursively generate themselves. Similarly, we removed the user's ability to move the avatars during the pipeline task. Moving the avatars with the handle provides inadequate support for minute tasks that involve interacting with objects; therefore developing an embodied way to edit the recordings to change the movements is recommended.

Apart from syncing, many avatars may cause visual clutter, as discussed with the experts. In practice, participants were sometimes overwhelmed by the avatars and had issues identifying them and recalling which avatars performed which motions. Humans are known to have limited multi-tasking capabilities [32] since tracking distinctly moving targets requires attention resources which are modulated by speed and proximity [99] with limits from three to five objects [21]. When looking at crowds of objects, people tend to focus on locations adjacent to these objects to represent them as a group [92]. When dealing with crowds, the crucial design requirement is to leverage selective attention and preattentive processing [242]. This means designers should minimize the mental load from processing any additional information besides what is needed. Indeed, for mixed reality contexts, finding a balance between *presented* information and *enough* information is key to preventing users from being overwhelmed [287]. While increasing the play area could help, however, VR outside the lab is limited by household objects and space [263]. Thus, managing avatar transparency, using fade-in and filtering mechanisms, and detecting user intent for displaying information could improve the handling of crowds. The design space offers opportunities to remove clutter by previewing partial avatars, a feature implemented in *AvatAR* [289]. Instead of displaying more avatars, the user may *blend* multiple inputs in the same avatars to reduce clutter.

Applications like *Mini-Me* Figure 35 C capture real-world use cases of *Blending* marking the descriptive power of the design space. More recently, Freiwald et al. [283] develop *Smart Avatar* visualizations that follow the user but switch to a different representation when users teleport, which facilitates social presence and increases *spacial awareness* in shared environments. Blending motion is used in the field of animation to interpolate between two motions of humanoid avatars, which has been used to replace the self-avatar movement, for example, to facilitate training and social inclusion when lacking input [229] or increase social presence [207]. Researchers have investigated similar concepts of playing rock-paper-scissors with oneself using EMS technology [274].

To implement multi-user avatars, the designer may need to develop custom avatar queuing systems that might involve priority. So, therefore, a user may queue to control an avatar or may instantly assume control of it if their input is prioritized. Schjerlund, Hornbæk, and Bergström [291]'s *OVRlap* and, more recently, Hoppe et al. [284]'s work exploring the perspective continuum for users who control an avatar from multiple views reveal possible design considerations for interacting with multiple avatars. While we conceptualize these designs, we leave their implementation and evaluation as future work. In this work, when we say users embody avatars, the avatar is their self-avatar, according to Gonzalez-Franco and Peck [200] definition. Users may control avatars from the third person that is located in the same place as their self-avatar, giving the impression of a first-person perspective. When users switch between avatars, they continuously embody one avatar, with the exception of multiplex avatars, which allow integrating multiple contexts at the same time and may describe multiple self-avatars.

5.7.2 *Using multiple avatars in mixed reality*

Multiple avatars may be used in similar ways in mixed and extended reality (MR/XR), expanding the focus of our work on virtual reality. We show selected works in Figure 35 where multiple avatars are used as recordings (Figure 35 A), for tutorials (Figure 35 B), to increase social presence (Figure 35 C) and for information visualization (Figure 35 D), which may facilitate facilitate experts in various fields who may lack technical skills [248].

Figure 35: This figure shows selected works from AR that use multiple avatars to record events enabling deep work A [282], record instructional tutorials in B [214], increase social presence with a *mini*-avatar in C [206], visualize immersive recordings of body movement D [286].

Developing interfaces that enable the generation of avatars without technical know-how is crucial to enable trainers to record their own tutorials or movements, which users may later train on. While we can map these interactions in the design space (as seen in Section A.3), the dimension of *Context* becomes problematic since XR includes the real-world context or other interfaces like tablets. In general, avatars are a substitute for actions that some related work displays by moving objects [241]. Humanoid representations of actions reveal the user's behavior in context and may increase social presence in some contexts, like in *Mini-Me* [206] (Figure 35 C), where a mini AR avatar follows the user's movements and blends their transform with the user's gaze, and in *AvatAR* where the user may visualize movement data. In *Mirror Fugue* [121], users may remotely join piano sessions together, either by mirroring the remote user's piano, offsetting it 90 degrees, or showing a pair of shadow hands on the same piano and additionally, they may play alongside recordings that could be used for learning and dueting. More recently, Grønbæk et al. [298] developed the concept of *Partially Blended Realities* for collaborative scenarios and defined how to define control explicitly and spatially.

5.7.3 *Appearance beyond multiple avatars*

While the usefulness of avatars for motion-related applications is clear, we discussed how appearance plays a role in generating different use cases within application scenarios. Avatars enable linking actions to expected outcomes, which is a key principle for *embodied learning* in immersive environments [142]. Some of the positive effects of realistic avatars are well-documented [112], from increasing the feeling of body ownership [143] to aiding in enhancing exercise [104]. More recently, Fitton et al. [297] found that allowing trainees to customize the trainers' avatars for dancing has the potential to increase the efficacy of the training.

Similar work shows that when users see themselves performing exergames in VR, it improves performance and motivation much better than when competing with an opponent [197]. A non-technical interface that enables controlling *virtual doppelgangers* [112] may allow researchers to investigate participants' choices and not just their reactions. Indeed, learning seems to be one of the most feasible use cases for multiple avatars since it could provide immersed users with a different type of *vicarious learning* [105]. As suggested by social learning theory, observing and imitating human behavior is a common learning tool [9]. Live streaming video games has evolved into an entire industry with pro-gamers, spectators, publishers, and sponsors [130]. These streams are also used for learning purposes and allow novices to teach as efficiently as experts [191]. Moreover, psychology research on *virtual representations of self* and *other* indicate positive outcomes for avatars that resemble users faithfully, particularly in appearance. For example, users may exercise more on a treadmill when seeing themselves instead of another [104].

However, increasing ownership over the avatars could negatively impact user experience and may not always be needed. For example, autonomous agents that reflect distinctive human traits may evoke the uncanny valley [194]. Since VR already poses *"cognitive, emotional and behavioral disturbances"* [**slater_2020**], further closing the gap between a user's avatar and their self-avatar may pose unique ethical challenges. More recently, Lee et al. [302] ran a study where 20 participants discussed the use of *AI Clones* in

8 scenarios. They found similar ethical challenges relating to how users may abuse *clones* by misrepresenting them or forming unsuitable relationships. For a more detailed overview of the ethical implications of using clones of the user as avatars, we refer to their work. In addition to the psychological disturbances, if the avatar resembles the user themselves, this exposes the users to privacy risks. Ens et al. [264] discuss how to ethically approach the usage of motion data and contextual data related to spatial visualizations in their work on *Grand Challenges in Immersive Analytics*. The researchers recognize that *"intimate knowledge"* may be revealed through this type of information, and therefore, designers should establish an ethical procedure for data handling and storage.

5.7.4 *Experts as a source of knowledge*

As Eriksson et al. [265] point out design spaces fall into what Höök and Löwgren [129] consider *"generative intermediate-level knowledge"*. For this related work, the *"core idea"*, which spans use cases and application domains, represents the concept of interacting through different avatars. This concept may be realized through different interfaces, one of which we provided with this research through the usability study. The design space for interacting through multiple avatars serves as a tool that helps the designer move from abstractions toward instances of interaction [129]. However, the *"lightweight nature"* of design spaces diminishes the designer's ability to retrieve relevant contextual usage from these instances [173]. For this, *design-space thinking* [269] methods include a process known as *design space filtering*, first introduced by Lim, Stolterman, and Tenenberg [97], whereby prototypes are developed to investigate particular dimensions of design.

It is common to involve experts in the design process of VR and XR systems both for the knowledge-building process and the filtering part. Experts may be used as a requirements gathering tool, for example, through interviews [237]. Such requirements may also be derived from analyzing literature as in Jetter et al. [241]. Design spaces are also commonly built on related work — a process that involves building a Zwicky box following a morphological analysis procedure [220, 237] or through open coding [261]. For the information-gathering process, experts may also be involved as a source. For example, in Márquez Segura et al. [272]'s work, 10 experts are paired up to play exergames and participate in follow-up interviews about their experience. Based on this study, the authors then develop a design space schema following a *Research through Design* process [128]. Designers may also be involved at a later stage, during the design space filtering process, where walkthrough-usability studies determine changes in the design space [305]. Eriksson et al. [265] generate a design space based on the implementing *4in1 games*, a process involving 50 developers. We have added affinity diagram-making in this process to formalize the data analysis from interviews and gain a broader understanding of contextual usages by looking at scenarios. This led us to the design space. In practice, we have differentiated between the *"academic"* view on multiple avatar interaction and the practical aspect of calling them *"clones"* to not overload participants, who might not even know what an avatar is. Many related works use avatar copies for interacting with multiple avatars [291, 211] since they are accessible and already preloaded in the application.

5.7.5 *Limitations*

We acknowledge some limitations in the implementation of the interface to generate multiple avatars in the usability study. First, the hand-tracking and the chosen library for representing avatars affected performance and did not support leg-tracking at the time of the study. Despite this, we chose the expressiveness of hand-tracking to enable integration with other XR mediums. Moreover, by focusing on hand-tracking, we create an opportunity to extend the concept beyond VR towards XR. The breadth-first approach of this research also has some limitations. While rich in information, the application scenarios and design space are not exhaustive. The formative workshop provided many unique ideas and discussion points, but the resulting interaction landscape and design space is mediated by the personal experiences and expertise of the participant and the interpreter.

With respect to the usability study, while we selected one of the most common use cases we did not evaluate all possible combinations that the design space may yield. Some of these combinations can be found in related work, so we focused on shared contexts for asynchronous avatars. While this allowed us to investigate the feasibility and usefulness of particular instances of the design space, other types of scenarios that were not included in the evaluation could yield different considerations.

5.8 CONCLUSION

This paper offers a comprehensive account of *why* users may interact with multiple avatars in VR by developing a design space for acting through multiple avatars, from theory to practice. Based on formative workshops with 12 VR experts, it introduces a design space that maps out four dimensions describing multiple avatar interactions in VR. The design space can be used to describe and evaluate existing works, including mixed reality. Furthermore, it can help researchers and designers generate interactions for multiple avatar use cases. Furthermore, we discuss system requirements and operations that are essential for realizing the *why* and *how* of acting through multiple avatars in VR. Instantiating the design space, we implemented a prototype covering four application scenarios and 14 basic operations to evaluate use cases in a usability study with 17 participants. This conveys the descriptive, evaluative, and generative powers of the design space that enable practitioners to uncover key application scenarios, such as using avatars as stand-ins, and gaps in related work, like blending control from the system and the user within one avatar. Our findings also shed light on the challenges of using multiple avatars to act in VR, such as syncing and clutter, and provide design recommendations to address them.

6

DISCUSSION

> *So, treating nonexistent things as if they're real?*
> *"Nonsense" probably isn't the right word for it.*
> *Being able to create and conceptualize a universe is a pretty amazing skill.*
>
> — Hank Green, Crash Course Philosophy, Episode 29

Chapter 2 has brought together knowledge from various fields to unfold a path from *concepts* to *future interfaces* through *design spaces* and their tangible counterpart —*prototyping*. Since concepts are the *building blocks* of thinking, they represent a way for researchers and practitioners to reason about design and abstract it to core parts. Concepts provide two crucial elements: *"what is being represented"* and *"how that information is typically used"*[67].

These elements are reflected in HCI through prototypes and design spaces. The prototype and exemplar views of concepts from cognitive psychology inform — the exemplars from bridging concepts [119, 152], — and reflect prototypes as tangible manifestations of HCI concepts. Prototypes represent a way for designers to reason about what the concept might do in the real world. While prototyping makes visible parts of this reasoning, what it *"leaves open is subject to more discussion and design space exploration"*[86]. This section takes an overarching view of conceptual-driven design as an approach to developing future VR interactions and interfaces. In the following, I present an account of the implications of the research presented in this thesis, together with reflections on the chosen methods.

6.1 AN APPROACH TO MANIFEST FUTURE INTERFACES FROM BUILDING BLOCKS

The general approach of this thesis has been a concept-driven approach. This approach borrows principles from:

1. Stolterman and Wiberg [119]'s concept-driven design rooted in future visions, grounded in theory and manifesting ideas in concrete, tangible designs;
2. Höök and Löwgren [129] view of intermediate-knowledge like strong concepts, which carry core ideas cutting across use cases and application domains;
3. Dalsgaard and Dindler [152]'s bridging concepts, which *span the gap between theory and practice* to reveal untried opportunities and account for exemplars and parameters that shape the concepts;
4. Beaudouin-Lafon and Mackay [86]'s view of prototyping as an idea generation exercise, which manifests visions for evaluation and leaves open design space exploration.
5. Beaudouin-Lafon et al.'s view of generative interactions as grounded in theory, actionable, and showcasing the descriptive, comparative, and generative powers of interaction through principles and tools [52, 70, 257];

So — what should designers do to innovate? How can they follow the path of conceptual-driven design in VR — and manifest visions as tangible designs, through an explorative nature, grounded in theory to inform future designs. The approach presented in this thesis centers *interaction models* — as grounded in theory, giving dimensions to *design spaces* — as a systematic and explorative approach to concept-driven interaction and prototyping — as tangible manifestations. In this approach, design spaces:

1. represent tools for generating interactions;
2. capture the interaction model based on HCI theory through its dimensions;
3. reveal the boundaries of interactions as parameters and values;
4. enable prototyping by revealing principles and parameters to consider;
5. reflect the prototype (or exemplar) as points or areas within the space;
6. record the design process.

The path *from theory to artifact* and *artifact to theory* is formalized through the design space-making process. This process results in an interaction model nested within a design space, which, upon evaluated empirically as a tool for design and as a manifested tangible artifact, serves to further inform the concept development. Figure 36 shows this approach.

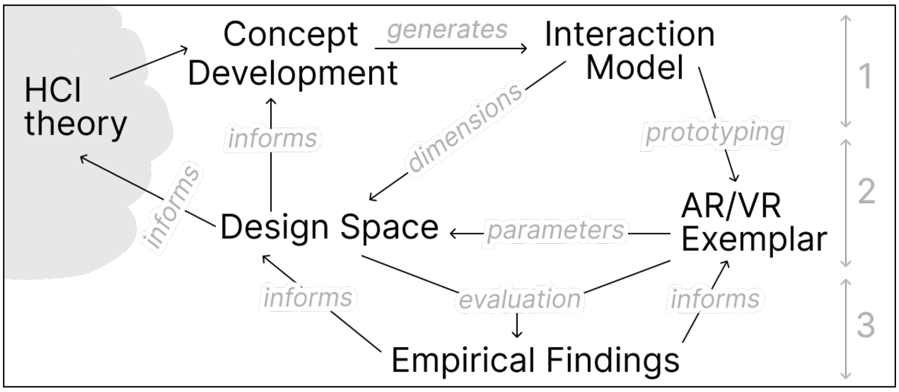

Figure 36: This figure depicts the approach to concept-driven design in VR in three phases: 1 — concept and interaction model development, 2 — prototyping VR/AR examplars that capture salient aspects of the model, determining the parameters and making the design space, 3 — gathering empirical data about the design space and the exemplars, which in turn inform the concept development.

Related design spaces in VR generate dimensions and parameters via literature reviews, which categorize related work through various lenses and reveal gaps as novel design opportunities [220, 261]. These design spaces classify related work according to some scope rather than present a model of generative interactions that may generate similar related work. For example, Hirzle et al. [220] derived dimensions and parameters from a literature review as technical qualities of VR headsets and human depth perception (e.g., monoscopic, stereoscopic, vergence, accommodation). The researchers used the design space to present related work from three perspectives: technology-based, application-based, and interaction-based. Danyluk et al. [261] reviewed VR works containing world-in-miniature techniques and derived a design space through open coding. The parameters of this design space captured observed empirical and technical qualities of the works (e.g., size, geometry, links, virtuality, etc.). This work suggested novel conceptual ways to implement world-in-miniature techniques by filling in combinations of parameters not captured by previous works. Drey et al. [237] generated a design space based on qualitative interviews and technical requirements with dimensions and parameters that abstracted 3D and 2D midair sketching in hardware (pent, tablet, etc.) and workflow for sketching (operations, types of objects, direct, indirect input types, etc.). The researchers proposed five metaphor groups (primitive extrusion, portal into space, etc.) based on the design space parameters to generate interaction ideas, categorize existing works, or reveal gaps. They also presented a prototype instantiating some of these groups.

The parameters and dimensions within VR design spaces often result from observed qualities of related works. The design space schemas do not necessarily abstract interaction models but provide various lenses or interpretations of the classified research. Instead, this thesis suggests that design space dimensions may abstract interaction models to serve as a design tool for generative interactions. The dimensions may reflect interaction models at a higher level of abstraction. The parameters may give a less abstract view of these models and give practical or technical options for implementing interactions in practice. Design spaces could nest a conceptual framework for thinking about interactions and a practical framework that reveals possibilities for implementing it in practice.

Presenting design spaces as interaction models has crucial advantages for design — they prepare researchers and practitioners for the stream of technologies. Lowgren and Stolterman [88] identified several principles of designing thoughtfully, one of which was *being prepared*. Crucially, being prepared involves discerning a trend from a paradigm-generating work. To navigate an increasingly evolving technological landscape, the researchers proposed conceptualizing ideas as variations of already familiar ideas. This would involve accumulating design exemplars and developing a vocabulary that describes and analyzes them. The conceptual-driven approach to interaction design accounts for the exemplars (prototypes) and the vocabulary required for researchers to understand the stream of technology (design spaces).

Design spaces make visible the conceptual —- they are a physical, semantic manifestation of potential designs and enable the exploration of these designs. Design spaces semantically describe interaction models and give a systematic account of options available to design prototypes. This principled overview would allow designers to identify which are the most salient characteristics of an interaction model. In the *"rapidly evolving technological landscape"* [180] of HCI, situating research and conveying the parameters it explores quickly could be a valuable tool for researchers and practitioners alike. This thesis captures research with concepts at the intersection of theoretical, empirical, and practical.

6.2 IMPLICATIONS FOR CONCEPT-DRIVEN DESIGN IN VR

In what follows, I describe the implications of each work and reflect on their importance within the VR landscape as — (1) *practical real-world* concerns for practitioners or social communities, and — (2) *methodological* concerns for researchers who might want to approach innovation as conceptual interaction design in VR. I also highlight the *future visions* that helped manifest the contributions of each work. In *The Social Life of Innovation*, Denning [71] posits that innovation socially transforms communities differently from inventions. Conceptually driven interaction design aims to manifest paradigm-shifting ideas into tangible prototypes. However, few concepts or prototypes reach the seminal status of *Dynabook* or *Sketchpad*. When their seminal status emerges, the community's perception of these inventions does not necessarily align with their creators' view [84]. Some upcoming observations give a 'social' account of how these concepts may shift perceptions.

6.2.1 *VRFails*

After analyzing and categorizing the YouTube videos from VRFails, it became clear that not all of the videos captured failures. Rather, the meaning of *fails* was formed by the YouTube community. Some examples were clear instances of breakdowns in interactions, while others captured joyful moments with friends. The most immediate outcome of using VR fails for qualitative analysis was, in my opinion, the press[1]. Without a doubt, this research generated mass media interest because it had a good hook and was, above all else, fun. **(I1)** It follows that having a good way to abstract research into a fun, simple concept has great potential for science communication.

Implication 1 — practical

Adopting the HCI concept of *seamful design* allowed us to take these two different perceptual lenses of VR fails as breakdowns or interaction opportunities. *Seamfulness*, the concept that breakdowns are a resource rather than a failure, was introduced in the context of wearable and ubiquitous interaction and stands in contrast to *seamlessness*, the view that computers and interactions with computers should aim towards being invisible [63]. By taking the lens of seamful design, we suggested interactions that were aimed at preventing fails and others, which enabled the positive shared experiences of VR. **(I2)** The implication here is that leveraging HCI concepts from different domains, like seamfulness, may determine novel insights for interaction design.

Implication 2 — method

The paper puts forward design implications sketched as interaction techniques grounded in specific scenarios (e.g., covering the headset to black out the view in the context of a VR horror game). The design implications of this work are relevant for designers wishing to design for VR use at home. The interactions and scenarios from this work research could help decrease injuries or damages as a result of using VR in the home and help further develop VR as a shared experience between users and spectators.

Future vision

The 'social lens' is revealed within the concept of *fail* itself, which shows how social platforms may appropriate and shape ideas that, upon further investigation, have a different meaning than initially thought. Since VR fails is seen from a social lens, its meaning is not static in time and will develop as the technology is shaped by the public and by the media. This means that a future VR fails classification might yield different results and implications, depending on their cultural, social, and technological context.

VRFails approached the generation of the VR fails concept systematically, grounding it in HCI theory through its use of seamfulness to manifest interaction designs that may inform VR designed for at-home use. This paper

Conceptual-driven interaction design approach

---

[1] I gave interviews and appeared on national television — for kids!

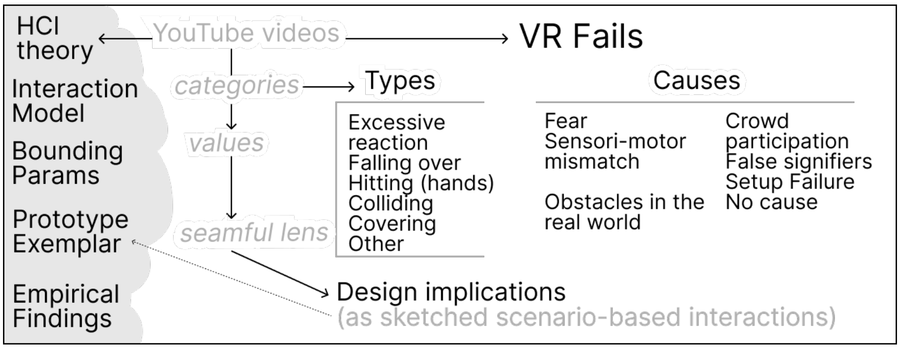

Figure 37: This figure depicts the classification of VR fails derived from quantitative content analysis and how it may fit in a concept-driven approach.

starts to develop the conceptual-driven approach to interaction design by first (1) generating the VR fails concept, (2) giving a systematic account of VR fails, and (3) proposing ways of practically applying it in the real-world. Figure 37 shows how VRFails may fit in the conceptual approach to VR interaction. This research does not fit the approach precisely as we have formalized it in terms of design spaces and interaction models. The next steps here, following the concept-driven design approach, would be (1) deriving an interaction model for VR-fails interactions, (2) making a design space based on those, and (3) manifesting a prototype based on the design space and key future use scenarios.

6.2.2 *Feedforward*

This work started at the same time with VRFails in 2020, from a concept my colleagues and I were going back and forth with, namely *embodied feedforward*. The idea was to create an interaction where the user's avatar would show the user what they can do. This work was a key feature of my ascent into design space thinking as a methodology. While an interesting concept, I could not find anything remotely similar to the concept of embodied feedforward. A notable exception here is Lopes et al.'s *Affordance++* [164].

Or descent, depending on the reader.

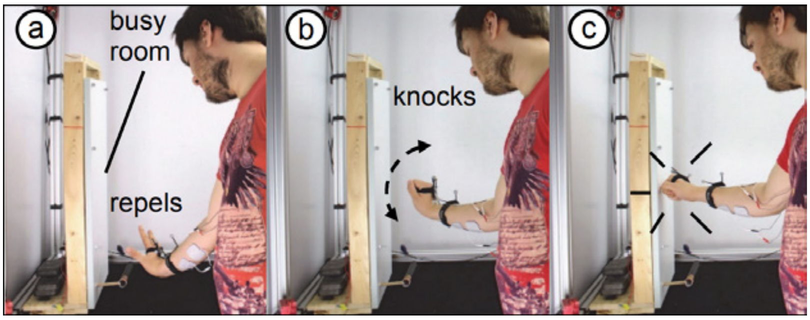

Figure 38: This figure represents *Figure 8* from Affordance++ [164], where the door rejects the user's knock in *a*, and encourages the user to knock once the room is no longer occupied in *b* and *c*.

The notion of feedforward in VR reflected concepts from Affordance++, so much so that parts of the challenges for VR interactivity, or interactivity

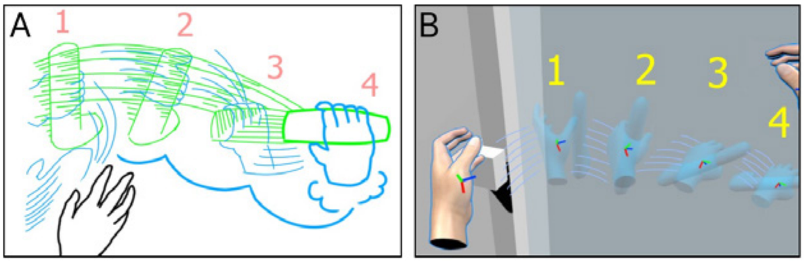

Figure 39: Subfigure A shows an earlier feedforward sketch draft, and B shows an earlier implementation of feedforward applied to a Norman door.

in general, were inspired by it in Section 4.5.4. We began the morphological analysis process for the design space based on one of those challenges. We chose the Norman door interaction as the best way to convey the concept of feedforward visually (also considered in Affordance++, Figure 38). Feedforward's teaser figure and implementation examples involved misleading Norman doors[2]. Figure 39 shows an early concept of Feedforward.

This project heavily relied on prototyping, both as sketching, through design spaces and as VR prototypes. Drawing out the theoretical concepts with the tangible implementation required much thoughtful design. Design spaces tell a story of how prototypes happen because they reveal the boundaries of design and the concepts underlying it. While prototyping instantiates parameters of the design space within one artifact. Feedforward though revealed so many parameters through prototyping that solely relying on one instance seemed inadequate. Therefore, I developed the feedforward prototype as an authoring tool that showed areas of the design space instead of values. The authoring tool enabled us to get feedback for more designs and more scenarios from the experts during the evaluation. **(I3)** The implication here is that prototyping as a process allows for concept exploration in practice while authoring tools enable design space exploration in practice. Authoring tools reveal areas of design spaces, not just points.

Implication 3 - practical

The conception of the feedforward design space involved a *participatory* 'social' aspect through expert evaluations. These evaluations changed the parameters of the design space (the schema) and revealed ways to improve the design space cheat sheet (the schema and the definition of parameters). For example, the experts suggested changing the order in which parameters they were presented to be more in line with their experiences designing for VR. **(I4)** The implication here is that developing design spaces as a tool means changing them to reflect the design process. As a tool in the real world, the feedforward design space should convey the principles of the models underlying the design space and aid in the interaction generation process. Therefore, the feedforward design space teaches its 'users' about what feedforward may do and how to design for it. This can help VR developers better design tutorials or training applications by considering the generative model of the interactions, which shows people what to do and how to do it. Instantiating feedforward as a design space helps bridge the gap between the HCI academic view of this interaction model and its practical applications.

Implication –4 - practical

Future visions

---

[2] Norman doors give misleading signals about how they may open. The reader might recall unfortunate door encounters themselves.

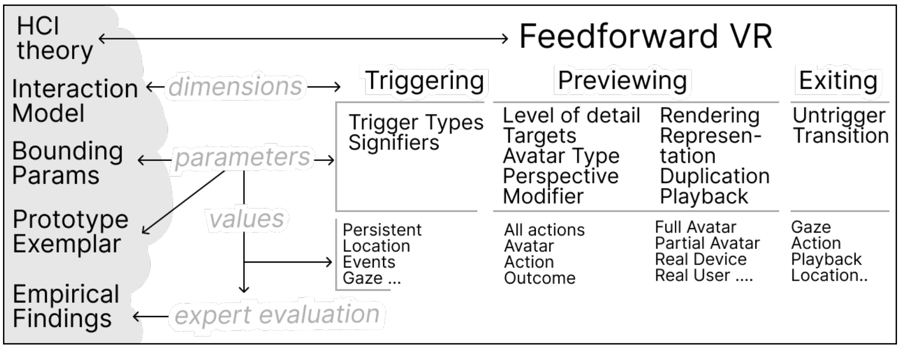

Figure 40: This figure depicts part of the feedforward design space mapped to the concept-driven approach principles.

Conceptual-driven interaction design approach

This work developed the concept for feedforward in VR as a model for generative interactions drawing from HCI theory. This represents the first work taking a concept-driven approach to VR interaction design through the design-space-making process. Figure 40 depicts Feedforward as a concept-driven approach. In doing so, it developed a design space to systematically explore the boundaries of an interaction model and manifest a tangible prototype that captures salient aspects of the model — the feedforward authoring prototype. Specifically, Feedforward proposed *triggering*, *previewing*, and *exiting* as a model for feedforward in VR, and through a morphological analysis process and prototyping, derived further parameters and values bounding these dimensions.

Experts evaluated the design space and prototype in two ways: first, as a semantic tool for concept generation, and second, as a prototype authoring various feedforward designs in three application scenarios (car, door, and kitchen). The experts helped develop novel parameters to design feedforward as a reflection of their real-world design process. They also helped improve the comprehension of the design space cheat sheet. The contributions of this research inform the design of tutorial or training VR applications in the real world. Furthermore, this work provides a method for developing design spaces as a tool for design and as generative interaction models grounded in expert evaluation and HCI theory.

6.2.3 *MultipleAvatars*

While the core idea of this paper was interacting through different avatars, this project began with the concept of cloning as an interaction technique. Later, we focused on the bigger picture — acting through other avatars. The semantic space of multi-avatar interactions includes several related concepts: copy-paste interactions, cloned interactions, motion-captured interactions, and recorded interactions. This conceptual space (Figure 41) helped me find related work I would have missed otherwise. **(I5)** The implication here is that when literature review searches yield few results, researchers may want to explore this conceptual space. Other related work might think of interactions using other metaphors than those used by the researchers. A systematic account of novel interactions given as design spaces may also help find related works instantiating novel interactions or interfaces in VR. Design spaces may reveal similar interactions described using different metaphors.

Implication 5 - method

A 'social' aspect of this work involves its cloning concept, which has a life of its own. It had a certain philosophical depth, and it posed questions such as — *what is a VR clone? where does the user end and the clone begin?* These questions, albeit interesting, were outside of the scope of MultipleAvatars. Using this term was conducive to running studies as it drew in participants and helped them verbalize their thoughts without using academic terms like 'self-avatar.' However, the crucial observation from the usability study was that some people verbalized understanding the concept, but in practice, they appeared not to understand the technique precisely. The avatars were *less intelligent* than some participants had expected. The avatars were not enhanced with artificial intelligence capabilities after all. Perhaps using the term *clone* in practice set some high expectations about autonomy. **(I6)** The implication here is that concepts have underlying meanings that might affect the expectations of people using VR technologies that leverage such concepts. So, when choosing a metaphor or concept to describe an interaction technique, the researcher should consider its implications and the expectations it might give rise to.

Implication 6 - practical

For me, MultipleAvatars was also an effort to improve prototyping feedforward interactions. The authoring tool produced by Feedforward required much too high technical skills. Making feedforward more approachable to non-skilled designers means developing a user-friendly interface to record interaction easily. Therefore, I approached this problem conceptually by abstracting it to its core parts — which was how to record or *clone* interactions. This revealed the multi-avatar concept that captured other practical applications and novel interaction techniques suited for training scenarios (e.g., blending controls of avatars for training). **(I7)** The implication here is that approaching problems from a conceptual lens might change the understanding of the problem itself (similar to *wicked problems*) and uncover new problems and solutions.

Future visions

Implication 7- method

MultipleAvatars developed the concept of acting through multiple avatars in VR based on expert brainstormings, proposed an interface that enacts concepts from the design space, and presents findings from a usability study that inform the design of future multi-avatar interfaces. This work assumes a concept-driven approach of interaction design as shown in Figure 42. This work relies on the theory following Beaudouin-Lafon [70]'s generative in-

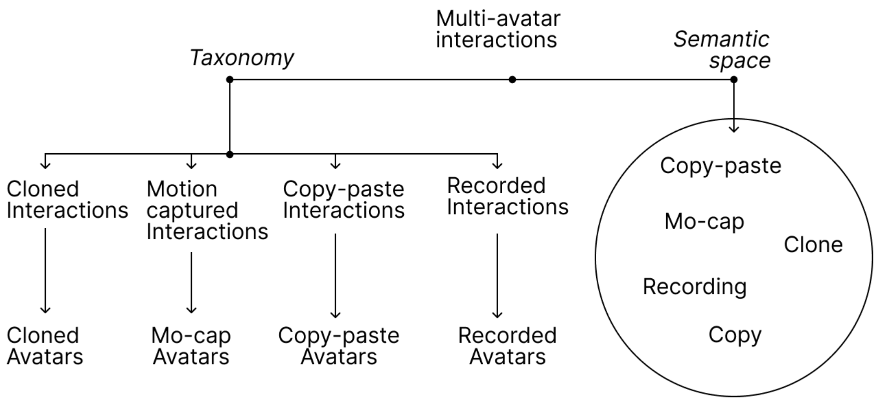

Figure 41: The taxonomy of concepts related to interacting through multiple avatars (left), and the semantic space of the concept of cloning (right, circled).

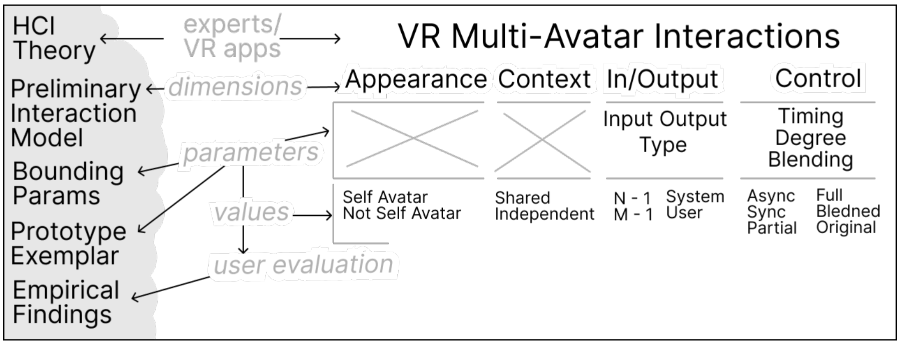

Figure 42: This figure depicts part of the multiple avatar interaction design space mapped to the concept-driven approach principles.

teraction principles. MultipleAvatars *indirectly* relies on HCI theory through the researcher experts who took part in producing the data that generated the design space. The aim of the formative workshops was to capture a systematic account of *why* users would want to interact through multiple avatars, not just how. This element of *why* appears in the design space as the *appearance* dimension, which defines two broad ways of using avatars. This helped embed some context and technical aspects in the design space and develop the concept.

This work produces a preliminary interaction model as dimensions. To validate this model, it requires a theory-triangulation process similar to the steps undertaken by Feedforward. The triangulation process helps ground the model in HCI theory. The next step involves performing a literature review starting from the related works that MultipleAvatars maps in the design space. The literature review would produce works that fit the multi-avatar interaction concept thus far and possibly extend the search to AR.

6.3 DESIGN SPACE REPRESENTATIONS AND FUNCTIONS

In this work, we borrowed the connected graph notion from Card et al.'s work [26] and build on the original Zwicky box notation [159]. Card et al.'s design space [26] captures composition operations through connected graphs like merge (union) and layout. Feedforward interactions are represented as dots in the design space. Each interaction has its own ID, allowing the designer to specify multiple designs on the same representation. Connections signify *AND* operators (similar with merge from Card et al.[26]), whereas disconnected dots refer to *OR* relationship. Figure 43 shows the classical graph notation from Feedforward, similar to Card et al. [26].

Because Feedforward contained the parameter of triggering, it required an *and* (or merge) operator representation. Otherwise, the design space schema could not represent all possible feedforward designs, especially for trigger and untrigger. These logical connections mark a very important distinction needed to generate interactions properly: one may trigger a preview using gaze *or* location, or one may trigger previews using gaze *and* location. These types of triggers are distinct conceptually and require different implementations. To reflect this type of practical and conceptual distinction on the design space schema, the connected graph notation was needed.

During the morphological analysis process, values are combined with *and*. Another type of composition operator is represented on the cross-consistency

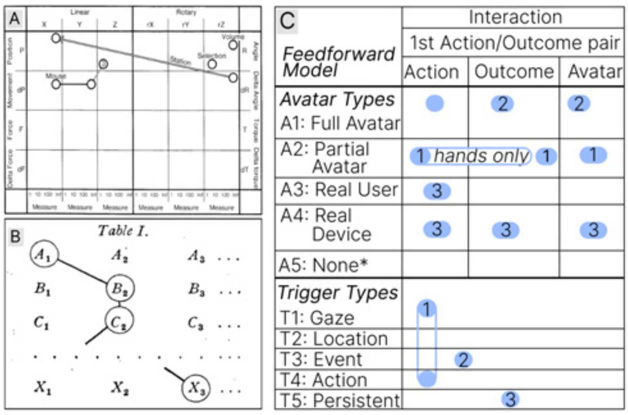

Figure 43: Subfigure A shows *Figure 3* from Card et al.'s representation of the design space of input devices [24]. B shows *Table 1* from Zwicky [136] original representation of the morphological box. C shows part of the feedforward design space, split by interaction outcomes and actions. The graphical notations allow selecting logical conditions with and (e.g., hands only) to be reflected in different types of feedforward design instances (e.g., 1, 2, 3).

matrix from MultipleAvatars. In morphological analysis, many times, the intersection between the same values of parameters is considered invalid or inconsistent. This was not the case for MultipleAvatars. Here, the same values could be combined differently — through an *and* and an *or*. Figure 44 shows an example of embedding *AND* and *OR* operations on a multidimensional matrix represented as a cross-consistency matrix for the MultipleAvatars design space.

These design space schemas or representations may have different advantages and disadvantages. The Feedforward design space schema is more user-friendly because it was developed as a conceptual design tool, later adjusted based on expert feedback. This makes it more suitable for practical uses, allowing practitioners to understand interaction models at a glance.

The MultipleAvatars representation captures the morphological analysis process. This representation is more useful for design space exploration because it reveals all possible combinations of parameters within the multiple avatar interaction model. However, its complexity may make it unsuitable as a user-friendly design tool. **(I8)** The implication here is that design-space schemas must account for their purpose, as each representation has advantages and disadvantages that make it more suitable for a task. This research suggests there could be a trade-off between the extent of theoretical information embedded in the design space, as HCI or VR concepts, and ease of understanding. This is important to consider if these design spaces would be used as tools by practitioners.

Implication 8 - method

6.4 IMPLICATIONS FOR USEFULNESS OF DESIGN SPACES

Greenberg and Buxton [94] remarked that usefulness is difficult to evaluate. Still, concepts are instances of visions about what interfaces may do, and while some may be inaccurate concerning the real-world instances they produce, vision has been critical. Throughout this thesis, design spaces have emerged as a structured approach to reasoning about novel interactions. So,

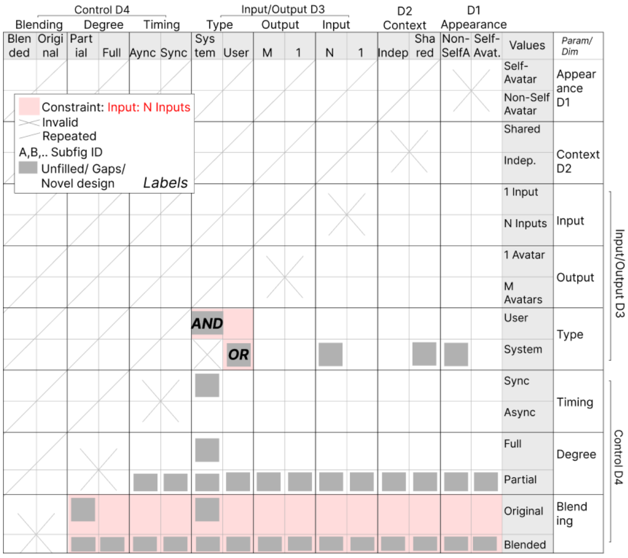

Figure 44: MultipleAvatars design space representation as a cross-consistency matrix. This figure highlights a way to represent *and/or* relationships in interactions. A multiple avatar interaction may be blended in a way in which a user and a system control another avatar simultaneously (*AND*), or each actor fully controls the avatars at certain moments in time (*OR*).

are design spaces useful as generative interaction tools? Their continuous use as a methodology indeed suggests so. However, in practice, it depends.

Feedforward embeds usefulness in the design space by starting from established theory and performing an expert evaluation in line with the proposed real-world usage of its research contribution. We evaluated the Feedforward design space with experts at different levels of their careers using virtual reality. During the interviews, one researcher could not find the usefulness in the design space. However, this number might be higher since participants might be biased towards preferring the artifacts if the researcher evaluating the results created them [127]. This suggests **(I9)** the learning curve of design spaces might prevent people from adopting them in practice.

Implication 9 -method

Experts may have already established an approach to idea generation (like brainstorming). Constricting this process to a design space creates tension. This tension emerges from a prerequisite of learning and understanding a design space before prototyping. Davies et al. [81] report similar concerns related to time trade-offs of conceptual modeling methods and remark that its advantages may depend on the analyst and the complexity of the task [81]. [3]. Echoing [173], this research suggests design space explorations are helpful for designers at the beginning of their careers since they make design decisions explicit. Moreover, the iterations before arriving at the final version of the design space allow designers to capture and reflect on the design process, similar to [173]'s *snapshots*.

Design spaces need not capture the full complexity of real-world practice [173] — they serve as a generative tool for designers. However, to fur-

---

[3] In practice, conceptual modeling is widely used in software development, especially the Unified Modeling Language [81].

ther close the *theory-practice gap* and make design spaces more practical[4], parameters may incorporate *"actionable insights"* [198] about how to design and implement an interface. Allowing experts to *explicitly* validate design spaces as tools for interaction generation helps embed relevant contextual parameters tied to practice. Starting from theoretical concepts like embodiment or immersion can also help embed practical considerations **(I10)** The implication here is that design spaces should embed technical aspects to aid in the generative design process. Considering the rise of scenario-based evaluation of VR research [298, 282, 203], defining usefulness at the beginning of an interaction design process is key to preventing design spaces from remaining *"point designs and empirical studies point studies"*, or design implications from being *if-then lists* [180].

Implication 10 - method

Feedforward and MultipleAvatars merge two current approaches to design-space making — the dimensions capture abstracted views of interactions [261, 220], whereas parameters embed technical or design aspects related to the development of these models as interfaces [261] (i.e., the concept of triggering previews vs how to trigger it in practice, and the concept of controlling avatars vs how to blend inputs in practice).

6.5 IMPLICATIONS FOR NOVELTY

It is difficult to foresee the impact of novel technologies or forecast consumer adoption [94]. In the initial stages of development, customers and designers may struggle to find uses for state-of-the-art technologies, especially if they underperform compared to existing tools [140]. Design spaces have a long-standing history of innovation — Zwicky used morphological analysis to make discoveries. For VR, design spaces innovate by revealing *gaps in knowledge* Danyluk et al. [261], Hirzle et al. [220], and Drey et al. [237]. This process involves assigning parameters to existing work and mapping it to the design space. Gaps then indicate novel opportunities for design. Changing parameters of existing work may also inform novel designs.

When requirements elicitation generates novel designs or techniques, they may suffer from the problem of *"unknown unknowns"*[147]. This problem may occur when researchers and practitioners lack domain knowledge — in this case, they miss relevant knowledge that might create dimensions or parameters of the interaction model. While establishing comprehensive analysis protocols might help, perhaps the way to overcome this might be to harness *"the power of human imagination"*[147]. MultipleAvatars attempted to tackle this issue by basing the design space on data from a brainstorming exercise with experts — the exercise constrained idea generation to VR domains and applications. **(I11)** The implication here is that adding a creative exercise in the design-space-making process might reveal novel dimensions.

Implication 11 - method

To reveal gaps and inform future design, MultipleAvatars mapped related works in the design space. Later, the search for related works was extended to AR — these AR works filled some of the initial gaps revealed by the MultipleAvatars design space (e.g., blending). While the design space's descriptive power was apparent, there is a caveat. If extending the scope of the related work fills in some of these gaps — this may question the generative power of the design space. How can we trust that design spaces put forward *"truly"* novel concepts and ideas? **(I12)** Perhaps accounting for commercial

Implication 12 - method

---

[4] The theory-practice gap is an HCI concept describing the space between academic research and its practical applicability [198].

applications when filling in gaps might be useful to strengthen the novelty claim of the design space. Systematic literature reviews do not account for commercial applications, which are a fast-developing VR resource of interaction design.

Moreover, for the field of VR, constrictive generative designs between types of extended realities might be a missed opportunity to discover related works that fill in the gaps. While having a common understanding of immersive concepts is important, terms like *AR/MR/VR* and even *VR fails* could fall out of use or change meaning completely [231]. Constraining the outcomes of this research solely to VR was relatively difficult as its boundaries were already overlapping with AR principles. The headsets that enable VR integrate a mixed-reality space beyond the play area boundary. Discussions of all works presented in this thesis expanded beyond what would be considered pure VR. **(I13)** Perhaps mapping generative interactions on the whole reality-virtuality continuum [35] is the way to design towards future immersive interactions. Discussions about the concept of virtuality might hinder this development.

Implication 13 - method

6.6 SUMMARY

This section presents a summary of the implications from the previous discussion. Based on this research, we propose that researchers should consider the following when adopting a concept-driven lens for interaction design in VR:

- **I1 - practical:** Having a good way to abstract research into a fun, simple concept has great potential for science communication.
- **I2 - method:** Leveraging HCI concepts from different domains, like seamfulness, may determine novel insights for interaction design in VR.
- **I3 - practical:** Prototyping as a process allows for concept exploration in practice, while authoring tools enable design space exploration in practice. Authoring tools reveal areas of design spaces, not just points.
- **I4 - practical** Developing design spaces as a generative tool for interactions means validating them and later changing them to reflect the design process of their target audience (practitioners, researchers, etc.).
- **I5 - method:** When literature review searches yield few results, researchers may want to explore this conceptual space of their interaction by abstracting it to a metaphor.
- **I6 - practical:** Concepts have underlying meanings that might affect the expectations of people using VR technologies that leverage such concepts. So, when choosing a metaphor or concept to describe an interaction technique, the researcher should consider its implications and the expectations it might give rise to.
- **I7 - method:** Approaching VR problems from a conceptual lens might change the understanding of the problem itself and uncover new problems and solutions.
- **I8 - method:** Design-space schemas must account for their purpose, as each representation has advantages and disadvantages that make

it more suitable for a task. Cross-consistency diagrams enable design space exploration, whereas connected graph representations of design spaces are more practical and user-friendly.

- **I9 - method:** The learning curve of design spaces might prevent people from adopting them in practice. Making design spaces more simple might aid in their adoption in practice.
- **I10 - method:** Design spaces should embed technical aspects to aid in generating interactions if used as generative tools in practice. This can be achieved by including VR-centred concepts in the interaction models (e.g., embodiment) or by running expert studies and adjusting the design space schema based on real-world design processes.
- **I11 - method:** The implication here is that adding a creative exercise in the design-space-making process might reveal novel dimensions.
- **I12 - method:** Accounting for VR commercial applications when filling in design space gaps might be useful to strengthen the novelty claim of the design space or generate novel dimensions.
- **I13 - method:** Mapping generative interactions on the whole reality-virtuality continuum grounds the interaction in future visions of the technology. Discussions about the concept of virtuality might hinder this development.

7

LIMITATIONS AND OUTLOOK

This section presents an outlook of conceptual design in VR and some limitations of this thesis and the research work it describes. The particular limitations and future opportunities of each work were presented in their respective chapters. Here, I focus on the broader limitations and future works relative to the conceptual-driven interaction design view of this thesis.

7.1 LIMITATIONS

Thoughtful designs should consider the real-world context in their design and evaluations (social, cultural, etc.) [88]. While we have discussed aspects of the 'social lives' of concepts, a potential limitation of this research is the lack of cultural and social grounding of concept generation. Including these aspects might improve the quality of conceptual interaction design. As future work, expanding to other concept design methods like future workshops [88] or inspiration cards [82] could reveal social, cultural, and political aspects of the futures envisioned by each work.

A limitation of design spaces as semantic structures is the presumed quality of the conceptual designs they generate or the gaps they reveal — design spaces do not ensure quality in the resulting design [70]. A measure of practical evaluation is required to make any claims about the quality of a design space-generated artifact. Even so, due to the large space of potentially unexplored designs, providing a meaningful evaluation of all variations in design spaces is difficult. Still, Feedforward presented a type of meaningful evaluation of a design space, considering it from a conceptual-tool generating perspective and practical design-generating perspective. Design spaces are also developed by the researcher and depend heavily on their interpretation of the data. This suffers from reliability issues as another designer may derive a different design space based on the same data.

Moreover, evaluating all possible parameter combinations within a design space would require considerable effort and time. While this endeavor is certainly possible, it involves a lifetime of work. Instead, the designer may reason about combining outcomes from deductive reasoning grounded in theory. This was also Card, Mackinlay, and Robertson [24]'s view concerning evaluating the points given by design spaces. Alternatively, designers may reason about what an exemplar of each dimension or design space area might look like and assess those particular prototypes. These exemplars capture the most salient aspects of a design space, like the values that make the most significant difference in appearance, usability, or technique relative to a particular use case scenario.

Design implications have their own caveats. Oulasvirta and Hornbæk [180] decry that *design implications are little more than incoherent lists of if-then rules. Their capacity to change design, which consists of multiple interrelated decisions, is limited. At the same time, they trivialize empirical findings .*". However, the purpose of these implications is to inform technological development and bring these findings to practitioners to address challenges outside of academia [295]. Beck and Ekbia [198] explain that practitioners often do not receive the *theoretical* message of HCI research outputs, which some have argued should be distilled in design implications as *"actionable insights"* [198]. Still, deriving novelty from existing work supposes that the research may account for much of the observed phenomena. To strengthen the novelty claims of design spaces, a second review of existing VR applications can attempt to fill in the gaps in the design space. Systematic literature reviews can provide taxonomies that classify related work, whereas design spaces offer valuable abstraction for generative interaction models. Using design spaces to account for generative models is a step towards overcoming *"research given by point designs"*, which cannot generalize to principles [180].

As Eriksson et al. [265] point out, design spaces fall into what Höök and Löwgren [129] consider *"generative intermediate-level knowledge"*. The *"lightweight nature"* of design spaces is a criticism of its underlying data, which is not consistent and varies across domains [173]. Dove, Hansen, and Halskov [173] explained that researchers have incorporated design spaces as a method to classify varying types of information for various purposes, from cognitive psychology to a design space of input devices. While design spaces reveal the grid of opportunities, contextualizing the resulting ideas is deferred to the designer, who uses "disciplined imagination" to decide on the feasibility and interest of the ideas. Note that *"knowledge is situated, so is that of the researchers studying them"*[87], which underscores how the outcome of research depends on and is limited by the interpreter's background.

Design spaces impose an abstraction level upon the data they represent. The abstraction may be a downside since it may diminish the designer's ability to retrieve relevant contextual usage from these instances [173]. This echoes a philosophical argument from Machery [109], who argues for *doing away with concepts* entirely from psychology research. His view is that concepts prevent an accurate characterization of various types of cognitive processes. Indeed, design spaces are not exhaustive sources of information and may never be. We echo Dove et al. [173]'s remark that design spaces are tools to abstract and systemize design [173], and provide helpful abstractions for interactions. The work within this thesis serves as an approach or a method for the generation of design spaces within a conceptual-driven lens of interaction design. This approach provides certain boundaries around design-space making, which help ground their generation in HCI theory and future visions.

Another downside of design spaces revolves around their failure to capture the *situated complexity* of interactivity, as they present snapshots of what interactions may involve. Berkel and Hornbæk [295] highlight that assuming the lens of real-world applicability is crucial when developing design implications. To add more ecological validity, the tools presented in this thesis should be evaluated in a real-world context. Dourish [57] emphasizes that interaction not only captures *"what is being done, but also as how it is being done. Interaction is the means by which work is accomplished, dynamically and in context."* Therefore, lab-based usability evaluations of the prototypes presented in this research need an account that could only be given by fieldwork. That is, investigating how people use these tools in the real world within a workplace scenario. Another limitation is the population sample within this work. Most participants for these studies were primarily people in their 20s and 30s and with a higher education background. Many were also experts in VR or generally knowledgeable with computers. More research needs to be done to account for how non-academic practitioners would use the design spaces and the artifacts they produced from this thesis.

Another argument that might prevent people from adopting design spaces, apart from the learning curve, is the argument for novelty. Putting forward an interaction technique asserts more novelty than building upon a different researcher's interactive model. Oulasvirta and Hornbæk [180] decried novelty as a damaging criterion for HCI and argued that papers need not be novel to improve the problem-solving capacity of research work. However, they remarked that novelty leads to one of the most prevalent HCI contributions and reflects a *"rapidly evolving technological landscape"* [180]. Mapping one's research on a pre-existing interaction model or design space may detract from this novelty, making the contribution less valuable to HCI. Therefore, researchers may be more inclined to leverage design spaces for novelty if they put the design space forward themselves.

7.2 FUTURE WORK

Conceptual review of VR interactions

As future work, I suggest a review that maps the conceptual avenues within VR interaction design. This review could delve into concept formation and result in a roadmap of research from the 90s to the state of the art. This review could take a conceptual lens of interaction design to reveal the blended spaces between interactions and roads not taken. Perhaps this road map could help identify qualities of seminal work like *Dynabook* or *Sketchpad*, answering questions about how paradigm-shifting work happens. This work could highlight the dimensions and parameters that made concepts more attractive to build upon. In Section 2.3, I described how recordings in VR [245] and the concept of worlds-in-miniature [37] resulted in a novel work that implements spatial design queries with direct manipulation [304]. A conceptual review of VR interactions could reveal the conceptual blends that give rise to novel, interesting, and important work.

While I briefly give examples of commercial work, extending design space reviews to include commercial VR applications and the reality-virtuality continuum might give novel insights about underlying interaction models Feedforward and MultipleAvatars. Borrowing from the VRFails methodology, a future feedforward paper could analyze videos of users playing games to look for clues about how practitioners design feedforward in the real world and perhaps give insights about whether these users understand the tutorials. To improve the validity of design space-making, we could borrow from VRFails qualitative content analysis and include inter-coder agreement procedures [74]. An avenue for future work would be to make design space exploration more systematic. This would involve conducting reviews of commercial applications or related work that map into design spaces after design spaces are developed. This would strengthen the novelty claim of design spaces.

Improving method

As future work, I suggest implementing the multi-avatar interface as an interface for a feedforward authoring tool. While feedforward also represents a contribution towards making VR authoring tools more accessible to non-technical users, recent research shows that even skilled VR developers face many barriers when developing VR and AR, from testing to implementation [236]. While many programming languages and environments exist for 3D and 2D development, most VR and AR applications are developed using Unity Engine or Unreal. This makes practitioners highly dependent

Accessibility

on these programs. If such businesses change their business model[1], it will make VR even less accessible. Thus, authoring tools are not only useful for prototyping and design but may serve as safety nets for users.

---

[1] A scenario very much grounded in recent events: https://www.pcworld.com/article/2081112/unity-walks-back-new-engine-pricing-after-protests.html

8

CONCLUSION AND SUMMARY

This thesis delves into concept creation and representation as *design spaces* and concept implementation and evaluation as *prototypes*. This thesis presents concepts as interaction models manifested through prototyping and design-space-making. The tenets followed throughout this work placed interaction design as rooted in the future, grounded in theory, and explorative in nature, aiming towards the concrete manifestation of visions. We found that particularly bridging concepts are suited to fulfill these ideas, which account for the tangible counterparts of these concepts (prototypes or exemplars) and the parameters they are shaped by. Design-space-making evolved as a valuable tool to systematize the concept-driven approach to exploration.

The research implications discussed underscore the importance of nuance and meaning when conceptualizing ideas at an early stage of interaction design. Representing novel interactions as interaction models and design spaces could be a valuable way to bridge the theory-practice gap. This could give practitioners and researchers the ability to situate work and determine its parameters quickly. In turn, they could make informed decisions concerning design before fully prototyping a model. This work highlights the importance of having a brief overview, or snapshot, of research not only related to a design process but, more broadly, to interaction models. And we have used design space in this manner — as snapshots of interaction models.

The research papers presented in this thesis have advanced the design of virtual reality interactions and interactions as follows:

1. VRFails developed the concept of VR fails from a seamful perspective,
2. VRFails proposed design implications as scenario-based sketches based on real-world use of design spaces in the home,
3. Feedforward developed the concept for feedforward in VR, instantiating it as an interaction model through a design space,
4. Feedforward presented an authoring tool prototype for immersive feedforward interaction design;
5. Feedforward presented findings from a qualitative evaluation of the design space as a design tool and of the designs enabled by the prototype to present design implications for future feedforward interactions;
6. MultipleAvatars developed the concept of acting through multiple avatars in VR based on expert brainstorming sessions grounded by real-world use cases,
7. MultipleAvatars proposed an interaction-generation model for the concept of acting through multiple avatars as a design space;
8. MultipleAvatars presented an interface that implements the concept of acting through multiple avatars and showed implications from a usability study for designing future multiple-avatar interactions;

9. MultipleAvatars interface may be used to record interactions by non-technically skilled designers and prototype tutorial and training applications more easily.

The theoretical aspects of interaction models and methods used to derive the design spaces that resulted in prototypes capture the *path from theory to practice*. The practicalities embedded in the parameters capture the *path from practice to theory*. These practicalities resulted from implementing prototypes that instantiate parts of the design space. Restructuring of the design space based on the expert feedback (as was the case for Feedforward) and prototyping serves as a *path from artifact to theory*. This research presents *design-space-making* as an evolving tool for systematizing concept generation in VR. The origins and purpose of design spaces mark it as a novelty generation tool. The implication is that innovation is not a coproduct of a design space but rather its purpose. Therefore, the design space should always presume innovation and *exploration*.

Taking a concept-driven approach to VR interaction design, this thesis has outlined some implications for future work and the importance of concepts for research. VRFails highlights the importance of simple concepts that can spill into popular culture and enhance scientific communication. Feedforward introduces a practical view of an established HCI concept, feedforward, that may be used as a tool by designers to implement training and tutorial applications. MultipleAvatars presents a preliminary model for interacting through multiple avatars in VR and showcases the descriptive, comparative, and generative power of concepts formalized as design spaces. These design spaces and their development methods emphasize the teleological aspect of design and highlight the tension between design and practice, which may contribute to the *light* perception of design spaces.

By assuming a concept-driven lens, this thesis highlights the versatility of concepts and the conceptual work that underlies producing novel concepts like *Sketchpad* or *Dynabook*, but for the case of VR. While thoughtful implementation supersedes any successful commercial product, concepts are the building blocks of innovative and future interfaces and a valuable currency for interaction generation. The works in this thesis have given an account of various types of intermediate knowledge and attempt to uncover the currency of an imaginary problem. The concepts considered evoked practical and tangible results in terms of prototypes. With the rise of VR as a prototyping medium, establishing a method that bridges the gap between concepts and prototypes at a glance to systematize this mountain of knowledge is crucial.

"The best way to predict the future is to invent it." – Alan Kay, probably

Part I

APPENDIX

A

APPENDIX

A.1 AR AND VR WORKS MAPPED INTO THE MULTIPLEAVATARS DESIGN SPACE

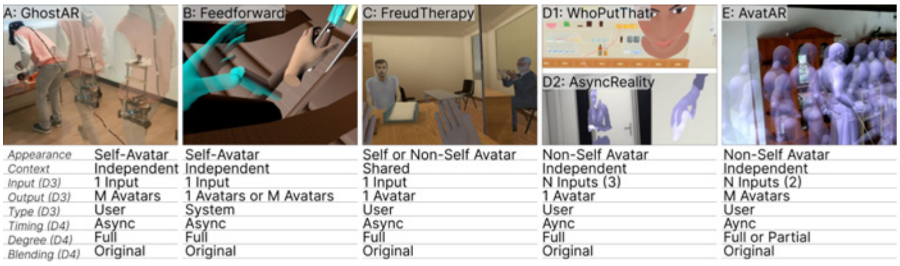

	A: GhostAR	B: Feedforward	C: FreudTherapy	D1: WhoPutThat D2: AsyncReality	E: AvatAR
Appearance	Self-Avatar	Self-Avatar	Self or Non-Self Avatar	Non-Self Avatar	Non-Self Avatar
Context	Independent	Independent	Shared	Independent	Independent
Input (D3)	1 Input	1 Input	1 Input	N Inputs (3)	N Inputs (2)
Output (D3)	M Avatars	1 Avatars or M Avatars	1 Avatar	1 Avatar	M Avatars
Type (D3)	User	System	User	User	User
Timing (D4)	Async	Async	Async	Aync	Aync
Degree (D4)	Full	Full	Full	Full	Full or Partial
Blending (D4)	Original	Original	Original	Original	Original

Figure 45: This figure shows selected related works characterized by dimensions and parameters of the MultipleAvatars Section 5.5.4. Subfigure A shows [214], B shows [305], C shows [230], D shows [282], E shows [289].

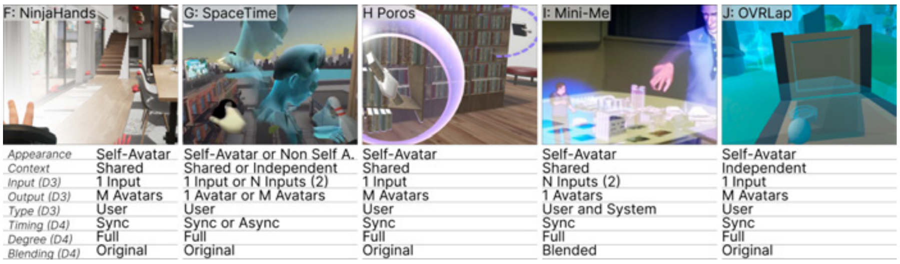

	F: NinjaHands	G: SpaceTime	H Poros	I: Mini-Me	J: OVRLap
Appearance	Self-Avatar	Self-Avatar or Non Self A.	Self-Avatar	Self-Avatar	Self-Avatar
Context	Shared	Shared or Independent	Shared	Shared	Independent
Input (D3)	1 Input	1 Input or N Inputs (2)	1 Input	N Inputs (2)	1 Input
Output (D3)	M Avatars	1 Avatar or M Avatars	M Avatars	1 Avatars	M Avatars
Type (D3)	User	User	User	User and System	User
Timing (D4)	Sync	Sync or Async	Sync	Sync	Sync
Degree (D4)	Full	Full	Full	Full	Full
Blending (D4)	Original	Original	Original	Blended	Original

Figure 46: This figure shows selected related works characterized by dimensions and parameters of the MultipleAvatars Section 5.6. Subfigure F shows [277], G shows [211], H shows [275], I shows [206], J shows [291].

A.2 MULTIPLEAVATARS USABILITY STUDY SET UP

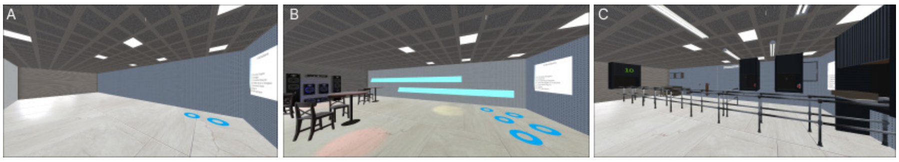

Figure 47: This figure contains camera views of the three setups from the usability study in Section 5.6. Subfigure A shows the setups of Phases 1 and 2 (*Multiplayer* and *Arguing*). Subfigure B shows the setup for Phase 3, *Dancing*. Subfigure C shows the setup of the last phase, *Pipeline*. Participants would place clones on the blue circles and the white panel contains instructions for each phase.

A.3 THE CROSS-CONSISTENCY MATRIX WITH AR WORKS FOR MULTIPLEAVATARS

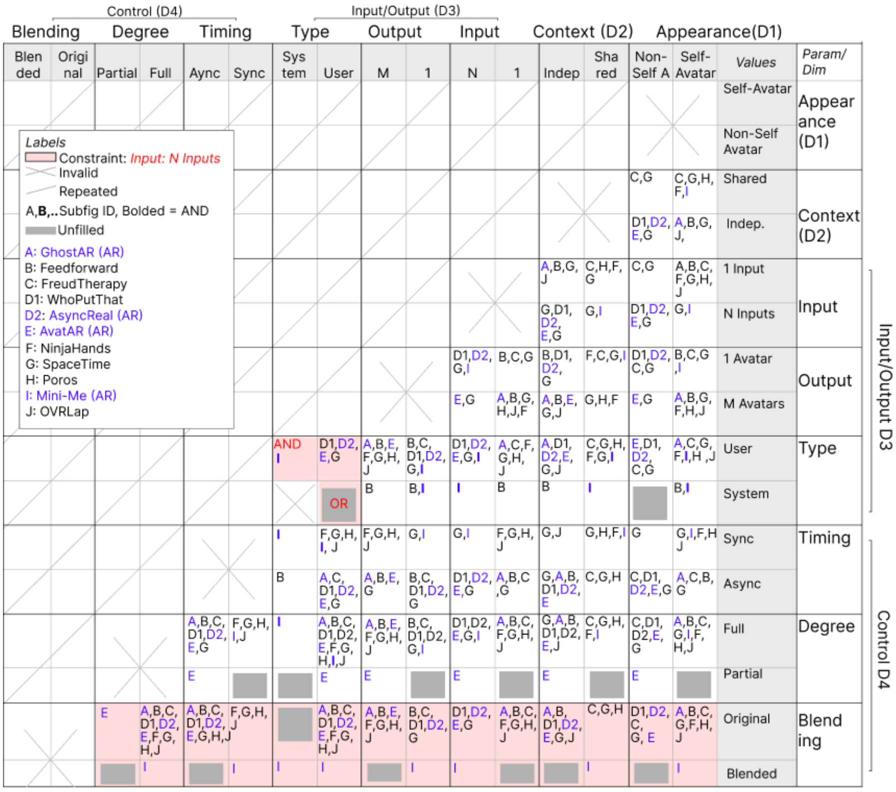

Labels
🟥 Constraint: *Input: N Inputs*
Invalid
Repeated
A,**B**,..Subfig ID, Bolded = AND
▬ Unfilled

A: GhostAR (AR)
B: Feedforward
C: FreudTherapy
D1: WhoPutThat
D2: AsyncReal (AR)
E: AvatAR (AR)
F: NinjaHands
G: SpaceTime
H: Poros
I: Mini-Me (AR)
J: OVRLap

Control (D4)						Input/Output (D3)												
Blending		Degree		Timing		Type		Output		Input		Context (D2)		Appearance(D1)				
Blended	Original	Partial	Full	Aync	Sync	System	User	M	1	N	1	Indep	Shared	Non-Self A	Self-Avatar	*Values*	*Param/ Dim*	
																Self-Avatar	Appearance (D1)	
																Non-Self Avatar		
														C,G	C,G,H,F,I	Shared	Context (D2)	
														D1,D2,E,G	A,B,G,J,	Indep.		
												A,B,G,J	C,H,F,G	C,G	A,B,C,F,G,H,J	1 Input	Input	Input/Output D3
												G,D1,D2,E,G	G,I	D1,D2,E,G	G,I	N Inputs		
										D1,D2,G,I	B,C,G	B,D1,D2,G	F,C,G,I	D1,D2,C,G	B,C,G ,I	1 Avatar	Output	
										E,G	A,B,G,H,J,F	A,B,E,G,J	G,H,F	E,G	A,B,G,F,H,J	M Avatars		
						AND **I**	D1,D2,E,G	A,B,E,F,G,H,J	B,C,D1,D2,G,**I**	D1,D2,E,G,**I**	A,C,F,G,H,J	A,D1,D2,E,G,J	C,G,H,F,G,**I**	E,D1,D2,C,G	A,C,G,F,**I**,H ,J	User	Type	
							OR ▬	B	B,**I**	**I**	B	B	**I**	▬	B,**I**	System		
						I	F,G,H,**I**, J	F,G,H,J	G,I	G,I	F,G,H,J	G,J	G,H,F,I	G	G,I,F,H J	Sync	Timing	Control D4
						B	A,C,D1,D2,E,G	A,B,E,G	B,C,D1,D2,G	D1,D2,E,G	A,B,C ,G	G,A,B,D1,D2,E	C,G,H	C,D1,D2,E,G	A,C,B,G	Async		
				A,B,C,D1,D2,E,G	F,G,H,I,J	**I**	A,B,C,D1,D2,E,F,G,H,**I**,J	A,B,E,F,G,H,J	B,C,D1,D2,G,I	D1,D2,E,G,I	A,B,C,F,G,H,J	G,A,B,D1,D2,E,J	C,G,H,F,I	C,D1,D2,E,G	A,B,C,G,I,F,H,J	Full	Degree	
				E	▬	▬	E	E	▬	E	▬	E	▬	E	▬	Partial		
		E	A,B,C,D1,D2,E,F,G,H,J	A,B,C,D1,D2,E,G,H,J	F,G,H,J	▬	A,B,C,D1,D2,E,F,G,H,J	A,B,E,F,G,H,J	B,C,D1,D2,G	D1,D2,E,G	A,B,C,F,G,H,J	A,B,D1,D2,E,G,J	C,G,H	D1,D2,C,G, E	A,B,C,G,F,H,J	Original	Blending	
		▬	I	▬	I	I	I	▬	I	I	I	▬	I	▬	I	Blended		

Figure 48: This figure shows selected related works characterized by dimensions and parameters of the multiple avatars design space in Section 5.5.4 with added works from AR (marked blue) to fill in some gaps.

A.4 IMPLEMENTED SYSTEM OPERATIONS DESCRIPTION FOR MULTIPLEAVATARS

We implemented the operations within the multiple avatar interaction prototype as follows:

1. Manage avatars
   We designed an avatar management system using a portable mid-air control menu, led by the handle (Figure 34 A). The handle contains some text that can be seen through objects to help users locate it when obstructed. Each avatar receives a unique identifier (id) when created, which appears on top of it when spawned (seen in Figure 34 C). Users can select avatars by the two arrows on the menu. When selected, the id on top is red and white otherwise. The menu displays the total number of avatars and the selected avatar's status (hidden, playing, or paused). Users can perform all operations (except for recording and setting replays) either for a selected avatar or for all avatars. Users can switch between crowd control (Figure 34 A), or single control (Figure 34 B) using a button on top.

2. Move and rotate avatars
   Users can move and rotate avatars by grabbing a handle that was placed at the avatar's chest. The handle is an object similar to the menu handle but located inside the avatar. It also has text visible when covered to signal its location to users (seen in Figure 34 C).

3. Remove avatars
   There are two types of avatar removals: *Hide*, to remove the avatars from the scene but not the system, and *Delete*, to remove the avatars completely.

4. Playback
   We add two syncing operations to *Pause* and *Play* avatars. These can be used to sync avatars together or to some other external event. Users can also set the number of repetitions for every newly spawned avatar in the *Crowd* menu under *Replays*. The button cycles through 5 options (repeating 1, 5, 10, and 100 times).

5. Copy-cat avatar
   We implement avatars that synchronously track the movements of the user. We use this feature to establish ownership in the introduction phase. Similar to a mirrored avatar, the copy-cat avatar in Figure 34 is rotated 180 toward the user and tracks their movements.

A.5 AFFINITY DIAGRAM PROGRESSION FOR MULTIPLEAVATARS

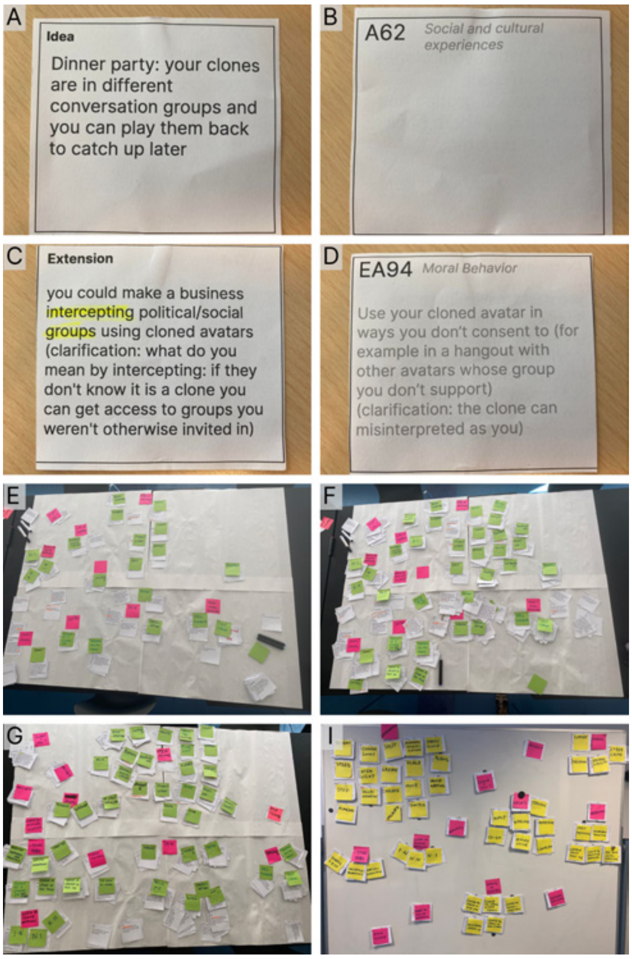

Idea

Dinner party: your clones are in different conversation groups and you can play them back to catch up later

A62 *Social and cultural experiences*

Extension

you could make a business intercepting political/social groups using cloned avatars (clarification: what do you mean by intercepting: if they don't know it is a clone you can get access to groups you weren't otherwise invited in)

EA94 *Moral Behavior*

Use your cloned avatar in ways you don't consent to (for example in a hangout with other avatars whose group you don't support) (clarification: the clone can misinterpreted as you)

Figure 49: This figure shows the affinity diagram notes and formation. A shows a workshop idea, and B shows the back side with its ID and category. E, F, G, I show the affinity diagram progression.

A.6 POST-EVALUATION DESIGN SPACE OF FEEDFORWARD IN VR

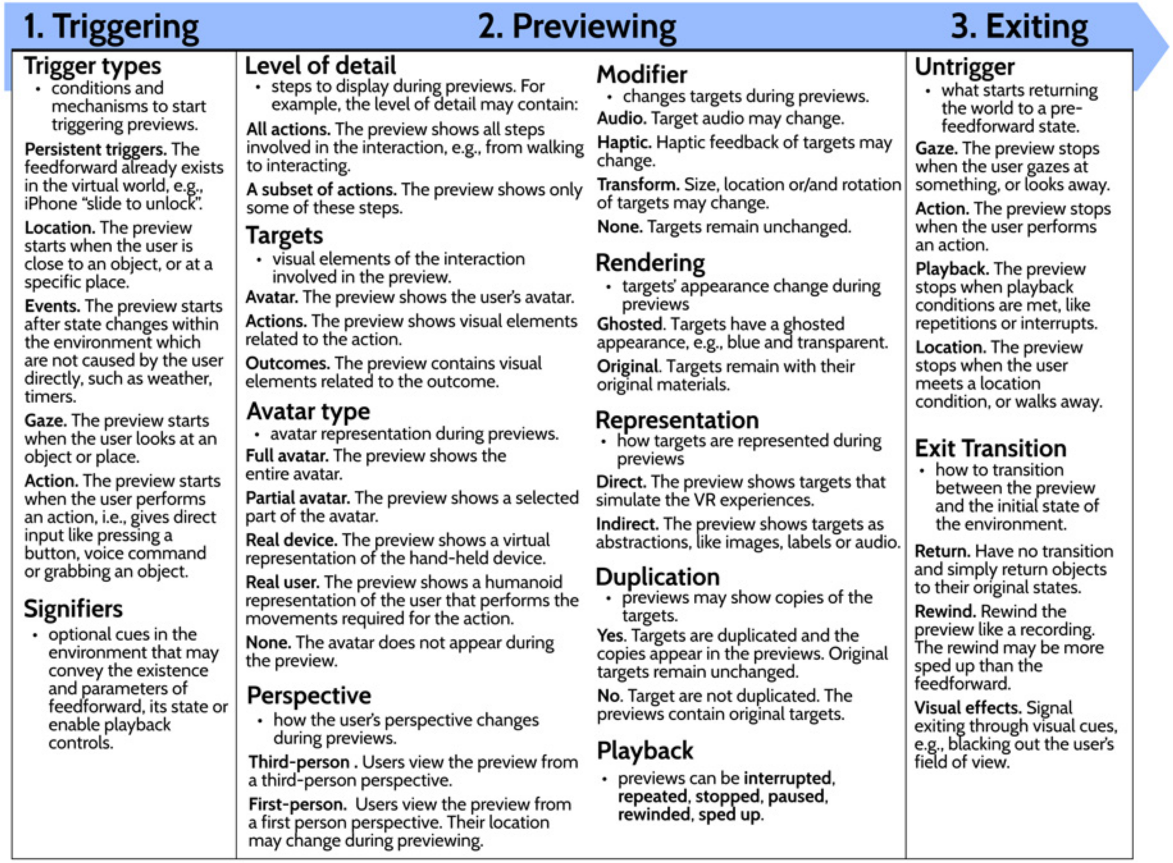

1. Triggering

Trigger types

- conditions and mechanisms to start triggering previews.

Persistent triggers. The feedforward already exists in the virtual world, e.g., iPhone "slide to unlock".

Location. The preview starts when the user is close to an object, or at a specific place.

Events. The preview starts after state changes within the environment which are not caused by the user directly, such as weather, timers.

Gaze. The preview starts when the user looks at an object or place.

Action. The preview starts when the user performs an action, i.e., gives direct input like pressing a button, voice command or grabbing an object.

Signifiers

- optional cues in the environment that may convey the existence and parameters of feedforward, its state or enable playback controls.

2. Previewing

Level of detail

- steps to display during previews. For example, the level of detail may contain:

All actions. The preview shows all steps involved in the interaction, e.g., from walking to interacting.

A subset of actions. The preview shows only some of these steps.

Targets

- visual elements of the interaction involved in the preview.

Avatar. The preview shows the user's avatar.

Actions. The preview shows visual elements related to the action.

Outcomes. The preview contains visual elements related to the outcome.

Avatar type

- avatar representation during previews.

Full avatar. The preview shows the entire avatar.

Partial avatar. The preview shows a selected part of the avatar.

Real device. The preview shows a virtual representation of the hand-held device.

Real user. The preview shows a humanoid representation of the user that performs the movements required for the action.

None. The avatar does not appear during the preview.

Perspective

- how the user's perspective changes during previews.

Third-person . Users view the preview from a third-person perspective.

First-person. Users view the preview from a first person perspective. Their location may change during previewing.

Modifier

- changes targets during previews.

Audio. Target audio may change.

Haptic. Haptic feedback of targets may change.

Transform. Size, location or/and rotation of targets may change.

None. Targets remain unchanged.

Rendering

- targets' appearance change during previews

Ghosted. Targets have a ghosted appearance, e.g., blue and transparent.

Original. Targets remain with their original materials.

Representation

- how targets are represented during previews

Direct. The preview shows targets that simulate the VR experiences.

Indirect. The preview shows targets as abstractions, like images, labels or audio.

Duplication

- previews may show copies of the targets.

Yes. Targets are duplicated and the copies appear in the previews. Original targets remain unchanged.

No. Target are not duplicated. The previews contain original targets.

Playback

- previews can be **interrupted**, **repeated**, **stopped**, **paused**, **rewinded**, **sped up**.

3. Exiting

Untrigger

- what starts returning the world to a pre-feedforward state.

Gaze. The preview stops when the user gazes at something, or looks away.

Action. The preview stops when the user performs an action.

Playback. The preview stops when playback conditions are met, like repetitions or interrupts.

Location. The preview stops when the user meets a location condition, or walks away.

Exit Transition

- how to transition between the preview and the initial state of the environment.

Return. Have no transition and simply return objects to their original states.

Rewind. Rewind the preview like a recording. The rewind may be more sped up than the feedforward.

Visual effects. Signal exiting through visual cues, e.g., blacking out the user's field of view.

Figure 50: This figure represents the updated version of the design space and cheat sheet after the expert evaluation.

[1] Erving Goffman. *The Presentation of Self in Everyday Life.* New York, USA: Anchor Books, 1959.

[2] F. Zwicky. "The Morphological Approach to Discovery, Invention, Research and Construction." In: *New Methods of Thought and Procedure*. Ed. by Fritz Zwicky and Albert G. Wilson. Berlin, Heidelberg: Springer Berlin Heidelberg, 1967, pp. 273–297. ISBN: 978-3-642-87617-2. DOI: `https://doi.org/10.1007/978-3-642-87617-2_14`.

[3] Joseph A Goguen. "The logic of inexact concepts." In: *Synthese* 19.3/4 (1969), pp. 325–373.

[4] Joseph L. Fleiss and Jacob Cohen. "The Equivalence of Weighted Kappa and the Intraclass Correlation Coefficient as Measures of Reliability." In: *Educational and Psychological Measurement* 33.3 (1973), pp. 613–619. DOI: `10.1177/001316447303300309`. eprint: `https://doi.org/10.1177/001316447303300309`. URL: `https://doi.org/10.1177/001316447303300309`.

[5] Horst WJ Rittel and Melvin M Webber. "Dilemmas in a general theory of planning." In: *Policy sciences* 4.2 (1973), pp. 155–169.

[6] Jesse L Greenstein and Albert G Wilson. "Remembering Zwicky." In: *Engineering and Science* 37.5 (1974), pp. 15–19.

[7] Eleanor Rosch and Carolyn B Mervis. "Family resemblances: Studies in the internal structure of categories." In: *Cognitive psychology* 7.4 (1975), pp. 573–605.

[8] John W Sutherland. "Systems: Analysis." In: *Administration, and Architecture* (1975).

[9] Albert Bandura and Richard H Walters. *Social learning theory*. Vol. 1. Englewood cliffs Prentice Hall, 1977.

[10] J. Richard Landis and Gary G. Koch. "The Measurement of Observer Agreement for Categorical Data." In: *Biometrics* 33.1 (1977).

[11] Douglas L Medin and Marguerite M Schaffer. "Context theory of classification learning." In: *Psychological review* 85.3 (1978), p. 207.

[12] Carolyn B Mervis and Eleanor Rosch. "Categorization of natural objects." In: *Annual review of psychology* 32.1 (1981), pp. 89–115.

[13] James D Foley, Victor L Wallace, and Peggy Chan. "The human factors of computer graphics interaction techniques." In: *IEEE computer Graphics and Applications* 4.11 (1984), pp. 13–48.

[14] Hideki Ishio, Junichiro Minowa, and Kiyoshi Nosu. "Review and status of wavelength-division-multiplexing technology and its application." In: *Journal of lightwave technology* 2.4 (1984), pp. 448–463.

[15] Robert M Nosofsky. "Choice, similarity, and the context theory of classification." In: *Journal of Experimental Psychology: Learning, memory, and cognition* 10.1 (1984), p. 104.

[16] Arthur B VanGundy. "Brain writing for new product ideas: an alternative to brainstorming." In: *Journal of Consumer Marketing* (1984).

[17] Jock Mackinlay. "Automating the design of graphical presentations of relational information." In: *Acm Transactions On Graphics (Tog)* 5.2 (1986), pp. 110–141.

[18] Ronald Baecker and William Buxton. "Readings in human-computer interaction: A multidisciplinary approach." In: (1987).

[19] Jock Mackinlay. "Applying a theory of graphical presentation to the graphic design of user interfaces." In: *Proceedings of the 1st annual ACM SIGGRAPH symposium on User Interface Software*. 1988, pp. 179–189.

[20] Donald A Norman. *The psychology of everyday things.* New York: Basic Books, 1988. URL: https://hal.science/hal-00692043.

[21] Zenon W Pylyshyn and Ron W Storm. "Tracking multiple independent targets: Evidence for a parallel tracking mechanism." In: *Spatial vision* 3.3 (1988), pp. 179–197.

[22] Allan MacLean, Richard M Young, and Thomas P Moran. "Design rationale: the argument behind the artifact." In: *ACM SIGCHI Bulletin* 20.SI (1989), pp. 247–252.

[23] Karl E Weick. "Theory construction as disciplined imagination." In: *Academy of management review* 14.4 (1989), pp. 516–531. DOI: doi.org/10.5465/amr.1989.4308376.

[24] Stuart K Card, Jock D Mackinlay, and George G Robertson. "The design space of input devices." In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1990, pp. 117–124.

[25] Jock Mackinlay, Stuart K Card, and George G Robertson. "A semantic analysis of the design space of input devices." In: *Human–Computer Interaction* 5.2-3 (1990), pp. 145–190.

[26] Stuart K Card, Jock D Mackinlay, and George G Robertson. "A morphological analysis of the design space of input devices." In: *ACM Transactions on Information Systems (TOIS)* 9.2 (1991), pp. 99–122.

[27] William W Gaver. "Technology affordances." In: *Proceedings of the SIGCHI conference on Human factors in computing systems Reaching through technology - CHI '91*. the SIGCHI conference. New Orleans, Louisiana, United States: ACM, 1991, pp. 79–84. ISBN: 978-0-89791-383-6. DOI: 10.1145/108844.108856. URL: http://portal.acm.org/citation.cfm?doid=108844.108856 (visited on 01/16/2023).

[28] Jiro Kawakita. "The original KJ method." In: *Tokyo: Kawakita Research Institute* 5 (1991), p. 1991.

[29] Allan MacLean, Richard M. Young, Victoria M.E. Bellotti, and Thomas P. Moran. "Questions, Options, and Criteria: Elements of Design Space Analysis." In: *Human–Computer Interaction* 6.3-4 (1991), pp. 201–250. DOI: 10.1080/07370024.1991.9667168. URL: https://www.tandfonline.com/doi/abs/10.1080/07370024.1991.9667168.

[30] Laurence Nigay and Joëlle Coutaz. "A design space for multimodal systems: concurrent processing and data fusion." In: *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. 1993, pp. 172–178.

[31] Kenneth D Bailey. *Typologies and taxonomies: An introduction to classification techniques*. Vol. 102. Sage, 1994.

[32] Harold Pashler. "Dual-task interference in simple tasks: data and theory." In: *Psychological bulletin* 116.2 (1994), p. 220.

[33] Susanne Bødker. "Applying Activity Theory to Video Analysis: How to Make Sense of Video Data in Human-Computer Interaction." In: *Context and Consciousness: Activity Theory and Human-Computer Interaction*. USA: Massachusetts Institute of Technology, 1995, 147–174. ISBN: 0262140586.

[34] Beverly L Harrison, Hiroshi Ishii, Kim J Vicente, and William AS Buxton. "Transparent layered user interfaces: An evaluation of a display design to enhance focused and divided attention." In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1995, pp. 317–324.

[35] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. "Augmented reality: A class of displays on the reality-virtuality continuum." In: *Telemanipulator and telepresence technologies*. Vol. 2351. Spie. 1995, pp. 282–292.

[36] John Rieman, Marita Franzke, and David Redmiles. "Usability evaluation with the cognitive walkthrough." In: *Conference companion on Human factors in computing systems*. 1995, pp. 387–388.

[37] Richard Stoakley, Matthew J Conway, and Randy Pausch. "Virtual reality on a WIM: interactive worlds in miniature." In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1995, pp. 265–272.

[38] Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. "The Go-Go Interaction Technique: Non-Linear Mapping for Direct Manipulation in VR." In: *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*. UIST '96. Seattle, Washington, USA: ACM, 1996, 79–80. ISBN: 0897917987. DOI: `10.1145/237091.237102`. URL: `https://doi.org/10.1145/237091.237102`.

[39] Doug A. Bowman and Larry F. Hodges. "An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments." In: *Proceedings of the 1997 symposium on Interactive 3D graphics - SI3D '97*. Providence, Rhode Island, United States: ACM Press, 1997, 35–ff. ISBN: 978-0-89791-884-8. DOI: `10.1145/253284.253301`. URL: `http://portal.acm.org/citation.cfm?doid=253284.253301` (visited on 01/16/2023).

[40] Tushar H Dani and Rajit Gadh. "Creation of concept shape designs via a virtual reality interface." In: *Computer-Aided Design* 29.8 (1997), pp. 555–563.

[41] Boris Velichkovsky, Andreas Sprenger, and Pieter Unema. "Towards gaze-mediated interaction: Collecting solutions of the "Midas touch problem"." In: *Human-Computer Interaction INTERACT '97*. Ed. by Steve Howard, Judy Hammond, and Gitte Lindgaard. Boston, MA: Springer US, 1997, pp. 509–516. ISBN: 978-1-4757-5437-7 978-0-387-35175-9. DOI: `10.1007/978-0-387-35175-9_77`. URL: `http://link.springer.com/10.1007/978-0-387-35175-9_77` (visited on 01/16/2023).

[42] Tobias Blickle, Jürgen Teich, and Lothar Thiele. "System-level synthesis using evolutionary algorithms." In: *Design Automation for Embedded Systems* 3 (1998), pp. 23–58.

[43] Ivan Poupyrev, Tadao Ichikawa, Suzanne Weghorst, and Mark Billinghurst. "Egocentric object manipulation in virtual environments: empirical evaluation of interaction techniques." In: *Computer graphics forum*. Vol. 17. 3. Wiley Online Library. 1998, pp. 41–52.

[44] Alistair G Sutcliffe and Michele Ryan. "Experience with SCRAM, a scenario requirements analysis method." In: *Proceedings of IEEE International Symposium on Requirements Engineering: RE'98*. IEEE. 1998, pp. 164–171.

[45] Ivan E. Sutherland. "Sketchpad—a Man-Machine Graphical Communication System." In: *Seminal Graphics: Pioneering Efforts That Shaped the Field*. New York, NY, USA: Association for Computing Machinery, 1998, 391–408. ISBN: 158113052X. URL: https://doi.org/10.1145/280811.281031.

[46] Doug A Bowman, Donald B Johnson, and Larry F Hodges. "Testbed evaluation of virtual environment interaction techniques." In: *Proceedings of the ACM symposium on Virtual reality software and technology*. 1999, pp. 26–33.

[47] Donald A Norman. "Affordance, conventions, and design." In: *Interactions* 6.3 (May 1999), pp. 38–43. ISSN: 1072-5520, 1558-3449. DOI: 10.1145/301153.301168. URL: https://dl.acm.org/doi/10.1145/301153.301168 (visited on 01/16/2023).

[48] Jeffrey S Pierce, Brian C Stearns, and Randy Pausch. "Voodoo dolls: seamless interaction at multiple scales in virtual environments." In: *Proceedings of the 1999 symposium on Interactive 3D graphics*. 1999, pp. 141–145.

[49] Ivan Poupyrev and Tadao Ichikawa. "Manipulating Objects in Virtual Worlds: Categorization and Empirical Evaluation of Interaction Techniques." In: *Journal of Visual Languages & Computing* 10.1 (Feb. 1999), pp. 19–35. ISSN: 1045926X. DOI: 10.1006/jvlc.1998.0112. URL: https://linkinghub.elsevier.com/retrieve/pii/S1045926X98901124 (visited on 01/16/2023).

[50] Karen O Solomon, Douglas L Medin, and Elizabeth Lynch. "Concepts do more than categorize." In: *Trends in cognitive sciences* 3.3 (1999), pp. 99–105.

[51] Kay M Stanney, Robert S Kennedy, Julie M Drexler, and Deborah L Harm. "Motion sickness and proprioceptive aftereffects following virtual environment exposure." In: *Applied ergonomics* 30.1 (1999), pp. 27–38.

[52] Michel Beaudouin-Lafon. "Instrumental interaction: an interaction model for designing post-WIMP user interfaces." In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2000, pp. 446–453.

[53] Mel Slater and Anthony Steed. "A Virtual Presence Counter." In: *Presence: Teleoper. Virtual Environ.* 9.5 (Oct. 2000), 413–434. ISSN: 1054-7460. DOI: 10.1162/105474600566925. URL: https://doi.org/10.1162/105474600566925.

[54] Albert Bandura. "Social cognitive theory: An agentic perspective." In: *Annual Review of Psychology* 52 (2001), pp. 1–26. ISSN: 1545-2085(Electronic),0066-4308(Print). DOI: 10.1146/annurev.psych.52.1.1.

[55] Victoria Bellotti and Keith Edwards. "Intelligibility and Accountability: Human Considerations in Context-Aware Systems." In: *Human–Computer Interaction* 16.2 (Dec. 2001), pp. 193–212. ISSN: 0737-0024, 1532-7051. DOI: 10.1207/S15327051HCI16234_05. URL: https://www.tandfonline.com/doi/full/10.1207/S15327051HCI16234_05 (visited on 12/01/2022).

[56] Seana Coulson. *Semantic leaps: Frame-shifting and conceptual blending in meaning construction*. Cambridge University Press, 2001.

[57] Paul Dourish. *Where the action is*. MIT press Cambridge, 2001.

[58] Mohammad El-Ramly, Eleni Stroulia, and Paul Sorenson. "Mining system-user interaction traces for use case models." In: *Proceedings 10th International Workshop on Program Comprehension*. IEEE. 2002, pp. 21–29.

[59] Chris Greenhalgh, Martin Flintham, Jim Purbrick, and Steve Benford. "Applications of temporal links: Recording and replaying virtual environments." In: *Proceedings IEEE Virtual Reality 2002*. IEEE. 2002, pp. 101–108.

[60] Anna Spagnolli and Luciano Gamberini. "Immersion/Emersion: Presence in hybrid environments." In: *Proceedings of the fifth annual international workshop, Presence*. Presence, 2002, pp. 421–434.

[61] Stanislav L Stoev and Dieter Schmalstieg. "Application and taxonomy of through-the-lens techniques." In: *Proceedings of the ACM symposium on Virtual reality software and technology*. 2002, pp. 57–64.

[62] Ungyeon Yang and Gerard Jounghyun Kim. "Implementation and Evaluation of "Just Follow Me": An Immersive, VR-Based, Motion-Training System." In: *Presence: Teleoper. Virtual Environ.* 11.3 (June 2002), 304–323. ISSN: 1054-7460. DOI: 10.1162/105474602317473240. URL: https://doi.org/10.1162/105474602317473240.

[63] Matthew Chalmers, Ian MacColl, and Marek Bell. "Seamful design: showing the seams in wearable computing." In: *2003 IEE Eurowearable*. 2003, pp. 11–16.

[64] Daniel Fallman. “Design-oriented human-computer interaction.” In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2003, pp. 225–232.

[65] Gilles Fauconnier and Mark Turner. “Conceptual blending, form and meaning.” In: *Recherches en communication* 19 (2003), pp. 57–86.

[66] Alex Olwal Steven Feiner. “The Flexible Pointer: An Interaction Technique for Selection in Augmented and Virtual Reality.” In: *Conference Supplement of UIST '03*. Vol. 3. Vancouver, BC, 2003, pp. 81–82.

[67] Robert L Goldstone, Alan Kersten, and Paulo F Carvalho. “Concepts and categorization.” In: *Comprehensive handbook of psychology* 4 (2003), pp. 599–621.

[68] Ann M Hickey and Alan M Davis. “Requirements elicitation and elicitation technique selection: model for two knowledge-intensive software development processes.” In: *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*. IEEE. 2003, 10–pp.

[69] Paddy McAloon. *I Trawl the Megahertz*. Digital Album. I Trawl the Megahertz. 2003.

[70] Michel Beaudouin-Lafon. “Designing interaction, not interfaces.” In: *Proceedings of the working conference on Advanced visual interfaces*. 2004, pp. 15–22.

[71] Peter J Denning. “The social life of innovation.” In: *Communications of the ACM* 47.4 (2004), pp. 15–19.

[72] Jun Fujima, Aran Lunzer, Kasper Hornbæk, and Yuzuru Tanaka. “Clip, connect, clone: combining application elements to build custom interfaces for information access.” In: *Proceedings of the 17th annual ACM symposium on User interface software and technology*. 2004, pp. 175–184.

[73] Ann M Hickey and Alan M Davis. “A unified model of requirements elicitation.” In: *Journal of management information systems* 20.4 (2004), pp. 65–84.

[74] Philipp Mayring. “Qualitative Content Analysis in Flickr. In Kardorff, U., Steinke, E. eds.” In: *A Companion to Qualitative Research*. London: Sage, 2004.

[75] Gregory Murphy. *The big book of concepts*. MIT press, 2004.

[76] Daniel Vogel and Ravin Balakrishnan. “Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users.” In: *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology*. UIST '04. Santa Fe, NM, USA:

ACM, 2004, 137–146. ISBN: 1581139578. DOI: 10.1145/1029632.1029656. URL: https://doi.org/10.1145/1029632.1029656.

[77] Stephan AG Wensveen, Johan Partomo Djajadiningrat, and CJ Overbeeke. "Interaction Frogger: A Design Framework to Couple Action and Function through Feedback and Feedforward." In: *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*. DIS '04. Cambridge, MA, USA: ACM, 2004, pp. 177–184. ISBN: 1581137877. DOI: 10.1145/1013115.1013140. URL: https://doi.org/10.1145/1013115.1013140.

[78] Stuart Reeves, Steve Benford, Claire O'Malley, and Mike Fraser. "Designing the Spectator Experience." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '05. Portland, Oregon, USA: Association for Computing Machinery, 2005, 741–750. ISBN: 1581139985. DOI: 10.1145/1054972.1055074. URL: https://doi.org/10.1145/1054972.1055074.

[79] Bo Westerlund. "Design space conceptual tool–grasping the design process." In: *Nordes* 1 (2005).

[80] Virginia Braun and Victoria Clarke. "Using thematic analysis in psychology." In: *Qualitative Research in Psychology* 3.2 (Jan. 2006), pp. 77–101. ISSN: 1478-0887, 1478-0895. DOI: 10.1191/1478088706qp063oa. URL: http://www.tandfonline.com/doi/abs/10.1191/1478088706qp063oa (visited on 08/07/2022).

[81] Islay Davies, Peter Green, Michael Rosemann, Marta Indulska, and Stan Gallo. "How do practitioners use conceptual modeling in practice?" In: *Data & Knowledge Engineering* 58.3 (2006), pp. 358–380.

[82] Kim Halskov and Peter Dalsgård. "Inspiration card workshops." In: *Proceedings of the 6th conference on Designing Interactive systems*. 2006, pp. 2–11.

[83] John Charles Mannone. "A Propensity for Genius: That Something Special About Fritz Zwicky (1898-1974)." In: *Swiss American Historical Society Review* 42.1 (2006), p. 2.

[84] John W Maxwell. "Tracing the Dynabook: A study of technocultural transformations." PhD thesis. University of British Columbia, 2006.

[85] Frank Steinicke, Timo Ropinski, and Klaus Hinrichs. "Object selection in virtual environments using an improved virtual pointer metaphor." In: *Computer Vision and Graphics: International Conference, ICCVG 2004, Warsaw, Poland, September 2004, Proceedings*. Springer. 2006, pp. 320–326.

[86] Michel Beaudouin-Lafon and Wendy E Mackay. “Prototyping tools and techniques.” In: *The human-computer interaction handbook*. CRC Press, 2007, pp. 1043–1066.

[87] Steve Harrison, Deborah Tatar, and Phoebe Sengers. “The three paradigms of HCI.” In: *Alt. Chi. Session at the SIGCHI Conference on human factors in computing systems San Jose, California, USA*. 2007, pp. 1–18.

[88] Jonas Lowgren and Erik Stolterman. *Thoughtful interaction design: A design perspective on information technology*. Mit Press, 2007.

[89] Dan R Olsen Jr. “Evaluating user interface systems research.” In: *Proceedings of the 20th annual ACM symposium on User interface software and technology*. 2007, pp. 251–258.

[90] Anne Adams, Peter Lunt, and Paul Cairns. “A qualititative approach to HCI research.” In: (2008).

[91] Olivier Bau and Wendy E. Mackay. “OctoPocus: A Dynamic Guide for Learning Gesture-Based Command Sets.” In: *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology*. UIST ’08. Monterey, CA, USA: ACM, 2008, 37–46. ISBN: 9781595939753. DOI: `10.1145/1449715.1449724`. URL: `https://doi.org/10.1145/1449715.1449724`.

[92] Hilda M Fehd and Adriane E Seiffert. “Eye movements during multiple object tracking: Where do participants look?” In: *Cognition* 108.1 (2008), pp. 201–209.

[93] CEL Green, D Freeman, E Kuipers, P Bebbington, D Fowler, G Dunn, and PA Garety. “Measuring ideas of persecution and social reference: the Green et al. Paranoid Thought Scales (GPTS).” In: *Psychological medicine* 38.1 (2008), pp. 101–111.

[94] Saul Greenberg and Bill Buxton. “Usability evaluation considered harmful (some of the time).” In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2008, pp. 111–120.

[95] Howell Istance, Richard Bates, Aulikki Hyrskykari, and Stephen Vickers. “Snap clutch, a moded approach to solving the Midas touch problem.” In: *Proceedings of the 2008 symposium on Eye tracking research & applications*. 2008, pp. 221–228.

[96] Robert J.K. Jacob, Audrey Girouard, Leanne M. Hirshfield, Michael S. Horn, Orit Shaer, Erin Treacy Solovey, and Jamie Zigelbaum. “Reality-based interaction: a framework for post-WIMP interfaces.” In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. Proceeding of the

twenty-sixth annual CHI conference. Florence, Italy: ACM, 2008, pp. 201–210. ISBN: 978-1-60558-011-1. DOI: 10.1145/1357054.1357089. URL: http://portal.acm.org/citation.cfm?doid=1357054.1357089 (visited on 08/09/2022).

[97] Youn-Kyung Lim, Erik Stolterman, and Josh Tenenberg. "The anatomy of prototypes: Prototypes as filters, prototypes as manifestations of design ideas." In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 15.2 (2008), pp. 1–27.

[98] Donald A. Norman. "THE WAY I SEE IT Signifiers, not affordances." In: *Interactions* 15.6 (Nov. 2008), pp. 18–19. ISSN: 1072-5520, 1558-3449. DOI: 10.1145/1409040.1409044. URL: https://dl.acm.org/doi/10.1145/1409040.1409044 (visited on 01/16/2023).

[99] Michael Tombu and Adriane E Seiffert. "Attentional costs in multiple-object tracking." In: *Cognition* 108.1 (2008), pp. 1–25.

[100] Christoph Bichlmeier, Sandro Michael Heining, Marco Feuerstein, and Nassir Navab. "The virtual mirror: a new interaction paradigm for augmented reality environments." In: *IEEE Transactions on Medical Imaging* 28.9 (2009), pp. 1498–1510.

[101] Mark Blythe and Paul Cairns. "Critical Methods and User Generated Content: The IPhone on YouTube." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '09. Boston, MA, USA: Association for Computing Machinery, 2009, 1467–1476. ISBN: 9781605582467. DOI: 10.1145/1518701.1518923. URL: https://doi.org/10.1145/1518701.1518923.

[102] Alex Cao, Keshav K Chintamani, Abhilash K Pandya, and R Darin Ellis. "NASA TLX: Software for assessing subjective mental workload." In: *Behavior research methods* 41.1 (2009), pp. 113–117.

[103] Jesse Fox and Jeremy N. Bailenson. "Virtual Self-Modeling: The Effects of Vicarious Reinforcement and Identification on Exercise Behaviors." In: *Media Psychology* 12.1 (Feb. 26, 2009), pp. 1–25. ISSN: 1521-3269, 1532-785X. DOI: 10.1080/15213260802669474. URL: http://www.tandfonline.com/doi/abs/10.1080/15213260802669474 (visited on 12/13/2022).

[104] Jesse Fox and Jeremy N Bailenson. "Virtual self-modeling: The effects of vicarious reinforcement and identification on exercise behaviors." In: *Media Psychology* 12.1 (2009), pp. 1–25.

[105] Jesse Fox, Jeremy Bailenson, and Joseph Binney. "Virtual experiences, physical behaviors: The effect of presence on imitation of an eating avatar." In: *Presence: Teleoperators and Virtual Environments* 18.4 (2009), pp. 294–303.

[106] Dustin Freeman, Hrvoje Benko, Meredith Ringel Morris, and Daniel Wigdor. "ShadowGuides: Visualizations for in-Situ Learning of Multi-Touch and Whole-Hand Gestures." In: *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*. ITS '09. Banff, Alberta, Canada: ACM, 2009, 165–172. ISBN: 9781605587332. DOI: `10.1145/1731903.1731935`. URL: `https://doi.org/10.1145/1731903.1731935`.

[107] Shahedul Huq Khandkar. "Open coding." In: *University of Calgary* 23 (2009), p. 2009.

[108] Brian Y Lim and Anind K Dey. "Assessing demand for intelligibility in context-aware applications." In: *Proceedings of the 11th international conference on Ubiquitous computing*. Ubicomp '09: The 11th International Conference on Ubiquitous Computing. Orlando Florida USA: ACM, Sept. 30, 2009, pp. 195–204. ISBN: 978-1-60558-431-7. DOI: `10.1145/1620545.1620576`. URL: `https://dl.acm.org/doi/10.1145/1620545.1620576` (visited on 01/16/2023).

[109] Edouard Machery. *Doing without concepts*. Oxford University Press, 2009.

[110] David Moher, Alessandro Liberati, Jennifer Tetzlaff, Douglas G Altman, and PRISMA Group*. "Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement." In: *Annals of internal medicine* 151.4 (2009), pp. 264–269.

[111] Bart Simon. "Wii are out of Control: Bodies, Game Screens and the Production of Gestural Excess." In: *SSRN Electronic Journal*. 2009. DOI: `http://dx.doi.org/10.2139/ssrn.1354043`.

[112] Jeremy N Bailenson and Kathryn Y Segovia. "Virtual doppelgangers: Psychological effects of avatars who ignore their owners." In: *Online worlds: Convergence of the real and the virtual*. Springer, 2010, pp. 175–186.

[113] Till Ballendat, Nicolai Marquardt, and Saul Greenberg. "Proxemic Interaction: Designing for a Proximity and Orientation-Aware Environment." In: *ACM International Conference on Interactive Tabletops and Surfaces*. ITS '10. Saarbrücken, Germany: ACM, 2010, 121–130. ISBN: 9781450303996. DOI: `10.1145/1936652.1936676`. URL: `https://doi.org/10.1145/1936652.1936676`.

[114] Jacky CP Chan, Howard Leung, Jeff KT Tang, and Taku Komura. “A virtual reality dance training system using motion capture technology.” In: *IEEE transactions on learning technologies* 4.2 (2010), pp. 187–195.

[115] Mar Gonzalez-Franco, Daniel Perez-Marcos, Bernhard Spanlang, and Mel Slater. “The contribution of real-time mirror reflections of motor actions on virtual body ownership in an immersive virtual environment.” In: *2010 IEEE virtual reality conference (VR)*. IEEE. 2010, pp. 111–114.

[116] Chin-Chang Ho and Karl F MacDorman. “Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices.” In: *Computers in Human Behavior* 26.6 (2010), pp. 1508–1518.

[117] Brian Y Lim and Anind K Dey. “Toolkit to support intelligibility in context-aware applications.” In: *Proceedings of the 12th ACM international conference on Ubiquitous computing*. Ubicomp ’10: The 2010 ACM Conference on Ubiquitous Computing. Copenhagen Denmark: ACM, Sept. 26, 2010, pp. 13–22. ISBN: 978-1-60558-843-8. DOI: `10.1145/1864349.1864353`. URL: `https://dl.acm.org/doi/10.1145/1864349.1864353` (visited on 12/15/2022).

[118] Donald A Norman. “Natural user interfaces are not natural.” In: *Interactions* 17.3 (May 2010), pp. 6–10. ISSN: 1072-5520, 1558-3449. DOI: `10.1145/1744161.1744163`. URL: `https://dl.acm.org/doi/10.1145/1744161.1744163` (visited on 12/02/2022).

[119] Erik Stolterman and Mikael Wiberg. “Concept-Driven Interaction Design Research.” In: *Human–Computer Interaction* 25.2 (May 28, 2010), pp. 95–118. ISSN: 0737-0024, 1532-7051. DOI: `10.1080/07370020903586696`. URL: `https://www.tandfonline.com/doi/full/10.1080/07370020903586696` (visited on 12/02/2022).

[120] S Shyam Sundar, Qian Xu, and Saraswathi Bellur. “Designing Interactivity in Media Interfaces: A Communications Perspective.” In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’10. Atlanta, Georgia, USA: ACM, 2010, 2247–2256. ISBN: 9781605589299. DOI: `10.1145/1753326.1753666`. URL: `https://doi.org/10.1145/1753326.1753666`.

[121] Xiao Xiao and Hiroshi Ishii. “MirrorFugue: communicating hand gesture in remote piano collaboration.” In: *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction*. 2010, pp. 13–20.

[122] Eunsuk Kang, Ethan Jackson, and Wolfram Schulte. "An approach for effective design space exploration." In: *Foundations of Computer Software. Modeling, Development, and Verification of Adaptive Systems: 16th Monterey Workshop 2010, Redmond, WA, USA, March 31-April 2, 2010, Revised Selected Papers 16*. Springer. 2011, pp. 33–54.

[123] Alan C. Kay. "A Personal Computer for Children of All Ages." In: *Proceedings of the ACM Annual Conference - Volume 1*. ACM '72. Boston, Massachusetts, USA: Association for Computing Machinery, 2011. ISBN: 9781450374910. DOI: 10.1145/800193.1971922. URL: https://doi.org/10.1145/800193.1971922.

[124] Tom Ritchey. *Wicked problems–social messes: Decision support modelling with morphological analysis*. Vol. 17. Springer Science & Business Media, 2011.

[125] Roland Aigner, Daniel Wigdor, Hrvoje Benko, Michael Haller, David Lindbauer, Alexandra Ion, Shengdong Zhao, and JTKV Koh. "Understanding Mid-Air Hand Gestures: A Study of Human Preferences in Usage of Gesture Types for HCI." In: *Microsoft Research TechReport MSR-TR-2012-111* 2.MSR-TR-2012-111 (Nov. 2012), p. 30. URL: https://www.microsoft.com/en-us/research/publication/understanding-mid-air-hand-gestures-a-study-of-human-preferences-in-usage-of-gesture-types-for-hci/.

[126] Tobias Blum, Valerie Kleeberger, Christoph Bichlmeier, and Nassir Navab. "mirracle: An augmented reality magic mirror system for anatomy education." In: *2012 IEEE Virtual Reality Workshops (VRW)*. IEEE. 2012, pp. 115–116.

[127] Nicola Dell, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies. ""Yours is better!": participant response bias in HCI." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12: CHI Conference on Human Factors in Computing Systems. Austin Texas USA: ACM, May 5, 2012, pp. 1321–1330. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2208589. URL: https://dl.acm.org/doi/10.1145/2207676.2208589 (visited on 01/16/2023).

[128] William Gaver. "What should we expect from research through design?" In: *Proceedings of the SIGCHI conference on human factors in computing systems*. 2012, pp. 937–946.

[129] Kristina Höök and Jonas Löwgren. "Strong concepts: Intermediate-level knowledge in interaction design research." In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 19.3 (2012), pp. 1–18.

[130] Mehdi Kaytoue, Arlei Silva, Loïc Cerf, Wagner Meira Jr, and Chedy Raïssi. "Watch me playing, i am a professional: a first study on video game live streaming." In: *Proceedings of the 21st international conference on world wide web*. 2012, pp. 1181–1188.

[131] Konstantina Kilteni, Raphaela Groten, and Mel Slater. "The sense of embodiment in virtual reality." In: *Presence: Teleoperators and Virtual Environments* 21.4 (2012), pp. 373–387.

[132] Nicolai Marquardt and Saul Greenberg. "Informing the Design of Proxemic Interactions." In: *IEEE Pervasive Computing* 11.2 (Feb. 2012), pp. 14–23. ISSN: 1536-1268. DOI: `10.1109/MPRV.2012.15`. URL: `http://ieeexplore.ieee.org/document/6127852/` (visited on 12/02/2022).

[133] Andrew YC Nee, SK Ong, George Chryssolouris, and Dimitris Mourtzis. "Augmented reality applications in design and manufacturing." In: *CIRP Annals* 61.2 (2012), pp. 657–679. ISSN: 00078506. DOI: `10.1016/j.cirp.2012.05.010`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0007850612002090`.

[134] Lance J Rips, Edward E Smith, and Douglas L Medin. "11 Concepts and Categories: Memory, Meaning, and Metaphysics." In: *The Oxford handbook of thinking and reasoning* (2012), p. 177.

[135] Rajinder Sodhi, Hrvoje Benko, and Andrew Wilson. "LightGuide: projected visualizations for hand movement guidance." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12: CHI Conference on Human Factors in Computing Systems. Austin Texas USA: ACM, May 5, 2012, pp. 179–188. ISBN: 978-1-4503-1015-4. DOI: `10.1145/2207676.2207702`. URL: `https://dl.acm.org/doi/10.1145/2207676.2207702` (visited on 12/01/2022).

[136] Fritz Zwicky. *Morphological astronomy*. Springer Science & Business Media, 2012.

[137] Fritz Zwicky and Albert George Wilson. *New methods of thought and procedure: Contributions to the symposium on methodologies*. Springer Science & Business Media, 2012.

[138] Lisa Anthony, YooJin Kim, and Leah Findlater. "Analyzing User-Generated Youtube Videos to Understand Touchscreen Use by People with Motor Impairments." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. Paris, France: Association for Computing Machinery, 2013, 1223–1232. ISBN: 9781450318990. DOI: `10.1145/2470654.2466158`. URL: `https://doi.org/10.1145/2470654.2466158`.

[139] Elizabeth Buie and Mark Blythe. "Meditations on YouTube." In: *Proceedings of the 6th International Conference on Designing Pleasurable Products and Interfaces*. DPPI '13. Newcastle upon Tyne, United Kingdom: Association for Computing Machinery, 2013, 41–50. ISBN: 9781450321921. DOI: `10.1145/2513506.2513511`. URL: `https://doi.org/10.1145/2513506.2513511`.

[140] Clayton M Christensen. *The innovator's dilemma: when new technologies cause great firms to fail*. Harvard Business Review Press, 2013.

[141] Richard Harper and Helena Mentis. "The Mocking Gaze: The Social Organization of Kinect Use." In: *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*. CSCW '13. San Antonio, Texas, USA: Association for Computing Machinery, 2013, 167–180. ISBN: 9781450313315. DOI: `10.1145/2441776.2441797`. URL: `https://doi.org/10.1145/2441776.2441797`.

[142] Robb Lindgren and Mina Johnson-Glenberg. "Emboldened by embodiment: Six precepts for research on embodied learning and mixed reality." In: *Educational researcher* 42.8 (2013), pp. 445–452.

[143] Antonella Maselli and Mel Slater. "The building blocks of the full body ownership illusion." In: *Frontiers in human neuroscience* 7 (2013), p. 83.

[144] Donald A. Norman. *The design of everyday things*. Revised and expanded edition. New York, New York: Basic Books, 2013. 347 pp. ISBN: 978-0-465-05065-9.

[145] Tobias Roehm, Nigar Gurbanova, Bernd Bruegge, Christophe Joubert, and Walid Maalej. "Monitoring user interactions for supporting failure reproduction." In: *2013 21st International Conference on Program Comprehension (ICPC)*. IEEE. 2013, pp. 73–82.

[146] Jonas Schild, Liane Bölicke, Joseph J LaViola Jr, and Maic Masuch. "Creating and analyzing stereoscopic 3D graphical user interfaces in digital games." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13: CHI Conference on Human Factors in Computing Systems. Paris France: ACM, Apr. 27, 2013, pp. 169–178. ISBN: 978-1-4503-1899-0. DOI: `10.1145/2470654.2470678`. URL: `https://dl.acm.org/doi/10.1145/2470654.2470678` (visited on 12/02/2022).

[147] Alistair Sutcliffe and Pete Sawyer. "Requirements elicitation: Towards the unknown unknowns." In: *2013 21st IEEE International Requirements Engineering Conference (RE)*. IEEE. 2013, pp. 92–104.

[148] Jo Vermeulen, Kris Luyten, Elise van den Hoven, and Karin Coninx. "Crossing the bridge over Norman's Gulf of Execution: revealing feedforward's true identity." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2013, pp. 1931–1940.

[149] Michael Mose Biskjaer, Peter Dalsgaard, and Kim Halskov. "A constraint-based understanding of design spaces." In: *Proceedings of the 2014 conference on Designing interactive systems*. 2014, pp. 453–462.

[150] Kristopher J Blom and Steffi Beckhaus. "The design space of dynamic interactive virtual environments." In: *Virtual Reality* 18.2 (June 2014), pp. 101–116. ISSN: 1359-4338, 1434-9957. DOI: `10.1007/s10055-013-0232-y`. URL: `http://link.springer.com/10.1007/s10055-013-0232-y`.

[151] Lung-Pan Cheng, Patrick Lühne, Pedro Lopes, Christoph Sterz, and Patrick Baudisch. "Haptic Turk: A Motion Platform Based on People." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. Toronto, Ontario, Canada: Association for Computing Machinery, 2014, 3463–3472. ISBN: 9781450324731. DOI: `10.1145/2556288.2557101`. URL: `https://doi.org/10.1145/2556288.2557101`.

[152] Peter Dalsgaard and Christian Dindler. "Between theory and practice: bridging concepts in HCI research." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14: CHI Conference on Human Factors in Computing Systems. Toronto Ontario Canada: ACM, Apr. 26, 2014, pp. 1635–1644. ISBN: 978-1-4503-2473-1. DOI: `10.1145/2556288.2557342`. URL: `https://dl.acm.org/doi/10.1145/2556288.2557342` (visited on 12/01/2022).

[153] J.J. Gibson. *The Ecological Approach to Visual Perception: Classic Edition*. Psychology Press & Routledge Classic Editions. Taylor & Francis, 2014. ISBN: 978-1-317-57937-3. URL: `https://books.google.dk/books?id=QReLBQAAQBAJ`.

[154] Shuai Hao, Bin Liu, Suman Nath, William GJ Halfond, and Ramesh Govindan. "Puma: Programmable ui-automation for large-scale dynamic analysis of mobile apps." In: *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*. 2014, pp. 204–217.

[155] Dennis L. Kappen, Pejman Mirza-Babaei, Jens Johannsmeier, Daniel Buckstein, James Robb, and Lennart E. Nacke. "Engaged by Boos and Cheers: The Effect of Co-Located Game Audiences on Social Player Experience." In: *Proceedings of the First ACM SIGCHI Annual Symposium on Computer-Human Interaction in Play*. CHI PLAY '14. Toronto, Ontario, Canada: Association for Computing Machinery, 2014, 151–160. ISBN: 9781450330145. DOI: `10.1145/2658537.2658687`. URL: `https://doi.org/10.1145/2658537.2658687`.

[156] André Kunert, Alexander Kulik, Stephan Beck, and Bernd Froehlich. "Photoportals: shared references in space and time." In: *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 2014, pp. 1388–1399.

[157] Kari Kuutti and Liam J Bannon. "The turn to practice in HCI: towards a research agenda." In: *Proceedings of the SIGCHI conference on human factors in computing systems*. CHI '14: CHI Conference on Human Factors in Computing Systems. Toronto Ontario Canada: ACM, Apr. 26, 2014, pp. 3543–3552. ISBN: 978-1-4503-2473-1. DOI: `10.1145/2556288.2557111`. URL: `https://dl.acm.org/doi/10.1145/2556288.2557111` (visited on 01/16/2023).

[158] Mathieu Nancel and Andy Cockburn. "Causality: A conceptual model of interaction history." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2014, pp. 1777–1786.

[159] Asunción Álvarez and Tom Ritchey. "Applications of general morphological analysis." In: *Acta Morphologica Generalis* 4.1 (2015).

[160] Lung-Pan Cheng, Thijs Roumen, Hannes Rantzsch, Sven Köhler, Patrick Schmidt, Robert Kovacs, Johannes Jasper, Jonas Kemper, and Patrick Baudisch. "TurkDeck: Physical Virtual Reality Based on People." In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software I& Technology*. UIST '15. Charlotte, NC, USA: Association for Computing Machinery, 2015, 417–426. ISBN: 9781450337793. DOI: `10.1145/2807442.2807463`. URL: `https://doi.org/10.1145/2807442.2807463`.

[161] Nicholas S. Dalton, Emily Collins, and Paul Marshall. "Display Blindness? Looking Again at the Visibility of Situated Displays Using Eye-Tracking." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Seoul, Republic of Korea: Association for Computing Machinery, 2015, 3889–3898. ISBN: 9781450331456.

DOI: 10.1145/2702123.2702150. URL: https://doi.org/10.1145/2702123.2702150.

[162] William Delamare, Céline Coutrix, and Laurence Nigay. “Designing guiding systems for gesture-based interaction.” In: *Proceedings of the 7th ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. EICS’15: ACM SIGCHI Symposium on Engineering Interactive Computing Systems. Duisburg Germany: ACM, June 23, 2015, pp. 44–53. ISBN: 978-1-4503-3646-8. DOI: 10.1145/2774225.2774847. URL: https://dl.acm.org/doi/10.1145/2774225.2774847 (visited on 12/02/2022).

[163] Juan Pablo Hourcade, Sarah L. Mascher, David Wu, and Luiza Pantoja. “Look, My Baby Is Using an IPad! An Analysis of YouTube Videos of Infants and Toddlers Using Tablets.” In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI ’15. Seoul, Republic of Korea: Association for Computing Machinery, 2015, 1915–1924. ISBN: 9781450331456. DOI: 10.1145/2702123.2702266. URL: https://doi.org/10.1145/2702123.2702266.

[164] Pedro Lopes, Patrik Jonell, and Patrick Baudisch. “Affordance++: Allowing Objects to Communicate Dynamic Use.” In: *Proceedings of the 33rd annual acm conference on human factors in computing systems*. CHI ’15. Seoul, Republic of Korea: ACM, 2015, pp. 2515–2524. ISBN: 9781450331456. DOI: 10.1145/2702123.2702128. URL: https://doi.org/10.1145/2702123.2702128.

[165] Andrés Lucero. “Using affinity diagrams to evaluate interactive prototypes.” In: *IFIP conference on human-computer interaction*. Springer. 2015, pp. 231–248.

[166] Mark McGill, Daniel Boland, Roderick Murray-Smith, and Stephen Brewster. “A Dose of Reality: Overcoming Usability Challenges in VR Head-Mounted Displays.” In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI ’15: CHI Conference on Human Factors in Computing Systems. Seoul Republic of Korea: ACM, Apr. 18, 2015, pp. 2143–2152. ISBN: 978-1-4503-3145-6. DOI: 10.1145/2702123.2702382. URL: https://dl.acm.org/doi/10.1145/2702123.2702382 (visited on 01/16/2023).

[167] Ohan Oda, Carmine Elvezio, Mengu Sukan, Steven Feiner, and Barbara Tversky. “Virtual replicas for remote assistance in virtual and augmented reality.” In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. 2015, pp. 405–415.

[168] Sofia Adelaide Osimo, Rodrigo Pizarro, Bernhard Spanlang, and Mel Slater. "Conversations between self and self as Sigmund Freud—A virtual body ownership paradigm for self counselling." In: *Scientific reports* 5.1 (2015), pp. 1–14.

[169] Stephanie Atherton, Angus Antley, Nicole Evans, Emma Cernis, Rachel Lister, Graham Dunn, Mel Slater, and Daniel Freeman. "Self-confidence and paranoia: an experimental study using an immersive virtual reality social situation." In: *Behavioural and cognitive psychotherapy* 44.1 (2016), pp. 56–64.

[170] Jeremy Boy, Louis Eveillard, Françoise Detienne, and Jean-Daniel Fekete. "Suggested Interactivity: Seeking Perceived Affordances for Information Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (Jan. 31, 2016), pp. 639–648. ISSN: 1077-2626. DOI: `10.1109/TVCG.2015.2467201`. URL: `http://ieeexplore.ieee.org/document/7192637/` (visited on 01/16/2023).

[171] Peter Dalsgaard, Kim Halskov, and Ole Sejer Iversen. "Participation Gestalt: Analysing Participatory Qualities of Interaction in Public Space." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: Association for Computing Machinery, 2016, 4435–4446. ISBN: 9781450333627. DOI: `10.1145/2858036.2858147`. URL: `https://doi.org/10.1145/2858036.2858147`.

[172] William Delamare, Thomas Janssoone, Céline Coutrix, and Laurence Nigay. "Designing 3D Gesture Guidance: Visual Feedback and Feedforward Design Options." en. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces*. Bari Italy: ACM, June 2016, pp. 152–159. ISBN: 978-1-4503-4131-8. DOI: `10.1145/2909132.2909260`. URL: `https://dl.acm.org/doi/10.1145/2909132.2909260` (visited on 12/02/2022).

[173] Graham Dove, Nicolai Brodersen Hansen, and Kim Halskov. "An Argument For Design Space Reflection." In: *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. NordiCHI '16. Gothenburg, Sweden: Association for Computing Machinery, 2016. ISBN: 9781450347631. DOI: `10.1145/2971485.2971528`. URL: `https://doi.org/10.1145/2971485.2971528`.

[174] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M Dollar, and Danica Kragic. "The grasp taxonomy of human grasp types." In: *IEEE Transactions on Human-Machine Systems* 46.1 (Feb. 2016), pp. 66–77. ISSN: 2168-2291,

2168-2305. DOI: 10.1109/THMS.2015.2470657. URL: http://ieeexplore.ieee.org/document/7243327/ (visited on 12/02/2022).

[175] Marco Gillies. "What is Movement Interaction in Virtual Reality for?" In: *Proceedings of the 3rd International Symposium on Movement and Computing*. MOCO'16: 3rd International Symposium on Movement and Computing. Thessaloniki GA Greece: ACM, July 5, 2016, pp. 1–4. ISBN: 978-1-4503-4307-7. DOI: 10.1145/2948910.2948951. URL: https://dl.acm.org/doi/10.1145/2948910.2948951 (visited on 12/14/2022).

[176] Jason Jerald. *The VR book: Human-centered design for virtual reality*. ACM books 8. New York] [San Rafael, California: ACM, Morgan & Claypool publishers, 2016. ISBN: 978-1-970001-12-9 978-1-970001-15-0 978-1-970001-13-6.

[177] Lorraine Lin and Sophie Jörg. "Need a hand? How appearance affects the virtual hand illusion." In: *Proceedings of the ACM Symposium on Applied Perception*. SAP '16: ACM Symposium on Applied Perception 2016. Anaheim California: ACM, July 22, 2016, pp. 69–76. ISBN: 978-1-4503-4383-1. DOI: 10.1145/2931002.2931006. URL: https://dl.acm.org/doi/10.1145/2931002.2931006 (visited on 10/25/2021).

[178] Keigo Matsumoto, Yuki Ban, Takuji Narumi, Yohei Yanase, Tomohiro Tanikawa, and Michitaka Hirose. "Unlimited corridor: redirected walking techniques using visuo haptic interaction." In: *ACM SIGGRAPH 2016 Emerging Technologies*. ACM. 2016, p. 20.

[179] Sergio Orts-Escolano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yury Degtyarev, David Kim, Philip L Davidson, Sameh Khamis, Mingsong Dou, et al. "Holoportation: Virtual 3d teleportation in real-time." In: *Proceedings of the 29th annual symposium on user interface software and technology*. 2016, pp. 741–754.

[180] Antti Oulasvirta and Kasper Hornbæk. "HCI research as problem-solving." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2016, pp. 4956–4967.

[181] Mel Slater and Maria V Sanchez-Vives. "Enhancing our lives with immersive virtual reality." In: *Frontiers in Robotics and AI* 3 (2016), p. 74.

[182] Anthony Steed, Sebastian Frlston, María Murcia Lopez, Jason Drummond, Ye Pan, and David Swapp. "An 'In the Wild' Experiment on Presence and Embodiment using Consumer Virtual Reality Equipment." In: *IEEE Transactions on*

Visualization and Computer Graphics 22.4 (2016), pp. 1406–1414. DOI: 10.1109/TVCG.2016.2518135.

[183] Frank Steinicke. *Being really virtual*. Springer, 2016.

[184] Jacob O Wobbrock and Julie A Kientz. “Research contributions in human-computer interaction.” In: *interactions* 23.3 (2016), pp. 38–44.

[185] Cristina Botella, Javier Fernández-Álvarez, Verónica Guillén, Azucena García-Palacios, and Rosa Baños. “Recent progress in virtual reality exposure therapy for phobias: a systematic review.” In: *Current psychiatry reports* 19 (2017), pp. 1–13.

[186] Jan Gugenheimer, Evgeny Stemasov, Julian Frommel, and Enrico Rukzio. “Sharevr: Enabling co-located experiences for virtual reality between hmd and non-hmd users.” In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2017, pp. 4021–4033.

[187] Kasper Hornbæk and Antti Oulasvirta. “What is interaction?” In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2017, pp. 5040–5052.

[188] Jaron Lanier. *Dawn of the new everything: Encounters with reality and virtual reality*. Henry Holt and Company, 2017.

[189] Thomas Lopez, Olivier Dumas, Fabien Danieau, Bertrand Leroy, Nicolas Mollet, and Jean-François Vial. “A playback tool for reviewing VR experiences.” In: *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*. VRST ’17: 23rd ACM Symposium on Virtual Reality Software and Technology. Gothenburg Sweden: ACM, Nov. 8, 2017, pp. 1–2. ISBN: 978-1-4503-5548-3. DOI: 10.1145/3139131.3141776. URL: https://dl.acm.org/doi/10.1145/3139131.3141776 (visited on 05/20/2022).

[190] Aske Mottelson and Kasper Hornbæk. “Virtual Reality Studies Outside the Laboratory.” In: *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*. VRST ’17. Gothenburg, Sweden: Association for Computing Machinery, 2017. ISBN: 9781450355483. DOI: 10.1145/3139131.3139141. URL: https://doi.org/10.1145/3139131.3139141.

[191] Katherine Payne, Mark J Keith, Ryan M Schuetzler, and Justin Scott Giboney. “Examining the learning effects of live streaming video game instruction over Twitch.” In: *Computers in Human Behavior* 77 (2017), pp. 95–109.

[192] Ken Pfeuffer, Benedikt Mayer, Diako Mardanbegi, and Hans Gellersen. “Gaze+ pinch interaction in virtual reality.” In: *Proceedings of the 5th symposium on spatial user interaction*. 2017, pp. 99–108.

[193] Dong-Hee Shin. "The role of affordance in the experience of virtual reality learning: Technological and affective affordances in virtual reality." In: *Telematics and Informatics* 34.8 (2017), pp. 1826–1836. ISSN: 07365853. DOI: 10.1016/j.tele.2017.05.013. URL: https://linkinghub.elsevier.com/retrieve/pii/S0736585317301223.

[194] Jan-Philipp Stein and Peter Ohler. "Venturing into the uncanny valley of mind—The influence of mind attribution on the acceptance of human-like characters in a virtual reality setting." In: *Cognition* 160 (2017), pp. 43–50.

[195] Burak S. Tekin and Stuart Reeves. "Ways of Spectating: Unravelling Spectator Participation in Kinect Play." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: Association for Computing Machinery, 2017, 1558–1570. ISBN: 9781450346559. DOI: 10.1145/3025453.3025813. URL: https://doi.org/10.1145/3025453.3025813.

[196] Rafael Ballagas, Sarthak Ghosh, and James Landay. "The Design Space of 3D Printable Interactivity." In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2.2 (July 5, 2018), pp. 1–21. ISSN: 2474-9567. DOI: 10.1145/3214264. URL: https://dl.acm.org/doi/10.1145/3214264 (visited on 12/01/2022).

[197] Soumya C Barathi, Daniel J Finnegan, Matthew Farrow, Alexander Whaley, Pippa Heath, Jude Buckley, Peter W Dowrick, Burkhard C Wuensche, James LJ Bilzon, Eamonn O'Neill, et al. "Interactive feedforward for improving performance and maintaining intrinsic motivation in VR exergaming." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–14.

[198] Jordan Beck and Hamid R Ekbia. "The theory-practice gap as generative metaphor." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–11.

[199] Kody R Dillman, Terrance Tin Hoi Mok, Anthony Tang, Lora Oehlberg, and Alex Mitchell. "A Visual Interaction Cue Framework from Video Game Environments for Augmented Reality." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18: CHI Conference on Human Factors in Computing Systems. Montreal QC Canada: ACM, Apr. 19, 2018, pp. 1–12. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173714. URL: https://dl.acm.org/doi/10.1145/3173574.3173714 (visited on 12/16/2021).

[200] Mar Gonzalez-Franco and Tabitha C Peck. "Avatar embodiment. towards a standardized questionnaire." In: *Frontiers in Robotics and AI* 5 (2018), p. 74.

[201] Melissa E Hudson, Karen S Voytecki, and Guili Zhang. "Mixed-reality teaching experiences improve preservice special education students' perceptions of their ability to manage a classroom." In: *Journal for Virtual Worlds Research* 11.2 (2018).

[202] Jarrod Knibbe, Jonas Schjerlund, Mathias Petraeus, and Kasper Hornbæk. "The Dream is Collapsing: The Experience of Exiting VR." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, 1–13. ISBN: 9781450356206. DOI: `10.1145/3173574.3174057`. URL: `https://doi.org/10.1145/3173574.3174057`.

[203] David Lindlbauer and Andy D Wilson. "Remixed reality: Manipulating space and time in augmented reality." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–13.

[204] Christian Mai and Mohamed Khamis. "Public HMDs: Modeling and Understanding User Behavior around Public Head-Mounted Displays." In: *Proceedings of the 7th ACM International Symposium on Pervasive Displays*. PerDis '18. Munich, Germany: Association for Computing Machinery, 2018. ISBN: 9781450357654. DOI: `10.1145/3205873.3205879`. URL: `https://doi.org/10.1145/3205873.3205879`.

[205] Michael Nebeling and Maximilian Speicher. "The trouble with augmented reality/virtual reality authoring tools." In: *2018 IEEE international symposium on mixed and augmented reality adjunct (ISMAR-Adjunct)*. IEEE. 2018, pp. 333–337.

[206] Thammathip Piumsomboon, Gun A Lee, Jonathon D Hart, Barrett Ens, Robert W Lindeman, Bruce H Thomas, and Mark Billinghurst. "Mini-me: An adaptive avatar for mixed reality remote collaboration." In: *Proceedings of the 2018 CHI conference on human factors in computing systems*. 2018, pp. 1–13.

[207] Daniel Roth, Constantin Klelnbeck, Tobias Feigl, Christopher Mutschler, and Marc Erich Latoschik. "Beyond replication: Augmenting social behaviors in multi-user virtual realities." In: *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE. 2018, pp. 215–222.

[208] Misha Sra, Aske Mottelson, and Pattie Maes. "Your Place and Mine: Designing a Shared VR Experience for Remotely Located Users." In: *Proceedings of the 2018 Designing In-*

teractive Systems Conference. DIS '18. Hong Kong, China: Association for Computing Machinery, 2018, 85–97. ISBN: 9781450351980. DOI: 10.1145/3196709.3196788. URL: https://doi.org/10.1145/3196709.3196788.

[209] Misha Sra, Xuhai Xu, Aske Mottelson, and Pattie Maes. "VMotion: Designing a Seamless Walking Experience in VR." In: *Proceedings of the 2018 Designing Interactive Systems Conference*. DIS '18. Hong Kong, China: Association for Computing Machinery, 2018, 59–70. ISBN: 9781450351980. DOI: 10.1145/3196709.3196792. URL: https://doi.org/10.1145/3196709.3196792.

[210] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. "Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: ACM, 2018, 86:1–86:12. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173660. URL: http://doi.acm.org/10.1145/3173574.3173660.

[211] Haijun Xia, Sebastian Herscher, Ken Perlin, and Daniel Wigdor. "Spacetime: Enabling fluid individual and collaborative editing in virtual reality." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 2018, pp. 853–866.

[212] Xiuming Zhang, Tali Dekel, Tianfan Xue, Andrew Owens, Qiurui He, Jiajun Wu, Stefanie Mueller, and William T Freeman. "Mosculp: Interactive visualization of shape and time." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 2018, pp. 275–285.

[213] Joanna Bergström, Aske Mottelson, Andreea Muresan, and Kasper Hornbæk. "Tool extension in human-computer interaction." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 2019, pp. 1–11.

[214] Yuanzhi Cao, Tianyi Wang, Xun Qian, Pawan S Rao, Manav Wadhawan, Ke Huo, and Karthik Ramani. "GhostAR: A time-space editor for embodied authoring of human-robot collaborative task with augmented reality." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 2019, pp. 521–534.

[215] Sven Coppers, Kris Luyten, Davy Vanacken, David Navarre, Philippe Palanque, and Christine Gris. "Fortunettes: Feedforward about the Future State of GUI Widgets." In: *Proceedings of the ACM on Human-Computer Interaction* 3.EICS (EICS June 13, 2019), pp. 1–20. ISSN: 2573-0142. DOI: 10.

1145/3331162. URL: https://dl.acm.org/doi/10.1145/3331162 (visited on 01/16/2023).

[216] Ceenu George, Julia Schwuchow, and Heinrich Hussmann. "Fearing Disengagement from the Real World." In: *25th ACM Symposium on Virtual Reality Software and Technology*. VRST '19. Parramatta, NSW, Australia: Association for Computing Machinery, 2019. ISBN: 9781450370011. DOI: 10.1145/3359996.3364273. URL: https://doi.org/10.1145/3359996.3364273.

[217] Kim Halskov and Aron Fischel. "The design space of media architecture displays." In: *Interactions* 26.6 (2019), pp. 60–63.

[218] Jeremy Hartmann, Christian Holz, Eyal Ofek, and Andrew D Wilson. "RealityCheck: Blending virtual environments with situated physical reality." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 2019, pp. 1–12.

[219] Masakazu Hirota, Hiroyuki Kanda, Takao Endo, Tomomitsu Miyoshi, Suguru Miyagawa, Yoko Hirohara, Tatsuo Yamaguchi, Makoto Saika, Takeshi Morimoto, and Takashi Fujikado. "Comparison of visual fatigue caused by head-mounted display for virtual reality and two-dimensional display using objective and subjective evaluation." In: *Ergonomics* 62.6 (2019), pp. 759–766.

[220] Teresa Hirzle, Jan Gugenheimer, Florian Geiselhart, Andreas Bulling, and Enrico Rukzio. "A Design Space for Gaze Interaction on Head-Mounted Displays." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300855. URL: https://doi.org/10.1145/3290605.3300855.

[221] ISO. *Ergonomics of human-system interaction — Part 210: Human-centred design for interactive systems*. en. Standard. Geneva, CH: International Organization for Standardization, 2019. URL: https://www.iso.org/obp/ui/en/#iso:std:iso:9241:-210:ed-2:v1:en.

[222] Kathleen Ingraham, Charles E Hughes, Lindsay A Taliaferro, Nicholas J Westers, Lisa Dieker, and Michael Hynes. "Using digital puppetry to prepare physicians to address non-suicidal self-injury among teens." In: *Universal Access in Human-Computer Interaction. Theory, Methods and Tools: 13th International Conference, UAHCI 2019, Held as Part of the 21st HCI International Conference, HCII 2019, Orlando, FL,*

USA, July 26–31, 2019, Proceedings, Part I 21. Springer. 2019, pp. 555–568.

[223] Aida Komkaite, Liga Lavrinovica, Maria Vraka, and Mikael B. Skov. "Underneath the Skin: An Analysis of YouTube Videos to Understand Insertable Device Interaction." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300444. URL: https://doi.org/10.1145/3290605.3300444.

[224] Andrey Krekhov, Sebastian Cmentowski, and Jens Krüger. "The Illusion of Animal Body Ownership and Its Potential for Virtual Reality Games." In: *2019 IEEE Conference on Games (CoG)*. 2019, pp. 1–8. DOI: 10.1109/CIG.2019.8848005.

[225] Jaeyeon Lee, Mike Sinclair, Mar Gonzalez-Franco, Eyal Ofek, and Christian Holz. "TORC: A Virtual Reality Controller for In-Hand High-Dexterity Finger Interaction." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: ACM, 2019, 71:1–71:13. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300301. URL: http://doi.acm.org/10.1145/3290605.3300301.

[226] Andreea Muresan and Henning Pohl. "Chats with bots: Balancing imitation and engagement." In: *Extended abstracts of the 2019 CHI conference on human factors in computing systems*. 2019, pp. 1–6.

[227] Duc-Minh Pham and Wolfgang Stuerzlinger. "Is the Pen Mightier than the Controller? A Comparison of Input Devices for Selection in Virtual and Augmented Reality." In: *25th ACM Symposium on Virtual Reality Software and Technology*. VRST '19: 25th ACM Symposium on Virtual Reality Software and Technology. Parramatta NSW Australia: ACM, Nov. 12, 2019, pp. 1–11. ISBN: 978-1-4503-7001-1. DOI: 10.1145/3359996.3364264. URL: https://dl.acm.org/doi/10.1145/3359996.3364264 (visited on 12/02/2022).

[228] Henning Pohl, Andreea Muresan, and Kasper Hornbæk. "Charting Subtle Interaction in the HCI Literature." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19: CHI Conference on Human Factors in Computing Systems. Glasgow Scotland Uk: ACM, May 2, 2019, pp. 1–15. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300648. URL: https://dl.acm.org/doi/10.1145/3290605.3300648.

[229] Daniel Roth, Gary Bente, Peter Kullmann, David Mal, Chris Felix Purps, Kai Vogeley, and Marc Erich Latoschik. “Technologies for Social Augmentations in User-Embodied Virtual Reality.” In: *Proceedings of the 25th ACM Symposium on Virtual Reality Software and Technology*. VRST ’19. Parramatta, NSW, Australia: Association for Computing Machinery, 2019. ISBN: 9781450370011. DOI: 10.1145/3359996.3364269. URL: https://doi.org/10.1145/3359996.3364269.

[230] Mel Slater, Solène Neyret, Tania Johnston, Guillermo Iruretagoyena, Mercè Álvarez de la Campa Crespo, Miquel Alabèrnia-Segura, Bernhard Spanlang, and Guillem Feixas. “An experimental study of a virtual reality counselling paradigm using embodied self-dialogue.” In: *Scientific reports* 9.1 (2019), p. 10903.

[231] Maximilian Speicher, Brian D Hall, and Michael Nebeling. “What is mixed reality?” In: *Proceedings of the 2019 CHI conference on human factors in computing systems*. 2019, pp. 1–15.

[232] Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, and Björn Hartmann. “TutoriVR: A Video-Based Tutorial System for Design Applications in Virtual Reality.” In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI ’19: CHI Conference on Human Factors in Computing Systems. Glasgow Scotland Uk: ACM, May 2, 2019, pp. 1–12. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300514. URL: https://dl.acm.org/doi/10.1145/3290605.3300514 (visited on 05/14/2022).

[233] Julius von Willich, Markus Funk, Florian Müller, Karola Marky, Jan Riemann, and Max Mühlhäuser. “You Invaded my Tracking Space! Using Augmented Virtuality for Spotting Passersby in Room-Scale Virtual Reality.” In: *Proceedings of the 2019 on Designing Interactive Systems Conference*. 2019, pp. 487–496.

[234] Yu Xiao and Maria Watson. “Guidance on conducting a systematic literature review.” In: *Journal of planning education and research* 39.1 (2019), pp. 93–112.

[235] Vero Vanden Abeele, Katta Spiel, Lennart Nacke, Daniel Johnson, and Kathrin Gerling. “Development and validation of the player experience inventory: A scale to measure player experiences at the level of functional and psychosocial consequences.” In: *International Journal of Human-Computer Studies* 135 (2020), p. 102370.

[236] Narges Ashtari, Andrea Bunt, Joanna McGrenere, Michael Nebeling, and Parmit K Chilana. "Creating augmented and virtual reality applications: Current practices, challenges, and opportunities." In: *Proceedings of the 2020 CHI conference on human factors in computing systems*. 2020, pp. 1–13.

[237] Tobias Drey, Jan Gugenheimer, Julian Karlbauer, Maximilian Milo, and Enrico Rukzio. "Vrsketchin: Exploring the design space of pen and tablet interaction for 3d sketching in virtual reality." In: *Proceedings of the 2020 CHI conference on human factors in computing systems*. 2020, pp. 1–14.

[238] Jeremy Hartmann, Yen-Ting Yeh, and Daniel Vogel. "AAR: Augmenting a Wearable Augmented Reality Display with an Actuated Head-Mounted Projector." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. Virtual Event, USA: Association for Computing Machinery, 2020, 445–458. ISBN: 9781450375146. DOI: `10.1145/3379337.3415849`. URL: `https://doi-org.ezproxy.lib.monash.edu.au/10.1145/3379337.3415849`.

[239] Xinyu Hu, Alec G Moore, James Coleman Eubanks, Afham Ahmed Aiyaz, and Ryan P McMahan. "The Effects of Delayed Interaction Cues in Virtual Reality Training." In: *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW). IEEE. Atlanta, GA, USA: IEEE, Mar. 2020, pp. 63–69. ISBN: 978-1-72816-532-5. DOI: `10.1109/VRW50115.2020.00019`. URL: `https://ieeexplore.ieee.org/document/9090424/` (visited on 12/16/2021).

[240] Pascal Jansen, Fabian Fischbach, Jan Gugenheimer, Evgeny Stemasov, Julian Frommel, and Enrico Rukzio. "ShARe: Enabling Co-Located Asymmetric Multi-User Interaction for Augmented Reality Head-Mounted Displays." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. Virtual Event, USA: Association for Computing Machinery, 2020, 459–471. ISBN: 9781450375146. DOI: `10.1145/3379337.3415843`. URL: `https://doi.org/10.1145/3379337.3415843`.

[241] Hans-Christian Jetter, Roman Rädle, Tiare Feuchtner, Christoph Anthes, Judith Friedl, and Clemens Nylandsted Klokmose. "" in vr, everything is possible!": Sketching and simulating spatially-aware interactive spaces in virtual reality." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–16.

[242] Árni Kristjánsson and Howard Egeth. "How feature integration theory integrated cognitive psychology, neurophysiology, and psychophysics." In: *Attention, Perception, & Psychophysics* 82.1 (2020), pp. 7–23.

[243] Jintae Lee and Kum-Yew Lai. "What's in design rationale?" In: *Design Rationale*. CRC Press, 2020, pp. 21–51.

[244] Germán Leiva, Cuong Nguyen, Rubaiat Habib Kazi, and Paul Asente. "Pronto: Rapid augmented reality video prototyping using sketches and enaction." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–13.

[245] Klemen Lilija, Henning Pohl, and Kasper Hornbæk. "Who put that there? temporal navigation of spatial recordings by direct manipulation." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–11.

[246] Juan Liu, Yawen Zheng, Ke Wang, Yulong Bian, Wei Gai, and Dingyuan Gao. "A real-time interactive tai chi learning system based on vr and motion capture technology." In: *Procedia Computer Science* 174 (2020), pp. 712–719.

[247] Thomas P Moran and John M Carroll. *Design rationale: Concepts, techniques, and use*. CRC Press, 2020.

[248] Arnaud Prouzeau, Yuchen Wang, Barrett Ens, Wesley Willett, and Tim Dwyer. "Corsican twin: Authoring in situ augmented reality visualisations in virtual reality." In: *Proceedings of the International Conference on Advanced Visual Interfaces*. 2020, pp. 1–9.

[249] Said A Salloum, Rehan Khan, and Khaled Shaalan. "A survey of semantic analysis approaches." In: *Proceedings of the International Conference on Artificial Intelligence and Computer Vision (AICV2020)*. Springer. 2020, pp. 61–70.

[250] Mel Slater et al. "The Ethics of Realism in Virtual and Augmented Reality." In: *Frontiers in Virtual Reality* 1 (2020), p. 1. ISSN: 2673-4192. DOI: `10.3389/frvir.2020.00001`. URL: `https://www.frontiersin.org/articles/10.3389/frvir.2020.00001`.

[251] Cosimo Tuena, Elisa Pedroli, Pietro Davide Trimarchi, Alessia Gallucci, Mattia Chiappini, Karine Goulene, Andrea Gaggioli, Giuseppe Riva, Fabrizia Lattanzio, Fabrizio Giunco, et al. "Usability issues of clinical and research applications of virtual reality in older people: A systematic review." In: *Frontiers in Human Neuroscience* 14 (2020), p. 93.

[252] Chiu-Hsuan Wang, Seraphina Yong, Hsin-Yu Chen, Hsin-Yu Ye, and Liwei Chan. "HMD Light: Sharing In-VR Experience via Head-Mounted Projector for Asymmetric Interaction." In: *To Appear in Proceedings of the 33rd Annual ACM Symposium on User Interface Software I& Technology*. UIST '20. New York, NY, USA: Association for Computing Machinery, 2020.

[253] Chunxue Wei, Difeng Yu, and Tilman Dingler. "Reading on 3D Surfaces in Virtual Environments." In: *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR). IEEE. Atlanta, GA, USA: IEEE, Mar. 2020, pp. 721–728. ISBN: 978-1-72815-608-8. DOI: `10.1109/VR46266.2020.00095`. URL: `https://ieeexplore.ieee.org/document/9089461/` (visited on 01/16/2023).

[254] Amal Yassien, Passant ElAgroudy, Elhassan Makled, and Slim Abdennadher. "A Design Space for Social Presence in VR." In: *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. NordiCHI '20: Shaping Experiences, Shaping Society. Tallinn Estonia: ACM, Oct. 25, 2020, pp. 1–12. ISBN: 978-1-4503-7579-5. DOI: `10.1145/3419249.3420112`. URL: `https://dl.acm.org/doi/10.1145/3419249.3420112` (visited on 12/01/2022).

[255] Christopher Ball, Kuo-Ting Huang, and Jess Francis. "Virtual reality adoption during the COVID-19 pandemic: A uses and gratifications perspective." In: *Telematics and Informatics* 65 (2021), p. 101728.

[256] Neil T Banerjee, Alex J Baughman, Shu-Yu Lin, Zoë A Witte, David M Klaus, and Allison P Anderson. "Development of alternative reality environments for spacecraft habitat design evaluation." In: *Virtual Reality* 25.2 (2021), pp. 399–408.

[257] Michel Beaudouin-Lafon, Susanne Bødker, and Wendy E Mackay. "Generative theories of interaction." In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 28.6 (2021), pp. 1–54.

[258] Rhonda Bondie, Zid Mancenido, and Chris Dede. "Interaction principles for digital puppeteering to promote teacher learning." In: *Journal of Research on Technology in Education* 53.1 (2021), pp. 107–123.

[259] Nathalie Bressa, Henrik Korsgaard, Aurélien Tabard, Steven Houben, and Jo Vermeulen. "What's the Situation with Situated Visualization? A Survey and Perspectives on Situatedness." In: *IEEE Transactions on Visualization and Computer Graphics* 28.1 (2021), pp. 107–117.

[260] Wolfgang Büschel, Anke Lehmann, and Raimund Dachselt. "Miria: A mixed reality toolkit for the in-situ visualization and analysis of spatio-temporal interaction data." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–15.

[261] Kurtis Danyluk, Barrett Ens, Bernhard Jenny, and Wesley Willett. "A design space exploration of worlds in miniature." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–15.

[262] Emily Dao, Andreea Muresan, Kasper Hornbæk, and Jarrod Knibbe. "Bad Breakdowns, Useful Seams, and Face Slapping: Analysis of VR Fails on YouTube." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21: CHI Conference on Human Factors in Computing Systems. New York, NY, USA: ACM, May 6, 2021, pp. 1–14. ISBN: 978-1-4503-8096-6. DOI: `10.1145/3411764.3445435`. URL: `https://doi.org/10.1145/3411764.3445435` (visited on 01/16/2023).

[263] Emily Dao, Andreea Muresan, Kasper Hornbæk, and Jarrod Knibbe. "Bad breakdowns, useful seams, and face slapping: Analysis of vr fails on youtube." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–14.

[264] Barrett Ens, Benjamin Bach, Maxime Cordeil, Ulrich Engelke, Marcos Serrano, Wesley Willett, Arnaud Prouzeau, Christoph Anthes, Wolfgang Büschel, Cody Dunne, et al. "Grand challenges in immersive analytics." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–17.

[265] Eva Eriksson, Gökçe Elif Baykal, Olof Torgersson, and Staffan Bjork. "The CoCe design space: Exploring the design space for co-located collaborative games that use multi-display composition." In: *Designing Interactive Systems Conference 2021*. 2021, pp. 718–733.

[266] João Marcelo Evangelista Belo, Anna Maria Feit, Tiare Feuchtner, and Kaj Grønbæk. "Xrgonomics: Facilitating the creation of ergonomic 3d interfaces." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–11.

[267] Katherine Fennedy, Jeremy Hartmann, Quentin Roy, Simon T Perrault, and Daniel Vogel. "OctoPocus in VR: Using a Dynamic Guide for 3D Mid-Air Gestures in Virtual Reality." In: *IEEE Transactions on Visualization and Computer Graphics* 27.12 (2021), pp. 4425–4438. ISSN: 1077-2626, 1941-

0506, 2160-9306. DOI: 10.1109/TVCG.2021.3101854. URL: https://ieeexplore.ieee.org/document/9507320/.

[268] Geoffrey Gorisse, Gizem Senel, Domna Banakou, Alejandro Beacco, Ramon Oliva, Daniel Freeman, and Mel Slater. "Self-observation of a virtual body-double engaged in social interaction reduces persecutory thoughts." In: *Scientific reports* 11.1 (2021), pp. 1–13.

[269] Kim Halskov and Caroline Lundqvist. "Filtering and informing the design space: Towards design-space thinking." In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 28.1 (2021), pp. 1–28.

[270] Germán Leiva, Jens Emil Grønbæk, Clemens Nylandsted Klokmose, Cuong Nguyen, Rubaiat Habib Kazi, and Paul Asente. "Rapido: Prototyping Interactive AR Experiences through Programming by Demonstration." In: *The 34th Annual ACM Symposium on User Interface Software and Technology*. 2021, pp. 626–637.

[271] Klemen Lilija, Søren Kyllingsbæk, and Kasper Hornbæk. "Correction of Avatar Hand Movements Supports Learning of a Motor Skill." In: *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 2021 IEEE Virtual Reality and 3D User Interfaces (VR). IEEE. Lisboa, Portugal: IEEE, Mar. 2021, pp. 1–8. ISBN: 978-1-66541-838-6. DOI: 10.1109/VR50410.2021.00069. URL: https://ieeexplore.ieee.org/document/9417640/ (visited on 01/16/2023).

[272] Elena Márquez Segura, Katja Rogers, Anna Lisa Martin-Niedecken, Stephan Niedecken, and Laia Turmo Vidal. "Exploring the Design Space of Immersive Social Fitness Games: The ImSoFit Games Model." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–14.

[273] Pedro Monteiro, Guilherme Gonçalves, Hugo Coelho, Miguel Melo, and Maximino Bessa. "Hands-free interaction in immersive virtual reality: A systematic review." In: *IEEE Transactions on Visualization and Computer Graphics* 27.5 (May 2021), pp. 2702–2713. ISSN: 1077-2626, 1941-0506, 2160-9306. DOI: 10.1109/TVCG.2021.3067687. URL: https://ieeexplore.ieee.org/document/9382869/ (visited on 03/29/2022).

[274] Rakesh Patibanda, Xiang Li, Yuzheng Chen, Aryan Saini, Christian N Hill, Elise van den Hoven, and Florian Floyd Mueller. "Actuating Myself: Designing Hand-Games Incorporating Electrical Muscle Stimulation." In: *Extended Abstracts of the 2021 Annual Symposium on Computer-Human Interaction in Play*. 2021, pp. 228–235.

[275] Henning Pohl, Klemen Lilija, Jess McIntosh, and Kasper Hornbæk. "Poros: configurable proxies for distant interactions in VR." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–12.

[276] Aniruddha Prithul, Isayas Berhe Adhanom, and Eelke Folmer. "Teleportation in Virtual Reality; A Mini-Review." In: *Frontiers in Virtual Reality* 2 (Oct. 18, 2021), p. 730792. ISSN: 2673-4192. DOI: `10.3389/frvir.2021.730792`. URL: `https://www.frontiersin.org/articles/10.3389/frvir.2021.730792/full` (visited on 01/16/2023).

[277] Jonas Schjerlund, Kasper Hornbæk, and Joanna Bergström. "Ninja hands: Using many hands to improve target selection in vr." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–14.

[278] Robert Soden, Pradnaya Pathak, and Olivia Doggett. "What we speculate about when we speculate about sustainable HCI." In: *ACM SIGCAS Conference on Computing and Sustainable Societies*. 2021, pp. 188–198.

[279] Barbara Rita Barricelli and Daniela Fogli. "Digital Twins in Human-Computer Interaction: A Systematic Review." In: *International Journal of Human–Computer Interaction* (2022), pp. 1–19.

[280] Jordan Aiko Deja, Sven Mayer, Klen Čopič Pucihar, and Matjaž Kljun. "A Survey of Augmented Piano Prototypes: Has Augmentation Improved Learning Experiences?" In: *Proceedings of the ACM on Human-Computer Interaction* 6.ISS (ISS Nov. 14, 2022), pp. 226–253. ISSN: 2573-0142. DOI: `10.1145/3567719`. URL: `https://dl.acm.org/doi/10.1145/3567719` (visited on 12/09/2022).

[281] Elena Dzardanova, Vlasios Kasapakis, Damianos Gavalas, and Stella Sylaiou. "Virtual Reality as a Communication Medium: A Comparative Study of Forced Compliance in Virtual Reality versus Physical World." In: *Virtual Reality* 26.2 (2022), 737–757. ISSN: 1359-4338. DOI: `10.1007/s10055-021-00564-9`. URL: `https://doi.org/10.1007/s10055-021-00564-9`.

[282] Andreas Rene Fender and Christian Holz. "Causality-preserving Asynchronous Reality." In: *CHI Conference on Human Factors in Computing Systems*. 2022, pp. 1–15.

[283] Jann Philipp Freiwald, Susanne Schmidt, Bernhard E Riecke, and Frank Steinicke. "The continuity of locomotion: Rethinking conventions for locomotion and its visualization in shared virtual reality spaces." In: *ACM Transactions on Graphics (TOG)* 41.6 (2022), pp. 1–14.

[284] Matthias Hoppe, Andrea Baumann, Patrick Chofor Tamunjoh, Tonja-Katrin Machulla, Paweł W Woźniak, Albrecht Schmidt, and Robin Welsch. "There is no first-or third-person view in virtual reality: Understanding the perspective continuum." In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 2022, pp. 1–13.

[285] Jennifer G Kim, Taewan Kim, Sung-In Kim, So-youn Jang, Eun Bin Lee, Heejung Yoo, Kyungsik Han, and Hwajung Hong. "The Workplace Playbook VR: Exploring the Design Space of Virtual Reality to Foster Understanding of and Support for Autistic People." In: *Proceedings of the ACM on Human-Computer Interaction* 6.CSCW2 (2022), pp. 1–24.

[286] Rune Klevjer. *What is the avatar?: Fiction and embodiment in avatar-based singleplayer computer games*. 2022.

[287] Andreas Künz, Sabrina Rosmann, Enrica Loria, and Johanna Pirker. "The Potential of Augmented Reality for Digital Twins: A Literature Review." In: *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE. 2022, pp. 389–398.

[288] Benjamin Lee, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny, and Tim Dwyer. "A design space for data visualisation transformations between 2d and 3d in mixed-reality environments." In: *Proceedings of the 2022 CHI conference on human factors in computing systems*. 2022, pp. 1–14.

[289] Patrick Reipschläger, Frederik Brudy, Raimund Dachselt, Justin Matejka, George Fitzmaurice, and Fraser Anderson. "AvatAR: An Immersive Analysis Environment for Human Motion Data Combining Interactive 3D Avatars and Trajectories." In: *CHI Conference on Human Factors in Computing Systems*. 2022, pp. 1–15.

[290] Yvonne Rogers. *HCI theory: classical, modern, and contemporary*. Springer Nature, 2022.

[291] Jonas Schjerlund, Kasper Hornbæk, and Joanna Bergström. "OVRlap: Perceiving Multiple Locations Simultaneously to Improve Interaction in VR." In: *CHI Conference on Human Factors in Computing Systems*. 2022, pp. 1–13.

[292] Mel Slater, Domna Banakou, Alejandro Beacco, Jaime Gallego, Francisco Macia-Varela, and Ramon Oliva. "A separate reality: An update on place illusion and plausibility in virtual reality." In: *Frontiers in Virtual Reality* 3 (2022), p. 914392.

[293] Kazuma Takada, Midori Kawaguchi, Akira Uehara, Yukiya Nakanishi, Mark Armstrong, Adrien Verhulst, Kouta Minamizawa, and Shunichi Kasahara. "Parallel Ping-Pong: Exploring Parallel Embodiment through Multiple Bodies by a Single User." In: *Augmented Humans 2022*. 2022, pp. 121–130.

[294] Kazuma Takada, Nanako Kumasaki, Tom Froese, Kazuhisa Shibata, and Shunichi Kasahara. "Shadow Clones: an Unattended Visual Interface to Enhance the Task Performance in Multiple Space using Gaze-Switching." In: *Proceedings of the 2022 ACM Symposium on Spatial User Interaction*. 2022, pp. 1–2.

[295] Niels van Berkel and Kasper Hornbæk. "Implications of Human-Computer Interaction Research." In: *Interactions* 30.4 (2023), pp. 50–55.

[296] Laura Crucianelli, Arran T Reader, and H Henrik Ehrsson. "Subcortical contributions to the sense of body ownership." In: *Brain* (2023), awad359.

[297] Isabel Fitton, Christopher Clarke, Jeremy Dalton, Michael J Proulx, and Christof Lutteroth. "Dancing with the Avatars: Minimal Avatar Customisation Enhances Learning in a Psychomotor Task." In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 2023, pp. 1–16.

[298] Jens Emil Sloth Grønbæk, Ken Pfeuffer, Eduardo Velloso, Morten Astrup, Melanie Isabel Sønderkær Pedersen, Martin Kjær, Germán Leiva, and Hans Gellersen. "Partially Blended Realities: Aligning Dissimilar Spaces for Distributed Mixed Reality Meetings." In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 2023, pp. 1–16.

[299] Kaden Gryphon, Veronica Vu, and Haeyong Chung. "A Design Space of Multi-Display Spatial Interactions for Visualization Tasks." In: *Proceedings of the 2023 ACM Symposium on Spatial User Interaction*. 2023, pp. 1–13.

[300] Ryo Hajika, Tamil Selvan Gunasekaran, Chloe Dolma Si Ying Haigh, Yun Suen Pai, Eiji Hayashi, Jaime Lien, Danielle Lottridge, and Mark Billinghurst. "RadarHand: a Wrist-Worn Radar for On-Skin Touch-based Proprioceptive Gestures." In: *ACM Transactions on Computer-Human Interaction* (2023).

[301] Yuan He, Lauren Buck, Brendan Rooney, and Rachel Mcdonnell. "Effect of Avatar Clothing and User Personality on Group Dynamics in Virtual Reality." In: *Proceedings of*

the 16th ACM SIGGRAPH Conference on Motion, Interaction and Games. 2023, pp. 1–8.

[302] Patrick Yung Kang Lee, Ning F Ma, Ig-Jae Kim, and Dongwook Yoon. “Speculating on risks of AI clones to selfhood and relationships: Doppelganger-phobia, identity fragmentation, and living memories.” In: *Proceedings of the ACM on Human-Computer Interaction* 7.CSCW1 (2023), pp. 1–28.

[303] Qiner Lyu, Kentaro Watanabe, Hiroyuki Umemura, and Akihiko Murai. “Design-thinking skill enhancement in virtual reality: A literature study.” In: *Frontiers in Virtual Reality* 4 (2023), p. 1137293.

[304] Karthik Mahadevan, Qian Zhou, George Fitzmaurice, Tovi Grossman, and Fraser Anderson. “Tesseract: Querying Spatial Design Recordings by Manipulating Worlds in Miniature.” In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 2023, pp. 1–16.

[305] Andreea Muresan, Jess McIntosh, and Kasper Hornbæk. “Using Feedforward to Reveal Interaction Possibilities in Virtual Reality.” In: *ACM Trans. Comput.-Hum. Interact.* (2023). ISSN: 1073-0516. DOI: `10.1145/3603623`. URL: `https://doi.org/10.1145/3603623`.

[306] Andreea Muresan, Jess Mcintosh, and Kasper Hornbæk. “Using Feedforward to Reveal Interaction Possibilities in Virtual Reality.” In: 30.6 (2023). ISSN: 1073-0516. DOI: `10.1145/3603623`. URL: `https://doi.org/10.1145/3603623`.

[307] Mel Slater, Carlos Cabriera, Gizem Senel, Domna Banakou, Alejandro Beacco, Ramon Oliva, and Jaime Gallego. “The sentiment of a virtual rock concert.” In: *Virtual Reality* 27.2 (June 2023), pp. 651–675. ISSN: 1434-9957. DOI: `10.1007/s10055-022-00685-9`. URL: `https://doi.org/10.1007/s10055-022-00685-9`.

[308] Jakob Carl Uhl, Helmut Schrom-Feiertag, Georg Regal, Linda Hirsch, Yannick Weiss, and Manfred Tscheligi. “When Realities Interweave: Exploring the Design Space of Immersive Tangible XR.” In: *Proceedings of the Seventeenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI ’23. Warsaw, Poland: Association for Computing Machinery, 2023. ISBN: 9781450399777. DOI: `10.1145/3569009.3571843`. URL: `https://doi.org/10.1145/3569009.3571843`.

[309] Andreea Muresan, Teresa Hirzle, and Kasper Hornbæk. “How and Why To Act Through Multiple Avatars in Virtual Reality.” In: Pending Peer-Review Process. Association for Computing Machinery, 2024.