Universiteit Leiden
The Netherlands

Arguably augmented reality : relationships between the virtual and the real

Schraffenberger, H.K.

Citation
Schraffenberger, H. K. (2018, November 29). *Arguably augmented reality : relationships between the virtual and the real*. Retrieved from https://hdl.handle.net/1887/67292

Version:	Not Applicable (or Unknown)
License:	Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden
Downloaded from:	https://hdl.handle.net/1887/67292

Note: To cite this publication please use the final published version (if applicable).

Cover Page

Universiteit Leiden

Leiden University Repository

The handle http://hdl.handle.net/1887/67292 holds various files of this Leiden University dissertation.

Author: Shraffenberger, H.K.
Title: Arguably augmented reality : relationships between the virtual and the real
Issue Date: 2018-11-29

Arguably Augmented Reality

Relationships Between the Virtual and the Real

Proefschrift

ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van Rector Magnificus prof. mr. C.J.J.M. Stolker,
volgens besluit van het College voor Promoties
te verdedigen op donderdag 29 November 2018
klokke 11:15 uur

door

Hanna Kathrin Schraffenberger
geboren te Stuttgart, Duitsland
in 1983

Promotores:	prof.dr. S. Haring prof.dr.ir. F.J. Verbeek
Promotiecommissie:	prof.dr. A. Plaat (voorzitter) prof.dr. M.E.H. van Reisen (secretaris) prof.dr. C.M. Jonker (Delft University of Technology) prof.dr.ir. P.P. Jonker (Delft University of Technology) dr. M. Postma (Tilburg University)

Copyright © 2018 Hanna Schraffenberger

The layout of this thesis is based on the Tufte-book LaTeXclass (see https://tufte-latex.github.io/tufte-latex/). This class is inspired by the books of Edward R. Tufte. For the cover, the Polyshaper tool at http://polyshaper.co was used to modify a picture of my former work environment at Leiden University.

Visit http://www.arguablyar.com for a digital copy of this thesis.

ISBN: 978-94-92679-67-3

First printing, November 2018

Contents

PART I WHAT IS AUGMENTED REALITY? 1

1 *A First Look* 3

1.1 *The Diversity of AR* 3

1.2 *Introduction to the Thesis* 5

2 *Existing Views* 19

2.1 *Common and Complementary Views* 19

2.2 *Less Common and Diverse Views* 23

2.3 *Conclusion* 30

3 *New Perspectives* 37

3.1 *From Technologies to Experiences* 39

3.2 *From Registration to Relationships* 51

3.3 *From Visuals to Multimodal and Interactive Environments* 57

3.4 *Synthesis, Discussion and Conclusion* 66

PART II WHAT FORMS CAN AUGMENTED REALITY TAKE? 71

4 *Relationships Between the Virtual and the Real* 73

4.1 *Coexistence: Independence of the Virtual and the Real* 78

4.2 *Presence: Spatial Relationships* 80

4.3 *Information: Content-Based Relationships* 84

4.4 *Extended Reality: The Virtual Supplements the Real* 89

4.5 *Hybrid Reality: The Virtual Completes the Real* 98

4.6 *Diminished Reality: The Virtual Removes the Real* 99

4.7 *Altered Reality: The Virtual Transforms the Real* 102

4.8 *Extended Perception: Translation-Based Relationships* 113

4.9 *Physical Relationships: The Virtual and the Real Affect Each Other* 116

4.10 *Behavioral Relationships: The Virtual and the Real Sense and React to Each Other* 121

4.11 *More Relationships* 126

4.12 *Summary, General Discussion and Conclusion* 128

5 *From Imitative to Imaginative Realities: Influences and Interactions Between the Virtual and the Real* 137
5.1 *Theoretical Exploration* 142
5.2 *Practical Exploration* 155
5.3 *Discussion, Conclusions and Future Directions* 166
6 *The Invisible Cube: Introducing Sonically Tangible Objects* 173
6.1 *The Sonically Tangible Cube* 175
6.2 *The Cube in Context* 187
6.3 *Reflection and Outlook* 193
PART III CONCLUSION 201
7 *Conclusion* 203
7.1 *What Is Augmented Reality?* 203
7.2 *What Forms Can AR Take?* 206
7.3 *New Forms of AR* 208
7.4 *Pending Questions* 209
7.5 *Limitations and Concerns* 209
7.6 *Creating AR* 211
7.7 *The Future of AR and AR Research* 211
References 215
INTERMEZZI
Chasing virtual spooks, losing real weight 33
Hitting imaginary walls, pulling virtual strings 133
Subject: Interview 197
ABOUT
Summary 233
Samenvatting 237
Zusammenfassung 241
Acknowledgements 247
Curriculum Vitae 249

List of Figures

1.1 Virtual information appears to exist in and extends the real environment. Reprinted from T. Caudell and D. Mizell (Jan. 1992). "Augmented reality: an application of heads-up display technology to manual manufacturing processes". In: *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*. Vol. 2. IEEE, pp. 659–669. Reprinted under fair use. 4

2.1 A simplified representation of the reality-virtuality continuum as shown in (Milgram, Takemura, et al., 1994). 19

3.1 The virtual beaver *Sphero* (2011) is not just overlaid onto our view but integrated into our view. The picture is a screenshot showing the image displayed on the iPad. (The screenshot was taken by the author.) 40

3.2 Due to inaccurate registration, virtual Pokémon creatures can appear at unrealistic positions or overlaid onto the live view, rather than as part of the real environment. (In order to amplify this effect, the author has manually moved the phone in space. However, Pokémon quite regularly appear 'detached' from the real world without trying to achieve this.) 42

3.3 Three typical fiducial markers that can be recognized by AR software, such as the popular open-source ARToolKit tracking library. (The displayed markers are part of the download of the *ARToolKit SDK* (1999).) 43

3.4 The projection-based AR project by Benko et al. (2014) can make it seem as if virtual objects existed in real space, rather than projected onto the world. The image shows two screenshots from the YouTube video about this project (Microsoft Research, 2014). 45

3.5 A mock-up of the Google Glass concept. Virtual content is overlaid onto the view of the real world but not registered with the real space. This image is a screenshot from a video demonstrating the concept behind Google Glass (Huzaifah Bhutto, 2012). The actual realization of the overlays looks quite a bit different and can be seen in figure 3.8. 46

3.6 A game similar to the *Mozzies* game. Virtual mosquitoes appear to be flying in the space before the phone's lens. The image appeared in a paper by López et al. (2014), and permission to use the image in this thesis was granted by Miguel Bordallo Lopez. 49

3.7 The NS Kids app shows Knoei flying next to the train (left) as well as next to the player showing off his muscles (right) on the camera feed. Screenshots by Jurriaan Rot and Hanna Schraffenberger. 50

3.8 Google Glass can overlay information that relates to our real surroundings and context. This image is a screenshot taken with the device, illustrating the user's view. Image created by and courtesy of Ben Collins-Sussman. 53

3.9 My cat shows no sign of awareness of the virtual beaver Sphero *(2011)*. According to our experience, this can harm the experience of Sphero being a part of the real environment. The picture is a screenshot showing the image displayed on the iPad. (The screenshot was taken by the author.) 56

3.10 The MetaCookie+ headset. Reprinted from Narumi, Nishizaka, et al. (2011a). Reprinted under fair use. 63

4.1 Virtual information completes a physical model of a zebrafish. Without the virtual component, the object is incomplete. Reprinted from M. A. Gómez-Maureira et al. (2014). "Illuminating Shadows: Introducing Shadow Interaction in Spatial Augmented Reality". In: *Creating the Difference: Proceedings of the Chi Sparks 2014 Conference*, pp. 11–18. Reprinted under fair use. 82
4.2 Virtual information removes advertisements on a billboard from the environment and replaces it with alternative content. Reprinted from S. Mann and J. Fung (2002). "EyeTap devices for augmented, deliberately diminished, or otherwise altered visual perception of rigid planar patches of real-world scenes". *Presence: Teleoperators and Virtual Environments*, 11(2), pp. 158–175. Reprinted under fair use. 82
4.3 Virtual information can transform the real world. Here, artist Valbuena (2008) alters the appearance of the *The Hague City hall* with his dynamic installation *N 520437 E 041900 [the hague city hall*. Images from `http://www.pablovalbuena.com/selectedwork/n-520437-e-041900`. Reprinted under fair use. 83
4.4 In this digital mirror, virtual information about the person in front of the mirror is acquired and presented in a comic-like thought-bubble (Gradman, 2010). Photograph by Bryan Jones. Printed with permission. 87
4.5 A physical information board informs visitors about the presence of virtual Pokémon in the environment. Photograph by ANP. Reprinted under fair use. 88
4.6 A combination of a tactile display and a holographic display allows participants to see and feel raindrops hit their palm. Reprinted from T. Hoshi, M. Takahashi, K. Nakatsuma, et al. (2009). "Touchable holography". In: *ACM SIGGRAPH 2009 Emerging Technologies*. ACM, p. 23. Reprinted under fair use. 95
4.7 A comparison between traditional augmented reality (left) and stylized AR (right) as implemented by Fischer et al. (2005). In both images, the teapot is a virtual object, while the cup and the hand are real. However, the stylized version uses an image filter and non-photorealistic rendering. Reprinted from J. Fischer et al. (2005). "Stylized augmented reality for improved immersion". In: *Proceedings IEEE Virtual Reality 2005*. IEEE, pp. 195–202. Reprinted under fair use. 103
4.8 A virtual clock informs customers of the Gusto Esporessobar in Winterswijk about the time. Image by Hanna Schrraffenberger. 111
4.9 The real world affects the appearance of the virtual glass. For instance, we can see the person's hand refracted in the glass. Reprinted from P. Kán and H. Kaufmann (2012). "High-quality reflections, refractions, and caustics in augmented reality and their contribution to visual coherence". In: *International Symposium on Mixed and Augmented Reality (ISMAR 2012)*. IEEE, pp. 99–108. Reprinted under fair use. 117

5.1 Real doors can open for virtual objects. Image © Hanna Schraffenberger and Edwin van der Heide. 138
5.2 Virtual and real domino stones can interact with the use of so-called portals. Reprinted from J. Leitner et al. (2008). "IncreTable, a mixed reality tabletop game experience". In: *Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology*. ACM, pp. 9–16. Reprinted under fair use. 146
5.3 In Sobecka's mirror, the viewer sees an animal overlaid on their own reflection. Image from `http://www.gravitytrap.com/artwork/perfect-creatures`. Printed under fair use. 151
5.4 A virtual ball is bouncing on a real table. Four snapshots from the live-view. Image © Hanna Schraffenberger and Edwin van der Heide. 157

5.5 A virtual ball approaches a row of dominoes (frame 1 and 2), hits the first stone (frame 3) and consequently changes its course (frame 4). Unlike in the real world, the stones do not fall. Image © Hanna Schraffenberger and Edwin van der Heide. 158
5.6 A virtual ball approaches a row of dominoes (frame 1) and hits the first stone (frame 2). This causes all domino stones to fall (frame 3 and 4). Image © Hanna Schraffenberger and Edwin van der Heide. 159
5.7 One of the physical spheres in a Newton's cradle is overlaid with a virtual sphere. The setup results in a realistic interaction between the spheres and allows for participant interaction with all spheres. Image © Hanna Schraffenberger and Edwin van der Heide. 160
5.8 A real door opens for a virtual sphere. Image © Hanna Schraffenberger and Edwin van der Heide. 161
5.9 The colored rubber balls attract virtual spheres of similar colors. Image © Hanna Schraffenberger and Edwin van der Heide. 163
5.10 Virtual colored spheres are attracted by light. The top row shows what happens when the lamp is turned on (from left to right) and the bottom row shows what happens when the light is turned off. Image © Hanna Schraffenberger and Edwin van der Heide. 164
5.11 Virtual colored spheres are attracted by the hands of the author/participant. Quickly moving the hands away from the virtual balls allows one to let go off them. Image © Hanna Schraffenberger and Edwin van der Heide. 166

6.1 A colleague explores the virtual, invisible and non-tactile cube. A Leap Motion Controller is used to track the position of his fingertips. 176
6.2 We recorded the sounds of squashing a tiny plastic bag filled with small crumbles of the aluminum foil (the image shows a recreation). 180
6.3 The sonically tangible cube was divided into 64 sub-cubes. A binaural recording was made at all 64 positions. Image of the cube contributed by Wim van Eck. 181
6.4 The author interacts with the cube. 184
6.5 The author shakes the hand to a variable extent to produce different types of rustling sounds. 185
6.6 During the development of the sound-design, a monitor was placed on the desk. As soon as a visual representation of the cube and the fingertips was presented, the cube no longer felt like a part of the real environment. Instead, it felt as if one were reaching into the space depicted on the monitor and as if one were touching the displayed cube in this space. 186

Part I

What Is Augmented Reality?

1 A First Look

With the advent of augmented reality (AR), virtual content has entered our everyday environment in a profoundly new way. Virtual objects no longer simply appear on the screen of computers, tablets, mobile phones, smart watches, digital information boards, advertisement screens or other displays. Rather, they seem to exist right here, in our physical space, just like real objects do: Wearing a head-mounted display (HMD), virtual, three-dimensional game characters appear to walk on real streets (e.g., Thomas et al., 2000). Looking at the environment through a mobile phone's screen, site-specific information, such as where to find nearby restaurants, metro stops and ATMs appears to be floating through the space in front of us (e.g., *Layar* 2009). Using AR technology, such as the HoloLens headset (Microsoft, n.d.), we can invite virtual characters into our house or turn our living room into the venue of a partially real and partially virtual adventure. With AR, the presence of virtual content in real space has gained a new dimension.

1.1 The Diversity of AR

Augmented reality research and development is usually traced back to 1968, when Sutherland (1968) introduced a head-mounted display (HMD) that allowed users to see both computer-generated images and the real surroundings at the same time.

When Caudell and Mizell (1992) coined the term *augmented reality* in the early 1990s[1], they built on this principle and proposed a heads-up display meant to "augment the worker's visual field of view" (p. 660) by overlaying virtual content onto the worker's view of the real world. Their proposed headset was intended to make the life of assembly and manufacturing workers easier—for instance, by presenting virtual arrows in real space, indicating where to drill (see figure 1.1).

[1] While some sources attribute the term to Caudell and refer to the year 1990 (Chien et al., 2010; K. Lee, 2012) others attribute it to Caudell and his colleague Mizell and refer to the year 1992 (Olsson and Salo, 2011; van Krevelen and Poelman, 2010).

In the meantime, many more researchers and developers have followed the examples of Sutherland (1968) and Caudell and Mizell (1992) and proposed AR systems that integrate computer-generated images into our view of the world by means of a visual display. For instance, the KARMA (Knowledge-based Augmented Reality for Maintenance Assistance) project uses a head-mounted display to superimpose virtual instructions on how to refill the paper tray right

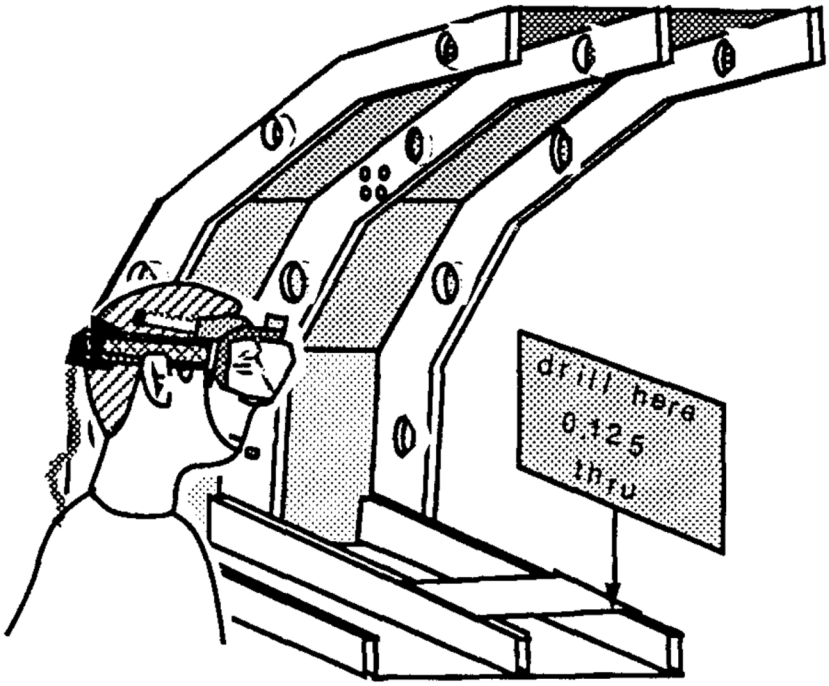

Figure 1.1: Virtual information appears to exist in and extends the real environment. Reprinted from T. Caudell and D. Mizell (Jan. 1992). "Augmented reality: an application of heads-up display technology to manual manufacturing processes". In: *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*. Vol. 2. IEEE, pp. 659–669. Reprinted under fair use.

onto the office printer (Feiner, Macintyre, et al., 1993). The AR version of the popular computer game Quake displays game characters in the real environment (Piekarski and Thomas, 2002; Thomas et al., 2000) by means of an HMD. Looking through an iPhone or iPad, the app Sphero turns a robot ball into a visual virtual beaver (*Sphero* 2011). In the MagicBook project, virtual 3D scenes come alive on the pages of a physical book when the book is viewed through special hand-held glasses (Billinghurst, Kato, and Poupyrev, 2001).

However, not everyone has taken the concept of "augment[ing] the worker's visual field of view with useful and dynamically changing information" (Caudell and Mizell, 1992, p. 660) so literally. Various so-called augmented reality applications approach AR more broadly. To mention just a few examples: The Disney Research team has used weak electric signals that are injected on the user's body to create a tactile-based form of AR that allows participants to feel virtual textures when running their fingers over real physical objects (Bau and Poupyrev, 2012). Visiting a museum, augmented reality audio guides can inform us about the art pieces we encounter by playing back matching pre-recorded sound-files when a visitor is close to certain artworks (Bederson, 1995). Looking at the environment through our phone's screen, we can see overlays of historic photographs (as opposed to 3D computer-generated objects), showing us how the area used to look like a long time ago (*Museum of London: Streetmuseum* 2014). Furthermore, in so-called *spatial augmented reality*, virtual con-

tent is integrated into the real world directly (rather than superimposed onto a participant's view) and, for instance, seemingly transforms the physical architecture of buildings by means of projections (e.g., Valbuena, 2008).

In addition, various AR projects use multi-sensory rather than solely visual displays. For instance, the MetaCookie headset (Narumi, Kajinami, et al., 2010a; Narumi, Nishizaka, et al., 2011b) not only changes the visual appearance of a neutral cookie into the appearance of a specific cookie (e.g. a chocolate flavored, almond or cheese cookie) but also presents the user with the matching olfactory information. (Reportedly, this can alter the taste of the cookie). Also, researchers have explored using force feedback devices such as the Phantom in order to give a tangible dimension to virtual visual objects (e.g., Bianchi et al., 2006).

By now, a wide variety of so-called AR applications exists. In many ways, these projects could not be more diverse. They make use of a broad variety of different technologies, such as headsets, projectors, headphones and tactile displays. In line with this, they present different kinds of sensory stimuli, like visuals, sounds and scents and provide various types of virtual content, among which textual information, photographs and sound recordings. They alter our experience of the real world in various ways; for instance, by seemingly removing physical objects from our view or by integrating additional elements into it. They are designed for many different contexts, such as work, entertainment and education. Accordingly, they serve a variety of purposes. For instance, some AR apps are here to inform us about our surroundings while others exist to keep us entertained.

On first sight, it is rather unclear what the various projects that go under the term *augmented reality* have in common. On the one hand, there seems to be a distinct group of projects that use technologies that overlay virtual content onto a participant's view, making it appear as if this content existed in real space rather than on a screen. On the other hand, there is a diverse group of projects that deviate from this principle, suggesting that there is more to AR. In their totality, the different forms of AR provide a rather blurry picture of AR that raises the questions: *"What is augmented reality?"* and *"What forms can AR take?"*. In this thesis, we address these questions and explore the fundamental characteristics, underlying principles and potential manifestations of AR.

1.2 *Introduction to the Thesis*

This thesis is about augmented reality (AR). AR is commonly considered a technology that integrates virtual images into a user's view of the real world. Yet, this thesis is not about such technologies themselves. Why is that?

We believe that a technology-based notion of AR is incomplete. In this thesis, we challenge the technology-oriented view, provide new perspectives on AR and argue for a different understanding. To be precise, we depart from widespread definitions of AR in three complementary ways.

First, we do not view AR as a technology. Instead, we claim that AR technology enables augmented reality. In this work, we focus on the resulting augmented reality environments, scenarios and experiences rather than on the technologies that enable them.

Secondly, we treat AR as a modalities-encompassing (not only visual) phenomenon and argue that AR engages all our senses. Instead of focusing on what a user or participant *sees*, we focus on non-visual and multimodal aspects of AR.

Third, we view AR as a result of the relationships between the virtual and the real. Whereas AR is generally assumed to involve the spatial integration of virtual content in (a participant's view of) the real world, we believe that other types of relationships between the virtual and the real are possible, potentially leading to other and new forms of AR.

In this thesis, we combine these different points of departure. We approach AR from a fundamental, conceptual, technology-independent, experience-focused, human-centered, modalities-encompassing view and explore the various relationships between the virtual and the real.

By approaching AR from this point of view, we create an unusually broad and diverse image of what AR is, or arguably could be. We learn about the fundamental characteristics of AR and the many possible manifestations it can take, including many forms that do not involve a technology that integrates virtual content in (our view of) the world.

1.2.1 *Aim*

The aim of this thesis is twofold, *theoretical* and *practical*. On the theoretical side, we strive for a better understanding of what augmented reality is and encompasses. On the practical side, we aim at facilitating, creating and exploring new forms of AR. In particular, we are interested in novel forms of AR that do not imitate reality but provide truly new experiences and interactions that have no equivalent in a purely physical world.

Both the practical and the theoretical aspects of this thesis serve a *fundamental* purpose: They address the underlying question of what AR is and what (else) it can be. More specifically, both address the various relationships between the virtual and the real that shape AR experiences.

This thesis fills a gap in existing AR research, which often either aims at (1) creating or improving AR systems *technologically* or (2) re-

alizing and exploring specific AR *applications*. For instance, there is plenty of research into technologies and techniques that enable or support the integration of virtual objects in our view of the real physical world, such as tracking or calibration techniques (cf. Zhou, Duh, and Billinghurst, 2008). In contrast, our research aims at advancing AR from a fundamental, experience-focused and conceptual, rather than applied or technological perspective. In line with this, technological aspects, such as tracking, fall out of the scope of this thesis. We are not so much interested in how things *are* or what a system does, but primarily interested in the perceptual result—in how things *appear* and what a participant can do, feel, see, hear, smell, taste or touch. Likewise, we are interested in the possible *manifestations* of AR, rather than *applications* of AR.

In existing technological and applied AR research, it is custom to talk about a *user* of AR. As we approach AR in a much broader and freer context, we speak of a *participant* who experiences and acts in AR, rather than a *user* who operates an AR system.

This thesis argues for a new and broader understanding of AR. However, our goal is to provide an *additional* and *complementary* perspective, rather than an *alternative* perspective.

While it might seem as if we aim to *define* the *term* AR, the primary focus of this thesis is not concerned with *terminology* and how the term is or should be used, but on actual AR experiences and scenarios. The question is not so much whether something should be called AR or not, but rather, what forms AR can take.

1.2.2 *Motivation*

Why are we so interested in learning about what AR is and can be? Our main motivation to answer these questions is personal curiosity. In our opinion, fundamental research does not need to be motivated or justified by any reason other than a researcher's desire to know or learn something about the world. However, this does not mean that we do not see any potential benefits of answering these questions. Most importantly, we believe that a better theoretical understanding of AR will inform AR research and practice and lead to novel manifestations of AR. In this regard, our work is motivated by the belief that current AR research and development is adopting an unnecessarily narrow view, and thus might be missing out on exciting opportunities. We hope to free practitioners and researchers alike from restricting ideas, such as the association of AR with *visual* overlays, and thereby inspire and facilitate new and different forms of both AR and AR research. Furthermore, we believe that in order to work and communicate in such a complex field as AR, we have to be able to clearly identify and single out specific forms of AR. We are convinced that a thorough theoretical understanding and an accurate definition will be beneficial

for related scientific disciplines that work with AR, such as medical and educational research.

In addition to the scientific and practical relevance of answering these questions, we are motivated by the social relevance of studying AR. Augmented reality plays an ever-increasing role in our everyday world, and we believe it is important to understand a phenomenon that has the potential to affect (or, as we will argue, *already* affects) our everyday lives.

1.2.3 *Methodology*

In order to learn more about the fundamental characteristics and possible manifestations of AR, we follow a multidisciplinary, topic-oriented, human-centered, partially practical, partially theoretical, philosophical, argumentative and most of all exploratory approach.

Multidisciplinary and Topic-oriented

This thesis is multidisciplinary in the sense that it draws from and contributes to many domains. The thesis follows the approach of "topic-oriented scholarschip" as defined by van Duijn (2016, p. 19):

> [...] it takes a topic as its starting point and then seeks for the right combination of methods and expertise across multiple disciplines for approaching it, instead of starting from the set of questions and assumptions customary in a particular discipline. Thereby, it aims at making progress not just by contesting existing findings, but also by adding new perspectives on these findings.

Accordingly, our research takes the topic "augmented reality" as a point of departure, and consequently incorporates knowledge and methods from different disciplines, such as engineering, philosophy, perception research, human-computer interaction and media studies in order to gain a better and multifaceted understanding of what AR is and potentially can be. Furthermore, we provide new perspectives on the topic. In particular, we approach augmented reality from a technology-independent, experience-focused, human-centered and modalities-encompassing perspective.

We have chosen for such a broad and multidisciplinary approach because augmented reality is a highly diverse and multidisciplinary research field. For instance, the primary AR conference ISMAR (International Symposium on Mixed and Augmented Reality) regularly featured both a "*Science and Technology*" track as well as a *"Media, Art, Social Science, Humanities and Design"* track).[2] In line with this, existing AR research combines, draws from and contributes to various technological research areas, such as engineering, computer vision, display development, human-computer interaction, wearable, ubiquitous and mobile computing, software engineering and information visualization. At the same time, it is a topic of interest in areas such as media

[2] These two tracks were offered as the two main tracks from 2009 until 2015. The "Media, Art, Social Science, Humanities and Design" first appeared under the name *Arts, Media, and Humanities*. In 2016, it was no longer offered as a second track, but instead took the form of a workshop.

art, design, psychology, communication studies, visual studies, media studies and philosophy.

When it comes to *applications* of AR, an even wider variety of disciplines is involved. Among others, AR has applications in areas such as medicine, manufacturing and education (for an overview, see, e.g., Azuma, 1997). Accordingly, many actual AR projects that we can study to learn about AR have been realized in the context of other research fields or in collaboration with other disciplines. In this thesis, we do not limit ourselves to research, projects or methods from one specific discipline or research direction, because we want to get an overview of the various possible manifestations of AR. This means that examples from a diverse range of domains are considered and included based on whether they reveal insights about what AR is and can be. However, as the focus of this thesis is on the possible AR *manifestations* rather than *applications*, we do not aim to give a comprehensive overview of AR applications. To some extent, we focus on art and entertainment examples (e.g., games). We do this for two reasons. First of all, because this thesis is realized in an art context. More specifically, this research is partially conducted at the Augmented Reality Lab (AR Lab) at the Royal Academy of Art, The Hague, where the author has a guest research position. This lab focuses, among others, on exploratory research in the artistic domain.[3] Secondly, we explore art and entertainment applications because we expect these to focus more explicitly on AR *experiences*. This, however, does not mean that other domains or disciplines were deliberately excluded. At times, the same examples are used repeatedly, to illustrate different points about AR.

[3] Unfortunately, the AR Lab has been closed in 2014.

Theoretical and Practical Research

As mentioned, our study has a theoretical and a practical aim. In line with this, we approach AR both from a practical as well as from a theoretical perspective.

Our theoretical approach includes a review of existing research literature as well as the use of arguments and ideas in order to arrive at a new and better understanding of what AR is and potentially can be. Regarding existing literature, we focus on influential views and descriptions of AR projects that have shaped current understandings of AR. Furthermore, we also pay attention to less common or commonly overlooked literature and AR projects.

In addition to this theoretical study, we also follow a practical approach. This involves actively engaging with existing AR projects, such as the mobile app "Pokémon GO" (*Pokémon GO* 2016) and, more importantly, building our own AR projects. This approach is used to arrive at new ideas and concepts for novel forms of AR. Furthermore, our practical approach is motivated by the belief that "by doing and creating, new scientific insights into the underlying question are encountered." (Media Technology MSc Programme - Leiden University,

n.d.).[4] We believe that creating our own AR scenarios (potentially) results in additional realizations about what AR is and what forms it can take. With respect to the practical exploration, we draw from our own *first-person experience*. Due to the constraints in time, experiments with participants are out of the scope of this thesis.

[4] This idea is at the basis of the Media Technology program at Leiden University, where this research was carried out.

With respect to the practical aspects, we build upon our experience in the field of Human-Computer Interaction research. Practical projects in this research are—as far as possible—realized with cheap everyday technology rather than typical or special AR equipment (e.g., we use our normal office computer, a webcam and a monitor to test ideas rather than a head-mounted display). Furthermore, projects are realized in a prototypical manner.

A Philosophical Approach

We believe that AR is more than just a technology that integrates virtual imagery into our view. Yet, our point of departure is not so much a hypothesis we can test objectively, rather than it is an attitude towards AR and an open question: What does AR entail if we broaden existing definitions and approach AR from a human-centered, technology-independent, modalities-encompassing and relationship-focused perspective? In other words, we are looking for a better understanding of what AR is and potentially can be, and are not concerned with testing an overreaching hypothesis.[5] Because we are interested in the qualities, fundamental characteristics and potential manifestations of AR, we have chosen an exploratory research approach.

[5] However, this does not mean that we have no assumptions or hypotheses at all. For instance, we address the assumptions that virtual objects do not have to behave like real objects in order to appear as a believable part of real space (see chapter 5).

Although our research is interested in the *qualities* of AR, our research approach does not incorporate common qualitative research methods such as focus groups, interviews and participant observation. Instead, it approaches the topic of AR in a rather playful manner. In terms of existing methods, our approach could best be described as *philosophical*. This is because our investigation into AR is driven by reasoning, and uses the instruments of what Sheffield (2004) calls "The Philosopher's Toolbox": we *analyze*, *clarify* and *criticize*. More specifically, we analyze the field of augmented reality with the goal of identifying defining characteristics, criticize existing notions of AR and clarify what else AR is and potentially could be. Moreover, our research shows similarities to dialectic research, which also often aims at developing new understandings rather than at testing hypotheses (Dialectical research, n.d.). Also, like dialectic investigation, we work "with arguments and ideas, rather than data" and examine competing notions and perspectives (Dialectical research, n.d.).

While our methods could be considered philosophical, we would like to emphasize that we do not view this work as philosophy. Likewise, the author does not see herself as a philosopher. Although the term 'augmented reality' might invite this, a philosophical discussion of the nature of reality is out of the scope of this thesis.

Aside from similarities with philosophical research, our research shares qualities with human-computer interaction (HCI) research, which also often incorporates a *human-centered* approach and focuses on the human experience. However, in contrast to much HCI research, we do not tie a human-centered approach to usability. Also, whereas *experience* is often addressed in the context of *User Experience (UX)* when it comes to new technologies, we deliberately do not focus on UX. Instead, we focus on the unique characteristics of AR experiences.[6] Our human-centered approach entails that we ask how things appear to the participant, that we question what a participant perceives and what a participant can do. We do focus on these aspects because we believe that AR is created *for humans*, with the goal of creating certain perceptual results and enabling certain experiences, rather than for technological purposes. It hence seems natural and necessary to look at AR experiences in order to understand the essence of AR.

[6] We would like to direct readers with an interest in UX to the seminal paper "User experience - a research agenda" by Hassenzahl and Tractinsky (2006).

A Cartographic Process

The nature of this research can be best summarized as exploratory. In a metaphorical way, it can be compared to a cartographic process. It explores the "AR landscape" in the hope of discovering "new places", but also with the goal to learn more about "known spaces" by looking at them from new perspectives. Furthermore, it re-evaluates where the lines between AR and other disciplines ought to be drawn and proposes a broader, more encompassing understanding of AR.

Limitations

This thesis does not focus on AR systems and technologies, but on the various forms AR can take, the different relationships between the virtual and the real that shape AR, and the many experiences that AR systems enable.

One limitation of this research is that our observations are based on our own, subjective experiences of AR. Of course, our experience might not fully represent how participants in general perceive AR and we cannot rule out the possibility that our experience is influenced by our expectations and beliefs about AR.

While experiments with participants would be desirable, these fall out of the scope of the thesis. This is because in order to systematically conduct experiments with participants, it is crucial to first understand what characterizes AR, and what types of experiments would foster a better understanding. In this regard, our exploratory study can be seen as a first fundamental step towards facilitating more directed experiments with participants in the future.

Another limitation of this research is that we draw from existing AR literature and other media, such as articles and videos rather than from

a first-hand experience of the documented projects. On the one hand, studying such mediated accounts of AR will surely allow us to learn about AR. On the other hand, it is not always possible to make inferences about AR *environments* and *experiences* from studying textual or visual descriptions that focus on other aspects, such as the workings of the *system*. In fact, our own research argues that descriptions of AR systems do not suffice for describing the resulting AR environments and experiences (what happens on a technological level is rather different from its perceptual result and similar systems can create many different environments and experiences). We hence have to be careful not to draw unsound conclusions about AR environments and experiences from such system-focused accounts of AR.

1.2.4 *Results*

There are two main contributions of this PhD research: On the theoretical side, it provides a better understanding of what augmented reality is and potentially can be. On the practical side, it suggests novel forms of AR.

1.2.5 *Structure and Outline*

The thesis is organized into three parts that contain seven chapters of varying length. Part 1 address the question *"What is augmented reality?"* and comes to the conclusion that relationships between the virtual and the real are decisive for AR. Part 2 investigates what forms AR can take and explores the relationships between the virtual and the real. Part 3 concludes the thesis, summarizes our results and presents suggestions for future AR research.

Part 1: What Is Augmented Reality?

Part 1 serves an introduction to the topic of augmented reality and addresses the question *"What is augmented reality?"*. We have a look at so-called AR applications, at definitions and descriptions and present our own perspective on AR.

In this chapter ("A first look"), we have taken an initial glance at examples of AR and illustrated the diversity of the AR landscape. On the one hand, we have encountered various AR works that use some sort of visual display to present virtual content and make it look as if this content existed in the otherwise real surroundings. On the other hand, we have seen examples that deviate from this typical setup, use different technologies (e.g., projectors or audio players), present us with different content (e.g., sound or tactile sensations) and create different experiences (e.g., alter how a real object feels or how a real cookie tastes). Together, the different examples of AR leave us with a rather blurry picture of the AR landscape and raise the question of

what augmented reality is.

In chapter 2 ("Existing views") we investigate how existing research answers this question. We review existing definitions and descriptions of AR and identify three common and intertwined ideas about augmented reality: First of all, AR is generally considered a technology. Second, AR is widely understood in terms of visual virtual overlays that are presented on top of a participant's view of the real world. Third, AR is considered to spatially align virtual content with the real world in three dimensions. These ideas are not at odds but complement each other well. Together, they draw a clear image of AR as a technology that integrates virtual content into our view of the real world.

At the same time, our review of existing AR literature also reveals many divergent and *broader* understandings of AR. For instance, we encounter research that also considers non-visual virtual content (such as sound) in the context of AR and researchers that explicitly argue against seeing AR as a technology. In addition, we notice that there are a variety of different claims about the qualities of the virtual content, the role of the real world in AR, the role of the user or participant and the question of what is augmented in AR.

In their totality, the partially agreeing and partially contradicting views on AR leave little doubt that AR *can* involve technologies that overlay virtual objects onto a participant's view and aligns them with the real world in 3D. At the same time, we get a strong sense that there is more to AR than such technologies. As such, the review leaves us wondering, what, if not just a technology, AR is or can be.

In chapter 3 ("New Perspectives), we respond to our initial findings, challenge commonly accepted views, and argue for new (or at least different) perspectives on AR. First, we depart from the understanding of AR as a technology. Instead, we claim that AR technology enables augmented reality. We focus on the resulting augmented reality environments and experiences rather than on the technologies that enable them. Second, we treat AR as a multimodal and interactive environment and argue that AR engages all our senses. Instead of focusing on what a user or participant sees, we focus on non-visual, multimodal and interactive aspects of both the real world and virtual content. Third, we see AR as a result of the relationships between the virtual and the real. Whereas AR is generally assumed to involve the spatial alignment of virtual content with the real world in 3D, we believe that other types of relationships between the virtual and the real are possible, potentially leading to other and new forms of AR. These three ideas are synthesized and culminate in our definition of AR as an interactive and multimodal environment where a participant experiences a relationship between virtual content and the real surroundings.

Part 2: What Forms Can AR Take?

Part 1 has identified relationships between the virtual and the real as crucial for AR. Part 2 discusses such relationships and explores what forms AR can take. In chapter 4 ("Relationships between the virtual and the real"), we illustrate the different ways in which the virtual and the real can relate to one another. On a fundamental level, we distinguish between *(1) coexistence* (participants do not experience any link between the virtual and the real), *(2) spatial relationships* (virtual content seemingly exists in real space) and *(3) content-based relationships* (the virtual relates to the real content-wise).

Subsequently, we question how virtual content can affect its real surroundings. Based on the role that the virtual content plays in the real space, we distinguish between five forms of AR:

1. *Extended reality:* scenarios where the virtual supplements the real environment.
2. *Diminished reality:* cases where virtual content seemingly removes real elements from the real environment.
3. *Altered reality:* environments where the virtual information changes the apparent qualities of the real world.
4. *Hybrid reality:* scenarios where the virtual completes a physical environment that would be considered incomplete without the virtual additions.
5. *Extended perception:* cases where unperceivable but real aspects of the real world are translated into virtual information that we can perceive with our senses.

We then focus on scenarios where virtual objects seemingly exist in and extend the real world. We notice that the presence of virtual objects in real space opens up possibilities for influences and interaction between the virtual and the real. On this level, we distinguish among two main forms of relationships between the virtual and the real: *(1) physical relationships*, where the virtual and the real seemingly affect each other physically and *(2) behavioral relationships*, where the virtual and the real sense each other and react to one another on a social or behavioral level.

Subsequently, we briefly discuss other possible relationships, such as temporal relationships between the virtual and the real and musical relationships between virtual and real instruments. We conclude the fourth chapter with a summary, general discussion and reflection.

Chapter 5 ("From Imitative to Imaginative Realities: Influences and Interactions Between the Virtual and the Real") is dedicated to the interaction between the virtual and the real. Based on the fact that virtual objects do not have to adhere to physical laws and cannot directly apply forces to real objects, we ask the following questions: What types

of interaction between the virtual and the real are both possible and believable? We explore (1) whether virtual objects can interact with physical objects in a realistic manner as well as (2) whether they can interact in imaginative but believable ways. In order to answer these questions, we follow both a theoretical and a practical approach. We review existing research and AR works, conduct our own initial series of practical experiments as well as reflect upon these experiments. We present a general discussion and conclude that virtual and real objects can believably simulate real-world influences as well as influence each other in imaginative ways that have no equivalent in the physical world.

Chapter 6 ("Sonically tangible objects") builds on the idea that virtual objects can differ from real objects and hence, also could be perceived differently from how we perceive real objects. In order to explore and illustrates such possibilities, we develop and present a prototype of what we call sonically tangible objects. More concretely, we present a virtual, invisible and non-tactile cube that is placed in a real, physical space. This cube can be experienced through exploratory hand gestures and sonic feedback. Touching the cube with one's fingers triggers binaural sounds that appear to originate from the exact spot where the object is touched. Our initial experimentation suggests that this sound- and movement-based approach can result in tactile-like experiences and convey the presence of virtual objects in real space. We discuss the concept behind, implementation of and our experience with the sonically tangible cube and place our research in a broader context.

Part 3: Conclusion

Part 3 concludes the thesis. It contains the final chapter of the thesis ("Conclusion"). In this chapter (7), we revisit our main questions (*"What is augmented reality?"* and *"What forms can AR take?"*) and reflect on the answers we have arrived at. Furthermore, we address pending questions that have surfaced during this trajectory (e.g., *"What is augmented in AR?"*) and that we can answer now that we have a thorough understanding of existing research, hands-on experience and our own comprehensive theory of AR. In addition, we summarize insights that can guide the design of AR experiences. (E.g., we suggest to incorporate both multimodal virtual content as well as multimodal qualities of the real world when working with AR and emphasize that designers can not only give shape to virtual content but also actively design the relationship between the virtual and the real.) Moreover, we discuss methodological and technological limitations of our study, and present possible directions for future AR research and development. For instance, we suggest researching the concept of believability (*"When is the behavior and appearance of virtual objects in real space believable?"*) and to systematically explore which factors contribute to virtual objects

being experienced as present in real space. Finally, we propose to focus less on mimicking our existing reality, and instead, to create new, imaginative and curious forms of AR that have no counterpart in a purely physical world.

1.2.6 *Intermezzi*

This thesis contains three intermezzi. These intermezzi are short independent articles about AR. Unlike the rest of this thesis, they are written in an entertaining, informal and personal way that is atypical for scientific publications and they provide a yet different perspective on AR. The included intermezzi have appeared in a slightly different form in the AR[t] magazine, a semi-annual magazine series about augmented reality, art and technology that has been edited by the author during her time has a PhD student.

Intermezzo 1 is a short essay that discusses the idea of audio-augmented reality in the context of going for a run with a mobile training application. Intermezzo 2 discusses the similarities between AR and urban dance and explores the idea of creating the impression of virtual objects existing in real space through movement. Intermezzo 3 is an open letter to media theorist Lev Manovich. It presents and discusses questions that have come up during reading Manovich's 2006 article "The Poetics of Augmented Space" and his 2001 book *The Language of New Media*.

The intermezzi are included in between thesis chapters. They can easily be recognized as they are printed on yellow paper and use a different page layout.

1.2.7 *Publications and Collaboration*

This thesis takes the form of a monograph rather than the increasingly popular form of an article thesis (also referred to "thesis by publication") that bundles independent research papers. Yet, the thesis is based on and includes material from the following published articles:

- H. Schraffenberger and E. van der Heide (2013a). "From Coexistence to Interaction: Influences Between the Virtual and the Real in Augmented Reality". In: *Proceedings of the 19th International Symposium on Electronic Art (ISEA2013)*. Ed. by K. Cleland et al. Sydney, pp. 1–3.

- H. Schraffenberger and E. van der Heide (2013b). "Towards Novel Relationships between the Virtual and the Real in Augmented Reality". English. In: *Arts and Technology*. Ed. by G. De Michelis et al. LNICST 116. Springer, pp. 73–80.

- H. Schraffenberger and E. van der Heide (2014). "The Real in Augmented Reality". In: *Proceedings of the Second Conference on Com-*

putation, Communication, Aesthetics and X (xCoAx 2014). Ed. by M. Carvalhais and M. Verdicchio, pp. 64–74.

- H. Schraffenberger and E. van der Heide (2014b). "Everything Augmented: On the Real in Augmented Reality". *Journal of Science and Technology of the Arts*, 6(1), pp. 17–29.

- H. Schraffenberger and E. van der Heide (2015). "Sonically Tangible Objects". In: *Proceedings of the Third Conference on Computation, Communication, Aesthetics and X (xCoAx 2015)*. Ed. by A. Clifford et al., pp. 233–248.

- H. Schraffenberger and E. van der Heide (2016). "Multimodal Augmented Reality: The Norm Rather Than the Exception". In: *Proceedings of the 2016 Workshop on Multimodal Virtual and Augmented Reality (MVAR '16)*. ACM, pp. 1–6.

- H. Schraffenberger and E. van der Heide (2018). "Reconsidering Registration: New Perspectives on Augmented Reality". In: *Interactivity, Game Creation, Design, Learning, and Innovation. ArtsIT 2017, DLI 2017*. Ed. by A. L. Brooks et al. LNICST 229. Springer, pp. 172–183.

The following published articles from the AR[t] magazine have been included as intermezzi in a slightly different form in this thesis:

- H. Schraffenberger (Nov. 2012). "Chasing virtual spooks, losing real weight". *AR[t], Augmented Reality, Art and Technology*, 2. Ed. by Y. Kolstee et al., pp. 48–51. URL: `http://arlab.kabk.nl/ar-magazines`. (Intermezzo 1)

- H. Schraffenberger (May 2014). "Hitting imaginary walls, pulling virtual strings". *AR[t], Augmented Reality, Art and Technology*, 5. Ed. by H. Schraffenberger et al., pp. 66–71. URL: `http://arlab.kabk.nl/ar-magazines`. (Intermezzo 2)

- H. Schraffenberger (May 2013). "Subject: Interview". *AR[t], Augmented Reality, Art and Technology*, 3. Ed. by H. Schraffenberger et al., pp. 18–23. URL: `http://arlab.kabk.nl/ar-magazines`. (Intermezzo 3)

As one can see, the scientific articles listed above all have been realized in collaboration with my colleague Edwin van der Heide. During this PhD trajectory, Edwin van der Heide has acted as an unofficial daily supervisor, and this work is strongly shaped by our regular discussions. The scientific articles have largely been restructured, rewritten, adapted and extended to accommodate the book format and to incorporate numerous new and additional insights. An exception is chapter 6, which is largely based on our paper "Sonically Tangible Objects" (Schraffenberger and van der Heide, 2015). Also, the second

half of chapter 5 is strongly based on "From Coexistence to Interaction: Influences Between the Virtual and the Real in Augmented Reality" (Schraffenberger and van der Heide, 2013a).

Although we present this thesis as a book, it is important to us that readers can focus on single chapters that raise their particular interest. In order to make sure the individual chapters are readable independently, some arguments and examples are repeated throughout the thesis. The downside of this approach is that the thesis contains some redundant parts. However, we believe the fact that every chapter can also stand on its own outweighs this disadvantage.

During my time as a PhD student, I was lucky to spend several years as a guest researcher at the AR Lab, which was based at the Royal Academy of Art in The Hague. This collaboration has resulted in the above-mentioned AR[t] magazine—a semi-annual magazine about augmented reality, art and technology that was aimed at the general public. We would like to direct the interested reader to this publication series, which also contains several more articles by the author. The AR[t] magazine publications can be found at `http://arlab.kabk.nl/ar-magazines`.

2 Existing Views

What is augmented reality? This is one of the key questions we address in this thesis. If we turn to existing answers, we can find many varying, often complementary, sometimes contradicting views on the subject. Yet, there are some notions of AR that have gained wide acceptance.[1]

2.1 Common and Complementary Views

First, there is the widespread understanding of AR in terms of Milgram and Kishino's (1994) much-cited reality-virtuality continuum (see figure 2.1) (see also Milgram, Takemura, et al., 1994).[2] The presented continuum ranges from purely real environments to entirely virtual environments.[3] The space in between these extremes is referred to as "Mixed Reality". The field of mixed reality includes both augmented reality and augmented virtuality. Augmented reality is placed somewhat closer to the real environment, and describes an (display of an)[4] otherwise real environment that is augmented by virtual objects. Similarly, Milgram, Takemura, et al. (1994) describe augmented virtuality as a principally virtual environment that is augmented through the addition of "real (i.e. unmodelled) imaging data" (p. 285). This happens, e.g., when a user's real hand is displayed in an otherwise virtual environment.

[1] AR has been actively addressed from a computer engineering perspective. Many definitions and descriptions that we review have been presented in an engineering context and have no ambition to make fundamental claims about the nature of AR. We nonetheless review such descriptions because collectively, they provide an overview of how AR is commonly approached.

[2] Originally, this continuum was referred to as "virtuality continuum" (Milgram and Kishino, 1994). However, by now the continuum is commonly known and referred to as the "reality-virtuality continuum".

[3] With their continuum, the authors focus on environments that are viewed via some sort of visual display.

[4] Milgram and Kishino (1994) refer to AR both as "all cases in which the display of an otherwise real environment is augmented by means of virtual (computer graphic) objects" (p. 1321) as well as "any case in which an otherwise real environment is 'augmented' by means of virtual (computer graphic) objects" (p. 1322).

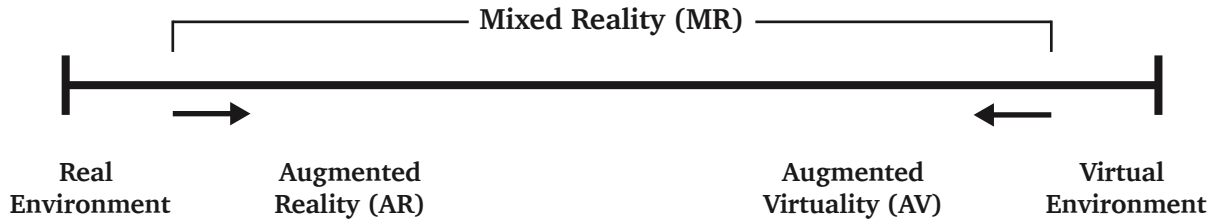

Figure 2.1: A simplified representation of the reality-virtuality continuum as shown in (Milgram, Takemura, et al., 1994).

Another view of AR that has gained an extremely wide acceptance is Azuma's (1997) much-cited definition from an early survey on AR. In this seminal survey, Azuma describes AR as a variation of virtual reality that "allows the user to see the real world, with virtual objects superimposed upon or composited with the real world" (p. 3567). Looking for a definition that does not restrict AR to a specific technol-

ogy, Azuma defines AR as a system that has three decisive characteristics. It:

1. Combines real and virtual
2. Is interactive in real time
3. Is registered in three dimensions

This definition resurfaced in a somewhat more elaborate and accessible form in a follow-up survey by Azuma et al. (2001), where AR is defined in terms of systems that embody the following three characteristics. They:

1. Combine real and virtual objects in a real environment
2. Run interactively, and in realtime
3. Register (align) real and virtual objects with each other.

In addition to these two often-cited views, we can identify three prevailing ideas about the nature and characteristics of augmented reality that complement and reaffirm the ideas stated above. First, AR is commonly seen as a *technology*. Secondly, AR is often understood in terms of *visual* additions that are overlaid onto our *view* of the real world. Thirdly, AR is generally considered to *spatially align* this virtual content with the real world.

2.1.1 *AR as a Technology*

One of the most prominent understandings of AR is the idea of AR as a technology. For instance, Zhou et al. (2008), in their review of 10 years' worth of AR research presented at the primary AR conference ISMAR (International Symposium on Mixed and Augmented Reality) and its predecessor, describe AR as "a technology which allows computer generated virtual imagery to exactly overlay physical objects in real time" (p. 193). Comparably, Reiners et al. (1998) claim that "Augmented Reality is a technology that integrates pictures of virtual objects into images of the real world" (p. 31). A similar description is given by Roberts, Evans, Dodson, Denby, Cooper, Hollands, et al. (2002), who describes AR as "a technology that allows information stored digitally to be overlaid graphically on views of the real world" (p. 1) as well as by Doyle, Dodge, and Smith (1998) who describe AR as "a technology in which a user's view of the real world is enhanced or augmented with additional information generated from a computer model" (p. 147).

While AR is often seen as a technology, usually, these views do not limit AR to a specific hardware technology, such as head-mounted displays.[5] Rather, these views focus on what AR technology does. This brings us to the other two often mentioned characteristics of AR:

[5] In fact, researchers have been very explicit about not limiting AR to a specific hardware. For instance, Azuma et al. (2001) emphasize that they do not restrict their definition (see above) to "particular display technologies, such as a head- mounted display (HMD)" (p. 34).

virtual content is (1) visually overlaid onto our view of the world and (2) spatially registered (aligned) with real 3D space.

2.1.2 *AR as Visual Virtual Overlays*

Existing notions of AR are commonly focused on what a user or participant *sees*. Accordingly, AR is commonly understood in terms of virtual *imagery* that is overlaid onto a user's or participant's *view* of the world. This idea has already surfaced in some of the previously cited views about AR technology. In addition, this understanding is, for instance, shared by Piekarski and Thomas (2002), who describe AR as "the process of overlaying and aligning computer-generated images over a user's view of the physical world" (p. 36). Likewise, Rosenblum (2000) describes AR as "the overlaying of computer-generated imagery atop the real world using a see-through display" (p. 39). The media theorist Manovich (2006) provides a yet similar description and summaries AR as "the laying of dynamic and context-specific information over the visual field of a user" (p. 222). A focus on vision is also predominant in the research by Milgram and Kishino (1994) (see above), who discuss AR in terms of visual displays.

2.1.3 *AR as the Registration of Virtual Content in Real Space*

In addition to the view that virtual content is *overlaid* onto the real world, there is also the common belief that virtual content is spatially integrated into or aligned with the real 3D space. This spatial alignment is commonly called *registration*.

Registration is a common process in image processing, where it refers to the process of "transforming different sets of data into one coordinate system" (Rani and Sharma, 2013, p. 288). In medical practice, for instance, registration is used to combine images obtained with different types of technologies, and/or images obtained at different points in time (L. G. Brown, 1992). For instance, two medical images of a patient that have been taken at different moments might be registered with each other to find changes (L. G. Brown, 1992).

In the case of AR, registration usually works similarly. However, here, one image contains virtual content and the other image is the participant's view of the real world. Both images are combined in a way that the virtual content appears to exist at the right position in the world. Using AR technology in the medical context, for instance, a virtual indicator might guide a surgeon in performing a surgery. This virtual indicator does not have to be aligned with a previously obtained image but has to appear at the right spot on the real patient. In this sense, virtual content is registered with the real world.[6]

In the context of AR, registration can be thought of as giving virtual content a position in the physical world. This not necessarily has to happen visually, but, for instance, could also involve aligning virtual

[6] One common form of AR that works a bit different is so-called spatial augmented reality. Here, virtual content is embedded into the real world directly (e.g., projected onto the world), and not just integrated into a participant's *view*.

sound sources with the real environment.

In AR research, the term registration is typically used to refer to *spatial* registration. However, registration also has a temporal component. For instance, Craig (2013), explains that in AR, the added information "is in both spatial and temporal registration with the physical world" (p. 20). Simply put, this refers to the fact that information has to appear at the right position at the right time. For virtual objects to appear at the intended location in space, AR systems often have to take the view of the participant into account, compute a corresponding view of the virtual object in real-time and display it with little latency. If a participant moves and there is too much latency, the virtual content will appear at the wrong position and 'lag behind'. This is why many definitions (e.g., Azuma, 1997; Azuma et al., 2001) argue that AR systems have to be run interactively and in real-time. One can argue that an accurate spatial alignment implies an accurate temporal alignment: Virtual objects are not displayed at the correct position in space if they are displayed at this position at the wrong moment.

If we are to believe existing research, registration is necessary for AR. Most importantly, the claim that AR requires registration is part of the often-cited definition of AR by Azuma (1997) and Azuma et al. (2001) (see above), which describes AR in terms of systems that, among other things, align/register virtual and real objects with each other. In his original review, Azuma (1997) illustrates the implications of this requirement and suggests that AR does not include "[t]wo-dimensional virtual overlays on top of live video" because "the overlays are not combined with the real world in 3D" (p. 356). By now, Azuma (1997)'s definition of AR is commonly accepted (cf. Zhou et al., 2008), and with it, so is the need for registration.

The notion that AR requires registration is, for instance, shared by Thomas (2009), who describes AR as "the process of a user viewing the physical world and virtual information simultaneously, whereby the virtual information is registered to the physical worldview" (p. 105). Craig (2013), too, sees AR as "a medium in which information is added to the physical world in registration with the world" (p. 15). He later argues that in AR, the added information "is in both spatial and temporal registration with the physical world" (p. 20).

According to existing research, registration is not only necessary for AR—it also distinguishes AR from other related phenomena. For instance, Piekarski and Thomas (2004) mention registration as a distinguishing factor between AR and VR: "Although AR and VR systems share some similarities, AR is unique in that it requires the registration of the physical and virtual worlds" (p.164). In a somewhat similar line of thought, Craig (2013, p.30) uses the requirement of registration in order to distinguish AR from the more general field of Mixed Reality:

> Many people use the term mixed reality interchangeably with augmented reality. However, in this book I consider mixed reality to be

> a broader interpretation that consists of anything of both the physical world and the digital world. The specific constraint of registration is relaxed.

Furthermore, Bimber and Raskar (2005), mention the lack of registration as a reason why a TV showing a cartoon or a radio playing music are no AR displays.

2.1.4 *Composite Views*

The three ideas about AR appear in different combinations. For instance, the current definition of AR in the Oxford English Dictionary (*Augmented Reality* 2005, accessed 07-05-2016) combines the encountered notions and defines AR as "a technology that superimposes a computer-generated image on a user's view of the real world, thus providing a composite view".[7]

On first sight, the various reviewed descriptions complement each another perfectly and, together, draw a clear picture of AR: Augmented reality is a technology that combines virtual content and the real world by overlaying virtual imagery onto our view and spatially registering it with the real environment. This process happens interactively and in real-time.

On second sight, however, this rather clear image of AR is somewhat simplified and generalized. In reality, ideas about AR are more varied and complex. If we take a second look at the AR research landscape, we can find many more descriptions of AR, many of which differ from or even oppose the previously stated, popular image of AR as a technology that superimposes virtual images on a user's view and provides a composite view.

[7] We wonder whether *composite view* refers to the fact that visual additions are not simply displayed on top of what a participant sees but also spatially integrated, resulting in one coherent seamless view rather than an additional layer on top of the world.

2.2 *Less Common and Diverse Views*

AR research generally agrees that AR involves both the real world and some kind of additional—so-called virtual—information. Aside from this fundamental agreement, opinions about AR vary.

In particular, views differ with respect to (1) what AR is, including the question whether AR indeed is a technology, about (2) the nature of the virtual content, including the question whether *visual overlays* are actually defining for AR, about (3) the way the virtual and the real relate to one another, including the question whether AR really requires registration. Furthermore, there are many different ideas about (4) the real in AR. E.g., there are different views about the role of the user/participant in AR, and different ideas about what is actually augmented in augmented reality.

2.2.1 *The Nature of AR: A Technology?*

If we take a closer look at previously reviewed works, AR has not only been described as a *technology* (cf., e.g., Doyle et al., 1998; Reiners et al., 1998; Roberts et al., 2002; Zhou et al., 2008), but also as a mixed reality *display environment* (Milgram and Kishino, 1994), as *a system's process* of overlaying computer-generated imagery onto (the view of) the real world (cf., e.g., Piekarski and Thomas, 2002; Rosenblum, 2000) and as *a user's process* of viewing the real world and virtual information (which is registered with the real world) at the same time (Thomas, 2009). If we take a look beyond the previously reviewed papers, we can find even more views: Wikipedia's current description sees AR as "a live direct or indirect *view* [italics added] of a physical, real-world environment whose elements are 'augmented' by computer-generated or extracted real-world sensory input such as sound, video, graphics or GPS data" (Augmented reality, n.d.). Furthermore, Spence and Youssef (2015) describe AR as "an *experience* [italics added] of a physical, real-world environment whose elements have been augmented, or supplemented, by computer-generated sensory input" (p. 1). In addition, Klopfer and Squire (2008) describe AR as "a *situation* [italics added] in which a real-world context is dynamically overlaid with coherent location or context sensitive virtual information" (p. 205) and Graham et al. (2013) refer to AR as "the material/virtual nexus mediated through technology, information and code, and enacted in specific and individualised space/time configurations" (p. 222).

While these views simply present different ideas about the nature of AR, Craig (2013) explicitly opposes the trend of approaching AR as a technology. In his book "Understanding Augmented Reality" (2013), Craig writes:

> Throughout the entirety of this book, I consider augmented reality to be a medium, as opposed to a technology. By medium, I mean that it mediates ideas between humans and computers, humans and humans, and computers and humans. [...] By taking the stance that augmented reality is a medium, it will become much clearer how the technologies involved can be used to create compelling applications for a variety of purposes instead of as a mere technological novelty. (p. 1)

2.2.2 *The Virtual in AR: Beyond Visual Overlays*

While researchers generally refer to *virtual* content that is added to the real world, few researchers take the trouble of explicitly defining what they mean with "virtual". Notable exceptions are Milgram and Kishino (1994), who explicitly discuss the differences between the virtual and the real on three dimensions. They distinguish between (1) real and virtual objects, (2) direct and non-direct viewing and (3) real and virtual images. With respect to the first distinction, they consider real objects to "have an actual objective existence" (p. 1324) and virtual objects to "exist in essence or effect, but not formally or actually"

(p. 1324). Their second distinction becomes clear when we consider the viewing of real objects. Real objects can be viewed directly (e.g., through the air) or they can be sampled (e.g. filmed with a camera) and then reconstructed via a display (e.g., played back on a monitor). According to their definition, virtual objects cannot be sampled directly and thus always have to be synthesized. The final distinction takes into account whether an image has luminosity at the location where it appears to be located in the space. In contrast to a real image, a virtual image has no luminosity at the spot where it appears. Virtual images can not only exist as contents of a digital display—other common examples are mirror images and holograms.

If we consider AR literature in its entirety, the term 'virtual' is generally used to refer to the intangible or non-physical. In contrast, the real stands for materiality and physical existence. More specifically, the term 'virtual' is typically used to refer to *computer-generated* content (see, e.g., Azuma et al., 2001; van Krevelen and Poelman, 2010). Aside from these general trends, we find many descriptions of what forms this virtual content can take.

Commonly, the virtual is considered a *visual* overlay. However, at the same time, many AR researchers point out that AR is *not* limited to just visual additions, and hence not constrained to virtual visual overlays. For instance, Azuma (1997), in his widespread review of AR (see above), mentions that "Augmented Reality might apply to all senses, not just sight" (p. 361) and suggests that "AR could be extended to include sound" (p. 361) by sensing the world with microphones and adding synthetic 3D sound. Azuma et al. (2001) reaffirm this and point out that they do not limit their definition to the sense of sight. They emphasize that "AR can potentially apply to all senses, including hearing, touch, and smell" (p. 34). (In line with this, the definition by Azuma (1997) and Azuma et al. (2001) does not refer to the overlay of visual content but more generally, the combination of the virtual and the real.)

Milgram and Kishino (1994), in their popular paper on Mixed Reality Visual Displays, also briefly refer to the possibilities of mixing and spatially aligning computer-generated spatial sounds with natural sounds in the environment, as well as mention the possibilities of haptic and vestibular AR, both of which provide non-visual additional stimuli.[8] Aside from these explicit statements, many other researchers indirectly suggest that non-visual and multimodal content can play a role in AR, for instance, by providing broader definitions. One example is Craig (2013), who speaks of "digital information" that is overlaid on the physical world rather than of visual content and our view of the world.

Like Craig (2013) (see above), many researchers see AR as a field that deals with *digital* additions to the real world. The possibility of virtual content taking on non-digital forms is usually not considered.

[8] With haptic AR, Milgram and Kishino (1994) refer to haptic displays where synthetic haptic information is superimposed on existing haptic sensations. With vestibular AR, they refer to synthesized information about the participant's acceleration that contends with existing gravitational forces.

An exception is the article "Pre-Digital Augmented Reality" by Lamers (2013) which discusses the *Pepper's Ghost* effect as an early (and non-digital) form of AR. This effect makes use of a second hidden room, glass (or comparable materials) and special lighting in order to let virtual objects appear or disappear in a room, to change their transparency or to morph different objects into one another (Pepper's ghost, n.d.).[9]

Although most researchers seem to associate virtual content with digital content, researchers generally have different ideas about what *exactly* is added to the real world in AR. To mention just a few examples: when speaking of AR, various researchers refer to the addition of virtual *objects* (e.g. Azuma, 1997; Milgram and Kishino, 1994). Others refer to the overlay of *computer-generated images/imagery* (e.g. Piekarski and Thomas, 2002; Rosenblum, 2000) or to *synthetic sensory* information (Vallino, 1998).[10] Often, the differences lie in the details: Whereas Azuma (1997) summarizes AR as a field "in which *3D virtual objects* [italics added] are integrated into a 3D real environment in real time" (p. 355), others refer to information that is not necessarily *3D* nor necessarily considered an *object*. For instance, Kounavis et al. (2012, p.1) refer to superimposed "computer-generated data, such as text, video, graphics, GPS data and other multimedia formats [...]".[11]

Furthermore, *some*—but certainly not all—authors share ideas about the behavior of virtual content. For instance, Craig (2013) claims that AR not only allows us to *perceive* virtual content, but also, that we can interact with it in the same way as we interact with physical objects (e.g., Craig, 2013), or in other words: that the virtual content is *interactive* content. Furthermore, it has been suggested that virtual objects should behave like real objects: Herling and Broll (2011) state "[i]deally, the virtual content would behave exactly like real objects" (p. 255) and (S. Kim et al., 2011) write "[i]n order to make virtual objects move as if they coexisted with real objects, the virtual object should also obey the same physical laws as the real objects, and thus create natural motions while they interact with the real objects" (p. 25). At the same time, AR and VR pioneer Sutherland (1965), points out that virtual objects can behave differently from real objects. In his vision of future computer displays from 1965, he claims that "[t]here is no reason why the objects displayed by a computer have to follow the ordinary rules of physical reality with which we are familiar" (p.2).[12]

In general, we can notice a quest for realism on the one hand and a pursuit of objects that are unlike real objects on the other hand. For instance, Azuma (1997) mentions that virtual images "do not necessarily have to be realistically rendered in order to serve the purposes of the application" but that "[i]deally, photorealistic graphic objects would be seamlessly merged with the real environment" (p. 366). According to him, "the ultimate goal will be to generate virtual objects that are so realistic that they are virtually indistinguishable from the real

[9] The *Pepper's Ghost* effect can be traced back to 1584, when Porta first described an illusion called "How we may see in a Chamber things that are not" as part of his 20 volume book "Magia Naturalis" (Pepper's ghost, n.d.; Porta, 1658). However, the phenomenon has been popularized by John Pepper in the 1860s and is nowadays known as the Pepper's Ghost effect.

[10] Of course, these opinions are not necessarily exclusive.

[11] However, Azuma (1997), too mentions the overlay of text in combination with 3D wireframes.

[12] It should be noted that Sutherland (1965) did write about virtual objects in general as opposed to virtual objects that are supposed to appear in real space in. We will later raise the question whether virtual objects have to behave like real objects in order for this illusion to work (see chapter 5).

environment" (p. 380). In contrast, Craig emphasizes the possibilities of creating things that are different from real objects: "Indeed, one of the more interesting aspects of AR is that anything that can be created digitally [...], whether permutations of physical objects or objects that could not exist in the physical world" (p.18).

Finally, there are different opinions about *the role* that virtual content can play in AR. It is generally assumed that the virtual augments or extends the real world. However, there are other possibilities. More specifically, the virtual can also 'diminish' the real world and seemingly remove real elements from the perception of the participant. This is typically referred to as "diminished reality". Diminished reality is sometimes seen as its own field of research (e.g., Herling and Broll, 2010). Yet, diminished reality is also considered a subset of AR (e.g., Azuma et al., 2001).

As we have shown, there are many views about the virtual. While there is no generally agreed upon definition, AR is commonly associated with digital or computer-generated content that appears in real space. One might wonder: Why then, are advertisement screens and public information boards that display digital content in many cities, not considered a form of AR? AR researchers commonly agree that AR requires more than the mere combination of the virtual and the real in real space—that a stronger link between the virtual and the real is necessary. In particular, many researchers believe that the virtual has to be spatially integrated or registered in real 3D space. But is this really the case?

2.2.3 *The Link Between the Virtual and the Real: AR Without Registration?*

Judging from existing research, the registration of virtual content in real 3D space is widely accepted as a defining and necessary characteristic of AR. However, at the same time, an explicit interest in AR without registration is beginning to surface: The call for papers of the 14th edition of the International Symposium on Mixed and Augmented Reality (ISMAR 2015)—the leading conference on AR—lists *"augmented reality without 3d registration"* as one of the two emerging areas of particular interest and states that "[l]ightweight eyewear such as Google Glass can be used for augmenting and supporting our daily lives even without 3D registration of virtual objects." (ISMAR2015, n.d.).

In addition, there are views on AR that, although requiring relationships between the virtual and the real, do not require a *spatial* relationship between the two. For instance, Manovich (2006), refers to "dynamic and context-specific information" (p. 222) and claims that "a typical AR system adds information that is directly related to the user's immediate physical space" (p. 225). Similarly, Klopfer and

Squire (2008) speak of "coherent location or context sensitive virtual information" (p. 205). While these claims suggest that the information has to be *contextually* related, they do not claim that the information has to be *spatially* embedded in real 3D space.

In its totality, existing literature suggests that registration plays a key role in AR. At the same time, the different opinions about registration leave us wondering whether a spatial alignment of virtual and real objects is indeed always necessary. The above-reviewed positions show that AR can be approached more broadly and suggest that relationships between the virtual and the real in general (rather than only spatial registration in particular) play an important role in AR. We will take up the question whether registration is always necessary in chapter 3.

2.2.4 *The Real in AR*

AR researchers agree that the real world plays an important role in AR. However, opinions differ with respect to what role exactly the real world plays and what aspects of the real world are of importance.

As shown before, existing AR research is very focused on vision. In line with this, AR is often discussed in terms of virtual additions that are added to our *view* of the real world (Piekarski and Thomas, 2002; Roberts et al., 2002) or to *images* of the real world (Reiners et al., 1998). However, renowned AR researchers like Azuma (1997) and Milgram and Kishino (1994) suggest that AR not only applies to the visual sense. This view not only entails that non-visual virtual content can be added to the real world—it also means that non-visual aspects of the real world can play a role in AR. For instance, Azuma (1997) mentions the possibility of AR technology that makes use of microphones to sense the real sound environment (and that adds synthetic sounds or cancels out real sounds from this real sound environment). Furthermore, he considers the possibility of augmenting *the feel* of a real desk (rather than its visual characteristics), for example, "making it feel rough in certain spots" (p. 361). Milgram and Kishino (1994), too, mention the possibility of synthetic haptic information mixing in with real haptic sensations as well as synthetic sound sources mixing in with real auditory signals from the environment. As mentioned previously, the authors briefly consider vestibular AR, where the participants are affected by a mix of real gravitational forces as well as synthesized (vestibular) information about their bodies' acceleration.

It is a common understanding that some AR technologies allow participants to perceive the real world directly (via air or glass), while other technologies allow participants to perceive the real world in a mediated form (e.g., on an electronic screen) (see, e.g., Azuma, 1997; Milgram, Takemura, et al., 1994).[13] However, opinions differ as to whether the perceiver of the real world is merely an outside observer

[13] For instance, users can see the real world directly (via glass) with a so-called optical see-through HMD. In contrast, video see-through HMDs make use of (a) camera(s) to provide the user with a video of the real world that contains virtual elements (see, e.g., Azuma, 1997).

or an active part of this world: Several researchers simply claim that AR allows a user to *see* the real world (or a live video thereof), while others claim that the user/participant is also present in the environment, and hence, can not only see it but also act in and interact with this world (and, as we would like to add, perceive it with all their senses rather than just *see* it). For instance, R. Silva et al. (2003) consider the technology responsible for the television weather report AR because the real image of the news reporter (who actually stands behind a blue screen in the studio) is augmented with a virtual map. Similarly, many researchers (e.g., Van Krevelen and Poelman, 2010) mention sports broadcasting systems that embed virtual "first down" lines in the live broadcasts of football matches as examples of AR technology.[14] These systems allow the viewer at home to passively see a real environment, with virtual content superimposed and integrated into the environment. However, the viewer is not part of this depicted space. In contrast, Craig (2013) emphasizes that AR allows participants to "engage in an activity in the same physical world that [they] engage with whether augmented reality is involved or not" (p. 1) and clearly state that "[a]ugmented reality is interactive, so it doesn't make sense to watch it or listen to it" (p. 2). Likewise, Azuma (1997) points out that (in contrast to VR) "AR allows the user to see the real world" (p. 356), but also that "AR requires that the user actually be at the place where the task is to take place" (p.366).[15]

[14] In football, the virtual first down line marks how far the offense has to advance to gain a so-called 'first down'.

[15] Azuma assumes that the virtual content helps a user to perform a real-world task.

Finally, while researches generally agree that AR combines the virtual and the real, there is surprisingly little consensus on what is actually augmented by this virtual content. Many argue that it is the *perception* of reality that is augmented (e.g., Normand et al., 2012; Ross, 2005). Furthermore, there is the notion that in AR, the *physical world* (Craig, 2013) or our real physical *environment* (Milgram and Kishino, 1994) is augmented. At the same time, Milgram and Kishino (1994) also refer to the augmentation of the *display* of an otherwise real environment. In addition, there is also the notion of augmented *space* (Manovich, 2006). Wikipedia's current definition of Augmented reality (n.d.), provides yet another different perspective and describes AR as "a live direct or indirect view of a physical, real-world environment whose *elements* [italics added] are 'augmented' [...] ". Furthermore, Mackay (1996) also approaches this question in another way and considers the carrier of the physical equipment as augmented (e.g., the user is augmented when he/she carries a helmet and an object is augmented when sensors are embedded in it). Consequently, she distinguishes between an augmentation of the *user*, an augmentation of the *physical object* and an augmentation of the *environment* surrounding the user/object.

2.3 *Conclusion*

We have reviewed existing AR literature with the goal of learning more about the nature of AR, asking the question "What is augmented reality". Did we find an answer? Yes and no! The review has provided an answer insofar as it has revealed several widely accepted views and common notions about AR. At the same time, our review has also revealed opposing and diverging views that make us doubt these widely accepted views. In many ways, our review mirrors the results of our first look at AR examples (see chapter 1): On the one hand, AR seems to involve technologies that integrate virtual images into our view of the real world. On the other hand, AR seems to take many other forms as well.

Existing research largely agrees that AR combines virtual content and the real world. More specifically, our review has revealed three common notions of AR. First, AR is generally considered a technology or system.[16,17] Second, AR is understood in terms of visual virtual overlays that are placed over a user's view. Third, AR is considered to align virtual and real objects in physical 3D space. These views are not at odds but complement each other well. Together, they reveal a rather clear image of AR as a technology that integrates virtual imagery into our view of the world.

[16] Whereas 'system' is a broader term than 'technology', we here use technology and system as interchangeable. This is because here, both refer to technological systems consisting of both hardware and software components

[17] Yet, AR is usually not limited to a specific hardware or device. Instead, AR is characterized by what the system *does*.

This concept of AR covers many of the actual AR examples encountered in the previous chapter. A typical example of a system that embodies all of these aspects is the see-through head-mounted display technology that Caudell and Mizell (1992) proposed when they coined the term *augmented reality* (see chapter 1).[18] As discussed, their proposed system helps manufacturing and assembly workers by overlaying virtual guides and instructions onto their view of the real world. At the same time, their system also registers/aligns the virtual content with the real world. As a result, the virtual information appears to have a location in the real, physical world. For instance, a virtual arrow might indicate where to drill a hole (see figure 1.1).

[18] Caudell and Mizell (1992) refer to so-called "HUDsets", which refers to heads-up (see-thru) display head set. We use the term HMD (head-mounted display) as it is more commonly used for the type of display depicted by the authors and for consistency reasons.

If we consider both our first look at AR works from the previous chapter as well as existing notions of AR reviewed in this chapter, there is no doubt that technologies that integrate virtual imagery into our view play an important role in AR. However, we also have encountered views that diverge from this notion and that suggest that AR can take different forms as well. First of all, not everyone agrees that AR is a technology. Among other things, AR has also been described as an environment (Milgram and Kishino, 1994), situation (Klopfer and Squire, 2008) and experience (Spence and Youssef, 2015). Secondly, several definitions convey a more encompassing idea about the virtual and suggest that AR technology can work with non-visual content and engage all senses. In contrast to definitions that focus on visual overlays, modalities-encompassing definitions capture a broader variety of

AR works. For instance, they include works that augment real objects with virtual tactile textures (e.g., Bau and Poupyrev, 2012). Third, not all views require 3D registration. Several descriptions and definitions of AR simply call for a relationship between the virtual and the real. Such definitions also describe a broader field. If we do not demand 3D registration, AR can, e.g., include audio guides that automatically inform us about art pieces in a museum (e.g., Bederson, 1995).

Altogether, we have encountered a variety of different views regarding the nature of AR, the qualities of virtual content, the link between the virtual and the real, the role of the real world in AR, the position of the participant and the question of what is augmented in AR. These varying descriptions and definitions do not provide a clear answer to our question of what AR is. Rather, the different reviewed views give rise to several more fundamental questions: Does AR require visual overlays? Is AR interactive? Can we experience AR remotely? What forms can virtual content take? What is the role of the participant in AR? How is the real world involved? Is AR something we experience with our eyes, or something that engages all our senses? What is actually augmented in augmented reality? Is registration really necessary?

Ultimately, our review leaves us with two options. Either (1) the widespread image of AR as a technology that overlays virtual imagery onto our view and aligns it with the world is correct. In this case, we are simply confronted with many works that have wrongfully been labeled AR, as well as many definitions and descriptions that have failed to capture the essence of AR. Or (2) this widespread picture of AR is incomplete: it captures common characteristics of AR, but also disregards many other possible manifestations of AR.

In this thesis, we consider the second option and take the position that AR goes beyond this—arguably—stereotypical image. We build on the belief that augmented reality is—or potentially can be—more than just a technology that overlays virtual imagery onto our view, thereby providing a composite view (*Augmented Reality* 2005).

Whereas the common image of AR as a technology is pretty clear, this broader conception of AR remains rather blurry. Although we have encountered a variety of concrete examples (see chapter 1), it is unclear what forms AR can take. For instance, it is unclear whether and how AR might work with virtual tastes. Is it actually necessary and possible to register "virtual tastes" in three dimensions and in real space, similarly to how visual objects are integrated into the environment? Intuitively, this seems weird. After all, taste is not something we experience in three dimensions.[19] More generally, we wonder what characterizes AR, if not a system that integrates virtual imagery into our view of the world.

In the following chapter, we will explore this question and propose an alternative understanding of AR. We will approach AR from a deliberately broad perspective that does not limit AR to a certain set of

[19] Yet, it might be possible to make a certain taste appear in a participant's mouth, whenever he or she crosses a certain spot in space or places a certain object in their mouth.

technologies, that does not require 3D *registration* of virtual and real elements and that does not limit AR to the addition of *visual* information. In line with this, we will also treat the virtual as a broad and elastic concept that encompasses stimuli that either have been synthesized or that do not directly originate from their original source. This notion of the virtual includes, but is not limited to, computer-generated simulations. For instance, it also includes the possibility of treating audio recordings or perfume as virtual stimuli.

Chasing virtual spooks, losing real weight

Augmented running and a side trip into the history of audio augmented reality

A strange voice tells me to run. My heartbeat rises as I follow the instructions without giving them a second thought. The voice's manner of speaking reminds me of my parents' TomTom. The only difference: instead of telling me to take a turn, I am instructed to accelerate, slow down, to run or—if I am lucky—to walk. I am running with my new mobile app and virtual trainer. The app tracks every move, knows when my heartbeat rises and is supposed to help me gain speed and lose weight. Today, I run to clear my head after a mentally exhausting but physically unchallenging day. However, trying to catch my breath, my thoughts return to work. More precisely, I pore over my research topic, non-visual augmented reality.

In augmented reality (AR), virtual content is added to our real environment. Most often, this happens visually. By now, probably all of us have seen some three-dimensional objects popping up upon designated markers, virtual pink bunnies above augmented cereal boxes or walking directions superimposed on real streets. However, AR does not have to be visual. Sound, in particular, has already brought forth some fascinating AR applications and artworks such as Edwin van der Heide's *Radioscape* (2000-) and Theo Watson's *Audio Space* (2005). Entering the latter, visitors can hear the sounds left by previous visitors, spatialized, as if they were actually still there. At the same time, they can leave their own audio messages at any point within a room. It is not just the fact that the physical space is augmented with the ghost-like presence of previous visitors that makes me term this work AR. Visitors can also relate their own sounds and messages to those left earlier by others; thereby establishing connections between the virtual and the real. I imagine walkers, cyclists and other runners leaving their sound-trails behind on the road, leaving it up to me to add my own sounds and follow their steps, which are spread across time and space.

My favorite mobile app, *RjDj* (Reality Jockey Ltd., 2013), can also be considered AR sound art. The app remixes the sounds of the surroundings and provides you with a soundtrack to your life that blends in, makes use of and accompanies your environment. Although it is certainly no typical AR application, the relation between the sounds of the real environment and those produced by the app is so strong that often, they seem to melt into a single soundscape.

I will have to try this app while running. I can already hear the sound of my steps on the asphalt evolving, blending into a rhythmical soundscape, slowly displaced by the wind or heavy breathing, interrupted by pitched variations of my sudden greetings whenever I meet another runner.

While *RjDj* (Reality Jockey Ltd., 2013) and successor apps like *Inception - The App* (2016) and *The app formerly known as H _ _ r* (2016) are a rather recent phenomenon, the idea of remixing the sonic environment is not new. The artist Akitsugu Maebayashi has worked with similar concepts for a long time. His portable *Sonic Interface* (Maebayashi, 1999) was built in 1999—years before mobile phones gained comparable sound-processing abilities. The custom built device consists of a laptop, headphones and microphones

and uses delays, overlapping repetitions and distortions in order to recompose ambient sounds in urban space. The resulting soundscapes break the usual synchronicity between what one hears and what one sees. Unsurprisingly, Maebayashi is not the only one who has been exploring sound-based augmentations of the environment early on. In fact, audio augmentations of our environment have quite a history of their own. Unfortunately, they are less known in the context of AR and are often not even considered to be part of AR history.

"Walk!", my virtual trainer gives in to my exhaustion and I slow down. However, my thoughts keep racing. Quickly, they approach the early 1990s: Tom Caudell is believed to have coined the term *augmented reality*. It describes a head-worn display that superimposes visual information onto real objects (Caudell and Mizell, 1992). In Caudell's case, the new AR system helps workers assemble cables into an aircraft at Boeing. What usually goes unnoticed is that around the same time, Janet Cardiff started recording her so-called audio walks. Those walks are designed for a certain walking route and confront the listener with instructions such as "go towards the brownish green garbage can. Then there's a trail off to your right. Take the trail, it's overgrown a bit. There's an eaten-out dead tree. looks like ants" (Cardiff, 1991). While the listener navigates the space, he gets to listen to edited mixes of pre-recorded sounds, which blend in with the present sounds of the environment. Cardiff's virtual recorded soundscapes mimic the real physical one "in order to create a new world as a seamless combination of the two" (Cardiff, n.d.). By superimposing an additional virtual world onto our existing one, and thereby creating a new, mixed reality, Cardiff's sound art explores one of the key concepts of AR. And Cardiff is not alone with this idea; as early as 1987, Cilia Erens introduced sound walks, soundscapes and sound panoramas in the Netherlands. In contrast to Cardiff, she forgoes spoken content and uses largely unmixed everyday sounds. Yet, the effect is similar; they create "a new reality within existing realms, a form of 'augmented reality'." (Erens, n.d.) Clearly, the developments in non-visual AR were in no way inferior to the development of their visual counterparts. Taking slow steps, I imagine being on such a walk right now... Listening to instructions on which route to take, where to look, superimposed footsteps, sounds recorded here, on this path earlier, maybe altered with special effects. I imagine those sounds mixing in with the naturally present sounds of the river, bikes, and the occasional mopeds passing by.

"Run!", my trainer, whom I decide to call Tom, puts an abrupt end to this walk. The fact that AR sound art like Cardiff's and Erens' walks are not usually mentioned in the context of AR leaves me wondering what else we miss.

After Tom's instruction, my music fades back in. The song is intended to get me to run even faster. After my footsteps have adapted to the new rhythm it hits me: these instructions about how fast to run, the information about my heart rate, distance covered and calories burned and options such as racing against a virtual running partner in real physical space—they are just like AR.

In fact, my virtual running trainer shares most of the characteristics commonly found in AR applications. It adds another layer of content to my running. It is interactive and operates in real-time (cf. Azuma, 1997). Just like many other GPS based AR applications, it reacts to my position in the world. Most importantly, Tom fulfills my own, personal requirements for an AR experience: there is a relationship between the additional layer of content (the information I receive) and the real world (my running).

When another runner passes me slowly, my heart rate drops. I wonder whether it might be his heart rate that is mistakenly reported back to me. I am astonished, that without the sensor's help, I cannot even accurately perceive such basic and vital facts as my very own heart rate. Maybe this is farfetched, but with respect to that, the running app relates to the kind of AR applications which allow us to perceive things about the world that we normally cannot perceive, such as seeing heat, feeling magnetic fields or hearing ultra-high frequencies. (This idea is also discussed in sec-

tion 4.8.) So why are virtual Tom and his colleagues not considered to be AR?

Perhaps because there are also numerous differences between running apps and typical AR applications. To begin with, this running app does not augment the environment. Rather, it augments an activity—my running. And to be honest, despite the fact that Tom follows my every move—chasing a virtual competitor or running with a virtual trainer—it still feels like they are running on my phone while I have to tackle the real road. What is more, location-based AR applications usually display content related to the user's absolute position in the world. Tom, on the other hand, is only interested in the change of my position over time.

"Stop!", apparently, my position has changed enough. My run is over. The result: more than 500 kcal burned, five miles run and the revelation that the combination of the virtual and the real encompasses much more than just adding virtual visual objects to the real physical environment. There is a whole field of augmented activities as well! I cannot wait to jam with virtual bands, to try augmented eating or to take an augmented nap. As if to approve, my heart rate makes a last excited jump. Who knows, in the future, Tom might learn from existing AR. He might then have a look at my environment and direct my turns so that I discover new routes, point out sights or, when needed, help me find a shortcut home. Considering current developments in lightweight AR glasses, I guess it cannot be long until we can also see our virtual competitor passing by, are asked to design avatars representing our personal best time in races against other runners and are challenged to chase visual virtual spooks. I would not mind that. And I bet that that is when augmented running will be truly considered to be AR.

3 *New Perspectives*

The previous chapter has revealed three prevailing ideas about the nature and characteristics of augmented reality. First, AR is commonly seen as a *technology*. Secondly, AR is often understood in terms of *visual* additions that are overlaid onto our *view* of the real world. Thirdly, AR is generally considered to spatially integrate this virtual content in the real world by aligning virtual and real content with each other in 3D. All three ideas contribute to the widespread notion of AR as a technology that integrates virtual imagery into our view of the real world (see, e.g. *Augmented Reality* 2005; Reiners et al., 1998; Zhou et al., 2008). There is no doubt that such technologies play an important role in the context of augmented reality. Yet, in our opinion, such common understandings of AR are incomplete and unnecessarily limit the AR research field. In this chapter, we challenge the focus on technology, the need for registration as well as the emphasis on vision. We address shortcomings in prevailing definitions and propose alternative perspectives on AR. The proposed shifts in perspective are outlined below and subsequently discussed in detail.

The first issue with prevailing notions is their focus on AR as a *technology*. Generally speaking, technology-based definitions inform us about what an AR system does but do not reveal much about the AR environments they create and the AR experiences they evoke in the participant. Yet, the underlying purpose of AR technologies is to allow participants to experience augmented environments. Considering this, it only seems natural to also take the participant's experience into account and explore the augmented environments that they perceive. In our opinion, what a system does and whether it fits a given definition is less important than whether it evokes the intended experience. We thus believe we need to take an environment- and experience- oriented perspective. We will discuss this shift in perspective and address both the workings of typical AR systems as well as the experiences they facilitate in section 3.1.

The second issue with common notions of AR is their focus on the alignment of virtual content with the real world in three dimensions and in real-time. There is no doubt that this so-called registration process plays an important role in creating the impression of virtual objects existing in real space. However, in our opinion, there are three

reasons to look beyond registration. First, virtual objects can seemingly appear to exist in real space, even if they are not aligned with the real world in 3D. Second and more fundamentally, virtual content can be part of, enhance and augment the real world even if it does not seem to exist in the physical space. For instance, an audio guide can augment our experience of an exhibition without seemingly existing in the museum space. In our opinion, this means that registration is not always necessary for creating AR experiences. Third, registration might not always be sufficient to evoke AR experience. For instance, when attempting to display a virtual ball in real space, it might matter whether this ball appears to be affected by real light sources and whether the ball moves when it is hit by a real object. It is easy to imagine that a lack of interactions between the real world and virtual objects can harm AR experiences and make virtual objects look "out of place" even when they are spatially registered with the world. We thus believe that other links between the virtual and the real aside from spatial registration need to be considered in the context of AR. We follow this line of thought in section 3.2. We propose that instead of defining AR in terms of registration between the virtual and the real on a technological level, to define it in terms of a relationship between the virtual and the real on an experiential level.

The third concern that applies to many common notions of AR is the emphasis on vision. As we have seen, many existing views approach AR in terms of visual imagery that is overlaid onto a participant's view. We see three main issues with this. First of all, AR environments are not just something the participant can see. Rather, they are environments that participants can perceive with all their senses, act in and interact with. Arguably, AR is inherently multimodal and interactive because AR environments include the multimodal and interactive real environment. A second reason to look beyond vision is that virtual content, too, can take non-visual and multimodal forms. In our opinion, there is no good reason to exclude non-visual virtual content from the domain of AR. Last but not least, a multimodal perspective is important because of the way our human perception works: Even if visual information is added to our view of the world, this information can affect how we perceive non-visual qualities of the real world. For instance, visual information can alter how a physical object feels. (This effect is called cross-modal interaction.) If we only consider a participant's view of the world, such effects will remain unnoticed. Based on these arguments, we propose to approach AR as a multimodal and interactive environment rather than as a visual phenomenon. Section 3.3 presents this move from a vision-focused view towards a multimodal perspective in detail.

We synthesize and discuss these three views in section 3.4. We propose to define AR in terms of interactive and multimodal environments where a participant experiences a relationship between virtual

content and the real world. Our proposed view of AR departs from common understandings of AR in three ways: (1) it focuses on the AR environments and experiences rather than on AR technologies (2) it argues that AR is based on relationships between the virtual and the real rather than on interactive/real-time 3D registration (3) it treats AR as an interactive and multimodal rather than visual phenomenon.

Although we present these three points one by one, they are related and interdependent. For instance, our idea of AR experiences without the use of traditional AR technologies is supported by projects that make use of non-visual forms of virtual content. E.g., we can find examples of classical AR experiences that are realized with simple iPods or MP3 players in the context of sound-based AR. At the same time, the possibility of working with non-visual information, such as tastes, challenges the need for registering information with the surrounding world in 3D. After all, taste is not something we experience in three-dimensions and in the surrounding world, but something we experience in our mouth. At the same time, the move towards an experience-based view suggests that we should let go of the focus on 3D registration. In this way, the different points work together, support each other and build upon each other.

Although we challenge prevailing views, we do not mean to critique them on an individual level. For instance, the view of AR as a technology can make sense in an engineering context. Similarly, the claim that AR technology overlays virtual images onto a user's view makes sense in the context of a project that works with visual overlays. It is only natural that many authors describe AR from the perspective of their own domain and emphasize forms of AR that are relevant in their own research. Our notion of AR is meant to provide an additional, complementary perspective from which we can study and explore AR. It is not meant to replace other perspectives altogether.

3.1 *From Technologies to Experiences*

As we have seen in the previous chapter, AR is often seen as a technology or system. Most prominently, AR is considered an interactive system that combines and aligns the virtual and the real in 3D and in real-time (Azuma, 1997). But what is the point of such an AR system? What is its purpose and what is in it for the user? This section addresses these questions.

3.1.1 *The Goal of AR Technologies*

Why do AR technologies exist? What is their purpose and what goals do they serve? A look at existing research reveals some common answers: AR technologies aim at creating the illusion of virtual objects existing in the real world, and more generally, try to make it appear

as if the virtual world and the real surroundings were one seamless environment. For instance, Vallino (1998) states that "[t]he goal of augmented reality systems is to combine the interactive real world with an interactive computer-generated world in such a way that they appear as one environment" (p. 1). Furthermore, Buchmann et al. (2004) propose that "[t]he goal is to blend reality and virtuality in a seamless manner" (p. 212). Billinghurst, Clark, et al. (2015), who survey almost 50 years of AR research and development, similarly state: "From early research in the 1960's until widespread availability by the 2010's there has been steady progress towards the goal of being able to seamlessly combine real and virtual worlds" (p. 73). More specifically, AR systems are commonly used to create scenarios where virtual objects appear to exist in real, physical space. E.g., Regenbrecht and Wagner (2002) state that "[t]he goal is to create the impression that the virtual objects are part of the real environment " (p. 504). Likewise, Azuma (1997, p. 356) mentions that "[i]deally, it would appear to the user that the virtual and real objects coexisted in the same space " (p.356).[1]

[1] Note that Azuma is using the word 'coexist' differently from how we use it. With coexist, we emphasize that there is no relationship between two things and that they exist independently. Azuma uses the term to refers to things that appear to exist in the same space, which implies a spatial relationship.

If we look at the AR landscape, indeed many so-called AR applications present us with virtual objects that seemingly exist in our otherwise real surroundings. To mention just a few examples: The IKEA Place app allows us to see virtual furniture in our physical environment (*IKEA Place* 2017). Likewise, the HoloLens by Microsoft (n.d.) seemingly fills our living rooms with visual virtual building blocks. Similarly, the app *Sphero* (2011) turns a robot ball into a visual virtual beaver that seemingly exists in our everyday surroundings. An example of the latter is shown in figure 3.1. This screenshot shows the little virtual beaver *Sphero* (2011), as seen through an iPad.

Figure 3.1: The virtual beaver *Sphero* (2011) is not just overlaid onto our view but integrated into our view. The picture is a screenshot showing the image displayed on the iPad. (The screenshot was taken by the author.)

As this example shows, the virtual content appears to exist in the space around us—the beaver seems to be standing on the authors living room floor, looking at the author's cat.[2]

[2] In AR literature, we often find claims that virtual content is (a) overlaid onto our view or (b) integrated into our view. This difference can be explained with the fact, that technically speaking the content is often overlaid. At the same time, however, it is also aligned with the real world in three dimensions and appears to exist in the space. In this sense, is integrated into the view.

In the following, we will discuss how AR systems achieve such effects. This look at the workings of AR technology is necessary for two

reasons. First, it will allow us to better understand common technological views on AR. Second, knowing how typical AR systems work allows us to show that different, alternative technologies can also be used to create AR experiences.

3.1.2 *How (Visual) AR Technologies Work*

AR systems can make it seem as if virtual objects were present in the real world. How does this work? Simply put, a typical AR system senses the participant's position in the world and consequently computes how a virtual object has to be presented so that it appears to exist in the real world. Once the virtual image is computed, it is displayed to the participant, e.g., on a head-mounted or hand-held display.

The Registration Problem

The process of giving a virtual object a position in the real space is called registration, and according to common notions (see section 2.1), characterizes AR. In his book *Understanding Augmented Reality: Concepts and Applications*, Craig (2013, p.17) explains registration like this:

> A key element to augmented reality rests with the idea of spatial registration. That is, the information has a physical space or location in the real world just like a physical counterpart to the digital information would have.

As discussed in subsection 2.1.3, registration is a common process in image processing, where it refers to the process of "transforming different sets of data into one coordinate system" (Rani and Sharma, 2013, p. 288). In the context of AR, registration typically refers to the alignment of virtual and real content. Strictly speaking, descriptions of this process vary slightly. For instance, Drascic and Milgram (1996) use registration to refer to the alignment of "the coordinate system of the virtual world with that of the real world" (p. 129). In addition, registration is also understood as aligning virtual and real objects with respect to each other (Azuma et al., 2001). However, in the end, registration always refers to a process that makes sure that virtual content has a position in the real world.

The challenge of properly aligning the virtual and the real is commonly referred to as "the registration problem" (e.g., Azuma, 1997), and regarded one of the key issues in AR research (e.g., Azuma, 1997; Bimber and Raskar, 2005; You and Neumann, 2001). Accurate alignment is considered important because improper registration of virtual and real objects can cause virtual objects to appear as if they existed separately from the real world, rather than in the real world. In other words, improper registration can compromise or break the illusion of virtual objects existing in real space (cf., e.g., Azuma, 1997; Bajura and Neumann, 1995; Vallino and C. Brown, 1999). In addition to breaking the illusion of virtual objects existing in real space altogether, inaccu-

rate alignment by an AR system might cause virtual objects to appear at a *wrong* position in real space. For instance, (Azuma, 1997) suggests that inaccurate registration could cause a virtual pointer to appear at an incorrect position: "[...] many applications *demand* accurate registration. For example, recall the needle biopsy application. If the virtual object is not where the real tumor is, the surgeon will miss the tumor and the biopsy will fail." (p. 367, italics in original). Likewise, Bajura and Neumann (1995) explain that "[i]f accurate registration is not maintained, the computer-generated objects appear to float around in the user's natural environment without having a specific 3D spatial position" (p. 52).

The effect of improper registration can, for instance, be seen when playing the game Pokémon GO. Here, virtual creatures often appear in unrealistic positions in the environment or look like an independent overlay that floats on top of the camera feed, rather than as part of the environment. Screenshots of such moments are presented in figure 3.2.

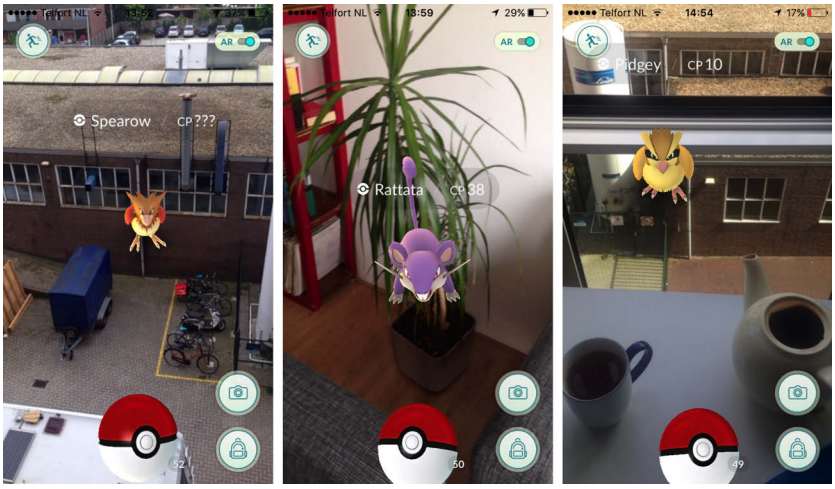

Figure 3.2: Due to inaccurate registration, virtual Pokémon creatures can appear at unrealistic positions or overlaid onto the live view, rather than as part of the real environment. (In order to amplify this effect, the author has manually moved the phone in space. However, Pokémon quite regularly appear 'detached' from the real world without trying to achieve this.)

Proper registration is particularly difficult because participants can move through AR environments and experience the world from different perspectives. For instance, we do not want a virtual cup of coffee to move in space, simply because we are moving our head. Also, if we stand up and look at the cup from above, we expect to see it from this particular perspective and, e.g., expect to see the cup's contents. Simply put: the virtual information has to dynamically adapt to our movement and perspective, in order to continuously appear correctly positioned in the real world.[3] What is more, for a virtual object to appear on top of, inside of, behind or otherwise related to a real object, the AR system needs to know the position of such real objects (Azuma et al., 2001).

[3] The possible movement of the participant also explains why many definitions (e.g., Azuma, 1997; Azuma et al., 2001) not only require registration but also point out that the AR system has to work interactively and in real-time.

Tracking

In order to accurately respond to changes in the participant's location and orientation, a variety of so-called tracking technologies are used to keep track of the participant's position in the real world (or of the position of a mobile device, through which the participant perceives the augmented environment). Often, computer-vision-based systems are used to determine the position and orientation of a participant (or of the intermediate device) (Craig, 2013). These systems make use of cameras in order to sense the world. Based on what the camera 'sees', software determines where the camera must be located and how it must be oriented in order to obtain this view of the world. For this to work, the environment must contain some cues that the software can recognize. These cues can take many forms. In the early days of AR, the cues typical took the form of so-called "fiducial markers" (see figure 3.3), which were physically integrated into the environment and specifically designed so that computers could easily recognize them. Currently, however, many efforts are put into markerless tracking and into using natural features of the environment, such as buildings and objects, as cues.[4]

Computer-vision-based tracking has the advantage that it is rather precise. Furthermore, the software can not only keep track of the location of the participant but also recognize and track objects of interest. As a result, virtual content can be positioned relative to real objects. For instance, a computer-vision-based system might be able to recognize a vase and display a flower in it. (Because of this, the stem of the flower can be hidden by the real vase. Furthermore, the flower can remain in the vase, even if the vase moves in space.)

[4] AR without markers is also referred to as markerless AR. The use of natural features for tracking is referred to as natural feature tracking (NFT). This concept can also be used to recognize magazine pages, photographs, posters or products and ultimately display virtual information on top of them. In these cases, the line between marker-based and markerless AR is blurry. For instance, a photograph can act both as an object that is augmented, as well as serve as a marker that is added to a scene to allow for tracking. Hence, NFT and markerless tracking overlap, but are not the same.

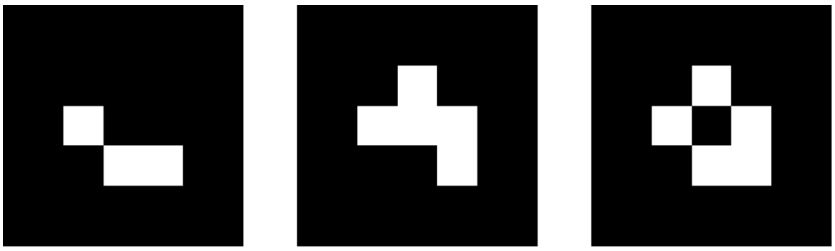

Figure 3.3: Three typical fiducial markers that can be recognized by AR software, such as the popular open-source ARToolKit tracking library. (The displayed markers are part of the download of the *ARToolKit SDK* (1999).)

Another common approach to tracking (and ultimately, registration) is the use of positioning systems, such as GPS (Global Positioning System) in order to obtain the location of the participant (or the location of the used device) in 3D, in combination with a compass, gyroscope and accelerometer to ultimately determine all six degrees of freedom of the participant.[5] This approach has the advantage that the required technologies are currently widely available, and integrated in many smartphones. Unfortunately, such smartphone-based solutions often also have several disadvantages. First of all, they can suffer from poor

[5] Six degrees of freedom (6DoF) refers to the six ways a rigid body can move in three-dimensional space. The possible movements include three ways of changing the *location*: (1) surging (moving forward and backward on the X-axis), (2) swaying (moving left and right on the Y-axis) and (3) heaving (moving up and down on the Z-axis) and three ways of changing the *orientation*: (4) rolling (tilting side to side on the X-axis), pitching (tilting forward and backward on the Y-axis), and yawing (turning left and right on the Z-axis) (Six degrees of freedom, n.d.).

accuracy. For instance, Blum et al. (2012) compared the accuracy of the orientation and location of iPhone 4, iPhone 4s and Samsung Galaxy Nexus phones. They found mean location errors of 10-30 meter as well as mean compass errors around 10-30°, both with high standard deviations that, according to the authors, render them unreliable in many settings. Also, GPS is especially unreliable indoors and in urban areas, where GPS signals can be blocked by high buildings (Cui and Ge, 2003). Furthermore, because the camera image is not analyzed, the application has no knowledge about the spatial structure of the physical environment and cannot track other objects of interest. As a result, they cannot be used to align virtual content with respect to a real object. In other words, GPS-based solutions are fine for displaying a virtual bird in the real sky, but not for showing a virtual flower in a physical vase (the accuracy would be too low), especially if this vase can be moved around (the system would not be able to recognize the vase and track its movement).

In addition to computer-vision and GPS-based approaches to tracking, other possibilities exist. For instance, the AR system by Feiner, Macintyre, et al. (1993), which helps with the maintenance of an office printer, make use of ultrasonic transmitters and receivers mounted on both the participant's head and on the printer in order to determine the spatial relationship between the two. Furthermore, many applications combine several different methods and sensors in order to obtain better (more accurate) results. E.g., Persa (2006) use a firewire webcam and a GPS receiver in combination with a radio data receiver to obtain position and orientation information. Similarly, the PhD thesis by Caarls (2009) focuses on fusing information from various sensors with different accuracies, update rates, and delays to address the challenge of real-time pose estimation of a user's eyes.

Computing Virtual Output

Once positioning data is obtained, the AR system typically uses this information to compute (or, in the case of images "render") the corresponding virtual output. If you are looking at my desk with an AR device, the information can, e.g., be used to compute a believable image of a virtual cup of coffee on my desk. If you are looking at the desk straight from above, the rim of the computed cup will have a circular shape. If you change your perspective slightly, the rim will have an elliptical shape. Ideally, the appearance of a virtual object changes depending on the participant's perspective, just like the appearance of real objects varies when one changes one's point of view. Commonly, this computed content takes a visual form.[6]

[6] However, as we emphasize throughout this thesis, virtual content can also take non-visual forms. For instance, the song of a bird could be synthesized in a way that it becomes louder if the participant gets closer and in a way that the song appears to originate from the same tree, even when the participant turns around and changes their orientation.

Display

As soon as the corresponding virtual output is computed, it is pre-

sented to the user. Most often, AR systems present virtual content in real space by means of a head-worn or hand-held display. However, other possibilities exist. Virtual content can, e.g., also be embedded into the world directly with projectors or flat panel displays. For instance, Benko et al. (2014), use three projectors to allow two participants to see virtual content in the real environment, and, for instance, toss a virtual (projected) ball back and forth through the space between them (see figure 3.4) Such forms of AR where virtual content is directly embedded into the real world is typically referred to as spatially augmented reality (Raskar, Welch, and Fuchs, 1998) or spatial augmented reality (Bimber and Raskar, 2005). Furthermore, in addition to visual displays, also other types of stimuli are sometimes used to convey the presence of virtual objects in real space. E.g., the Sound-Pacman game by Chatzidimitris et al. (2016) makes use of synthesized 3D sound played back on headphones in order to give virtual ghosts a position in the real physical environment and communicate their location to the player.

Figure 3.4: The projection-based AR project by Benko et al. (2014) can make it seem as if virtual objects existed in real space, rather than projected onto the world. The image shows two screenshots from the YouTube video about this project (Microsoft Research, 2014).

In order to accurately register the virtual and the real even when a participant moves, these processes have to happen in real-time and with very little latency. If the registration process takes too long, the delay can cause registration errors. For instance, if you were to turn your head very fast, the virtual cup of coffee on my desk might not be able to keep up with you. In the time it would take the system to figure out your perspective and compute and display the cup of coffee, your perspective would have already changed so much that the resulting output would no longer match your view. Simply put, virtual content has to appear at the right position at the right time. This is why it is sometimes stated that an AR system has to operate interactively and *in real-time* (e.g., Azuma, 1997), and that the virtual not only has to be registered with the real world spatially but also *temporally* (e.g., Craig, 2013).

The Greatest Common Factor

As the examples above indicate, AR systems take many different forms. They can, e.g., present various forms of virtual content (e.g., visual or auditory content) and use different information displays to convey this information (e.g., screens, projectors or headphones).

Furthermore, displays can be placed in the environment statically or carried by the user (and in the latter case, can be head-mounted or hand-held). Different setups go hand in hand with different system requirements. For instance, tracking the participant's position is not necessary in cases where virtual content is projected onto the real world directly with a projector in order to change *surface* attributes of physical objects, such as their texture or color because here the rendering is independent of the viewers position (Raskar, Welch, and Chen, 1999).

Although AR systems differ, they generally make use of a computer system that registers the virtual with the real world interactively, in real-time and in three dimensions (Azuma, 1997). In the following, we refer to this type of technology as *traditional AR technology* or *traditional AR systems*.

Given the common goal of making it seem as if virtual objects existed in real space, defining AR in terms of traditional AR systems can seem like a natural choice. After all, traditional AR systems can enable this illusion. More than that, without an AR system that registers the virtual and the real, virtual content typically appears to exist independently from its real surroundings as opposed to as part of the world. Without registration, a virtual character might, e.g., appear on a screen, a voice might appear "on a sound recording", or a text might simply overlay what we see—rather than seemingly exist in the real surroundings. This happens, for instance, with the virtual overlays presented by the Google Glass device (see figure 3.5). The overlays are not registered with the real world in 3D, and appear *on top of our view*, rather than *integrated into* space.

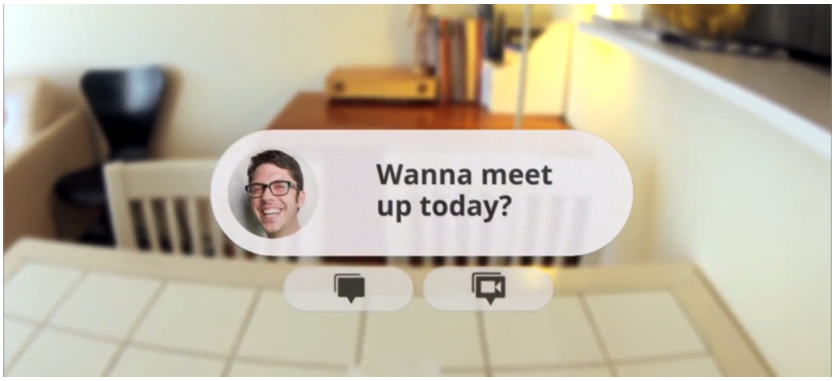

Figure 3.5: A mock-up of the Google Glass concept. Virtual content is overlaid onto the view of the real world but not registered with the real space. This image is a screenshot from a video demonstrating the concept behind Google Glass (Huzaifah Bhutto, 2012). The actual realization of the overlays looks quite a bit different and can be seen in figure 3.8.

If typical AR systems create the desired illusion of virtual objects existing in real space while other types of systems do not create this illusion, why not define AR in terms of typical AR systems? In our opinion, there are two answers to this question. First, *alternative AR technologies exist*: While rare, different types of technologies can also make it seem as if virtual objects existed in real space. In other words,

the assumption that other types of systems cannot create the desired AR experience is wrong. Second, *alternative AR experiences exist*: Although this illusion is commonly desired, virtual content does not have to seemingly exist in real space in order to contribute to, enhance or otherwise augment this environment. We will demonstrate the first point in the following, and pick up the second point in section 3.2.

3.1.3 *Typical AR Experiences With Alternative Technologies*

There is no doubt that interactive AR systems that align virtual and real content in 3D can make it look as if virtual objects were part of real space and merge virtual and real worlds. However, if we only understand AR in terms of traditional AR systems, we miss one crucial aspect: other types of technologies likewise can create the desired effect. In the following, we will discuss three examples that illustrate that we do not need a typical AR system to blend the virtual and the real and to make virtual objects appear in real space. Next to visually augmented reality, we will also consider sound-based forms of AR. We do this because different types of virtual content might blend in with the real world in different ways.

Forest Walk

Early examples of AR experiences and environments that work without the use of traditional AR systems include Janet Cardiff's audio walks (Cardiff, n.d.), such as *Forest walk*.[7] *Forest walk* can be described as a "soundtrack" to the real world, specifically recorded and mixed for a pre-determined walking route. The track includes multiple layers of recordings, such as the sounds of Cardiff walking in the forest, her footsteps, the sound of her hand brushing tree bark, the sounds of the forest, such as crows, voices and in particular, Cardiff's voice, talking about the environment, giving walking instructions and describing her surroundings. For instance, one can hear Cardiff say "Go towards the brownish green garbage can. Then there's a trail off to your right. Take the trail, it's overgrown a bit. There's an eaten-out dead tree. Looks like ants." (Cardiff, 1991), while navigating the particular environment Cardiff is talking about.

One thing that makes Cardiff's recordings special is that her virtual soundscape relates to the real environment. This relationship happens on several levels: For one, Cardiff's recordings describe the real space. Instructions such as "Ok, there's a fork in the path, take the trail to the right." refer to the real surroundings and lead the way. Furthermore, the used sounds have been recorded on the same site where they are later on experienced by the participant. Consequently, the recorded sounds are similar to the real surrounding soundscape. According to Cardiff, this similarity is important for the soundscape to mix in with the real environment. As Cardiff herself puts it: "The virtual recorded

[7] Cardiff is neither the only nor the first artist to work with audio walks. For instance, Celia Erens, a sound artist from the Netherlands, has realized a series of works that present pre-recorded 3D soundscapes in the real sound environment. Also, "Forest Walk" is not the only walk by Cardiff that illustrates our point. However, as it is the first in Cardiff's series of audio walks, and, unlike Erens' work, also includes spoken text and instructions, we have chosen this particular example.

soundscape has to mimic the real physical one in order to create a new world as a seamless combination of the two" (Cardiff, n.d.).

Another aspect that characterizes Cardiff's soundscape, is that the used sounds have been recorded in binaural audio. Binaural audio is a recording technique that captures the spatial characteristics of the sound in 3D and consequently provides a 3D audio experience (rather than the usual stereo distribution of the sound) when the recording is played back on headphones.[8] Binaural audio often results in a very realistic impression. To quote Cardiff: "it is almost as if the recorded events were taking place live" (Cardiff, n.d.). Cardiff mixes her main walking track with several layers of sound effects, music, and voices, in order to create "a 3D sphere of sound" (Cardiff, n.d.). Judging from Cardiff's descriptions and our own experience with binaural audio, the pre-recorded sounds appear to originate in the real environment.[9]

So what does this work have to do with AR? Little, if we take a conventional, technology-based perspective on AR. Instead of using an AR system, Cardiff's work makes use of a simple CD player (or iPod/MP3 player). There is no system that aligns or registers the virtual sound sources in real three-dimensional space.[10] Instead, the sounds are placed in the space more loosely: The participant is told where to start the walk and press play. Also, the audio mix includes instructions that tell the participant where to go and that guide their attention. Indirectly, these instructions affect the participant's position in and movement through the environment, and consequently, also roughly determine where the virtual sound sources appear in space.[11] However, although Cardiff's walk does not make use of typical AR technology, it yet shares fundamental similarities with typical AR projects: It allows us to experience the real environment, supplemented with virtual content. More than that, it makes us experience a seamless, mixed, partially virtual, partially real environment. It is such a seamless combination of the virtual and the real, which is commonly considered to be the *goal* of AR (cf. section 3.1).

[8] Binaural audio is based on the fact that hearing makes use of two signals: the sound pressure at each eardrum (Møller, 1992). If these two signals are recorded in the ears of a listener (or a dummy head), the exact 3D hearing experience can be reproduced by playing the signals back on a headset.

[9] This claim that sounds indeed seemingly originate in the real surrounding was confirmed by Zev Tiefenbach, the studio manager of Cardiff/Miller, who in turn confirmed this with Janet Cardiff (personal communication).

[10] If the listener turns their head, the recorded sounds will move along—they have no fixed position in real 3D space but are always relative to the position and head of the listener.

[11] One potential reason why this loose alignment suffices is that the recorded sounds not necessarily have to appear at a specific position in the surrounding space. For instance, no exact 3D registration is necessary when dealing with flying elements such as crows, as it does not matter where exactly they appear in the environment.

Mozzies

Another application that makes virtual objects appear in real space without a traditional AR system is the early mobile game *Mozzies*. This game was installed on the Siemens SX1 cell phone that launched in 2003 (López et al., 2014). The mobile application used to show flying mosquitos, overlaid on the live image of the environment captured by the phone's camera. Players could shoot the virtual mosquitoes by moving the phone and pressing a button when aiming correctly (Siemens SX1, n.d.). In contrast to Cardiff's work, the game makes use of an interactive system. However, the application does not make use of registration in the traditional sense, but instead, 'only' uses the camera as a motion sensor (Siemens SX1, n.d.) and applies 2D motion detection (Reimann and Paelke, 2006). Yet, judging from the images

that can be found of this (and similar) games online, it appears as if mosquitoes were flying through the space in front of the phone's lens. An impression of this can be seen in figure 3.6, which shows a similar game running on a Nokia N95.

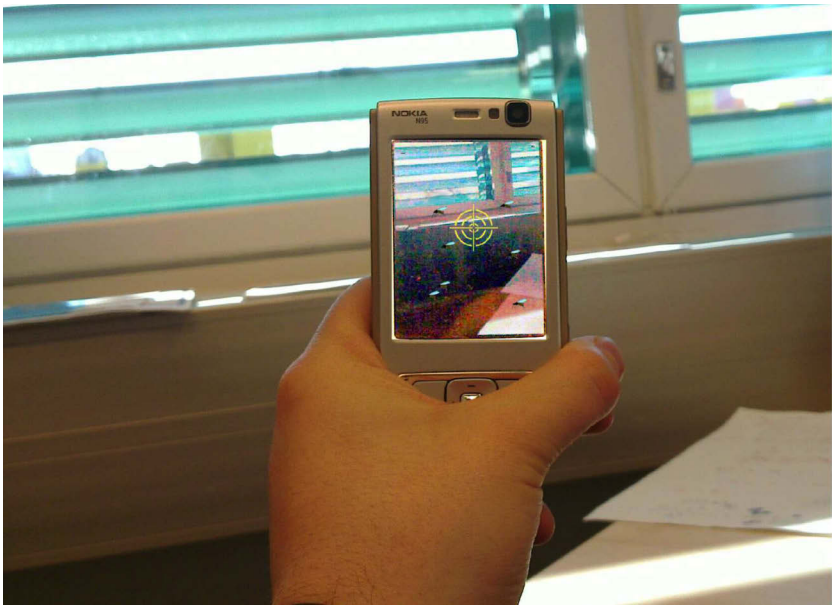

Figure 3.6: A game similar to the *Mozzies* game. Virtual mosquitoes appear to be flying in the space before the phone's lens. The image appeared in a paper by López et al. (2014), and permission to use the image in this thesis was granted by Miguel Bordallo Lopez.

Presumably, this works because mosquitoes 'only' have to appear to be flying *somewhere* in the surrounding space rather than at an exact position. To achieve this, exact registration seems not to be necessary. However, because the creatures are not registered in 3D, is not possible to walk around the virtual insects and look at them from all directions and angles. Furthermore, the virtual mosquitoes cannot disappear behind real objects. Due to the lack of first-hand experience, it remains open how one experiences these issues when quickly moving and turning the device.

NS KidsApp

A third example of AR experiences that work without typical AR systems is the NS KidsApp. This mobile application by the Dutch railway operator *Nederlandse Spoorwegen* (NS) is primarily aimed at children (and their parents) and it introduces a short story with the two characters *Oei* and *Knoei*. When starting up the application, it becomes clear that Knoei has missed the train, and that as a result, Oei and Knoei are not traveling together. It is then up to the user of the application to spend time with Knoei during the train journey.

There are several playful assignments for the player that allow them to make videos with Knoei appearing in the otherwise real environment. In these assignments, the player is asked to point their phone at a particular spot or have someone else point the phone at them and

film them, while they are at a certain location. For instance, one assignment asks the players to put the camera against the window and film the outside. As a result, one can see Knoei flying next to the train in a superman kind of fashion on the phone's screen. Another assignment asks users to point the camera at the typical place-name signs that can be found on Dutch train stations. The resulting view of the scene on the phone shows Knoei swinging on the place-name sign. Yet another assignment invites the player to sit next to Knoei, while someone else is pointing the phone at them and filming. When doing so, the one filming can see Knoei hovering over a train chair, showing off his muscles to his neighbor (see figure 3.7).

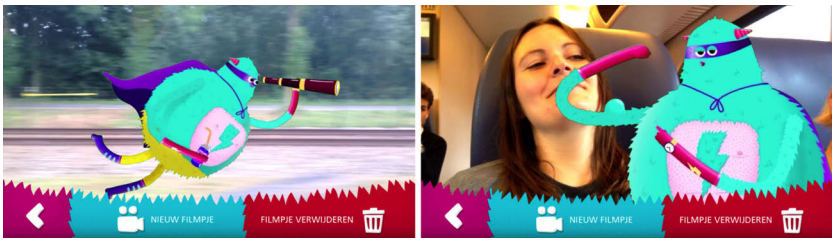

Figure 3.7: The NS Kids app shows Knoei flying next to the train (left) as well as next to the player showing off his muscles (right) on the camera feed. Screenshots by Jurriaan Rot and Hanna Schraffenberger.

This application, too, creates the illusion of virtual content existing in real space, without the use of a traditional AR system. Instead of a system, the *participant* can align the virtual and the real. Like in Cardiff's case, instructions are part of the game. These instructions make sure that what the participant sees will serve as a fitting background for the virtual overlay.

3.1.4 *AR Technologies Versus AR Experiences*

The previous sections have shed some light on the workings of traditional AR systems and the kind of experiences they aim to create. It has become clear that AR systems can make it seem as if virtual content existed in real space and as if the virtual and the real were one seamless environment. The creation of such mixed virtual-real environments and the presence of virtual content in an otherwise real environment seem to be primary goals of AR practice. However, we have seen that similar experiences and environments can also be achieved without traditional AR technologies. For instance, instead of an AR system, participants can align virtual and real content.

A question that we thus have to ask ourselves is what actually is defining for AR—the augmented environments and unique experiences that we hope to create or the technologies we develop in order to create them? Do we unnecessarily limit AR, if we only consider scenarios where an AR system registers virtual content with the real world in 3D? Do system-based definitions actually capture what we are ultimately interested in?

The answer to these questions remains a matter of opinion. Independent of one's individual position, it is clear that there are two sides to augmented reality: On the one hand, AR systems and on the other hand, the participant's experiences when using the system. If we want to truly understand and advance AR, a focus on either one alone will not suffice.

Personally, we take an experience- focused point of view. This is because AR technologies are meant to create augmented environments that a participant can experience. Arguably, the sole purpose of AR systems is for participants to use them and to experience augmented environments. Accordingly, we believe what matters most, is not what an AR system does, but what the participant experiences. If we ultimately aim at creating certain environments and experiences, why define the field in terms of the technologies that enable them rather than in terms of the environments and experience we are actually interested in? An environment- and experience- focused definition will hold, even if enabling technologies change or take unforeseen forms. We thus propose to define AR in terms of the unique environments a participant experiences, rather than in terms of certain types of systems.

So far, we have identified one key form of AR, namely otherwise real environments in which a participant experiences the presence of additional virtual objects. However, other types of AR experiences might exist as well. In fact, we suspect that virtual content can augment the real world even when it does not appear to exist in the physical space. We will address this possibility in the following section.

3.2 *From Registration to Relationships*

Registration is widely seen as a defining and necessary characteristic of AR (see, e.g., Azuma (1997); Azuma et al. (2001); Bimber and Raskar (2005)). There is no doubt that registration is important to AR. The previous section has shown that it can play a key role in making it seem as if virtual objects existed in real space. However, we believe there are three reasons to look beyond registration and to challenge the common focus on spatial alignment. First of all, making virtual content seemingly exist in real space does not always require 3D registration. The previous section has already shown that alternative approaches to placing virtual content in real space exist: For instance, Janet Cardiff's audio walks (Cardiff, n.d.) do not incorporate 3D registration, yet communicate the presence of virtual content in real space. Also, some settings require less strict forms of registration. E.g., an exact alignment might not be necessary when dealing with flying objects. Second and more fundamentally: The illusion of virtual content existing in real space, which motivates the need for registration, might not be necessary for AR in the first place. Arguably, not all forms of AR require

for virtual objects to seemingly exist in real space! Simply put, other types of relationships (aside from spatial registration) between the virtual and the real are possible, potentially facilitating other forms of AR experiences. For instance, virtual content can inform us about the real world, and by doing so supplement and augment (our experience of) the real world. Third, we have to look beyond registration because registration alone might not always suffice to create the intended AR experience. For instance, it might not only be necessary to present a virtual object at the right position but also necessary to apply a realistic illumination in order for virtual objects to appear as if they existed in real space. Because the first argument has been discussed in detail (see subsetion 3.1.3), we will focus on the second and third point in the following.

3.2.1 *Alternative AR Experiences*

In this section, we challenge the need for registration and explore alternative forms of AR experiences that are not based on 3D registration and that do not entail the apparent existence of virtual objects in real space. In particular, we explore the idea of augmentation through content-based relationships between the virtual and the real. We present two examples that illustrate this concept and where the virtual contributes to, extends and augments our environment by informing us about it.

Audio Guides

The idea of virtual additions that inform us about the real world is common in the cultural sector. For instance, many museums provide additional information in the form of audio tours that guide the visitor through a museum, and which supplement the real world and ideally, enhance our experience of the exhibition. In our opinion, such audio tour guides can accompany a user and augment a user's experience of their real surroundings, even if they do not appear to be spatially present.

We are not alone with the opinion that audio tours and audio guides can be considered AR. For instance, Bederson (1995) argues "[o]ne place a low-tech version of augmented reality has long been in the marketplace is museums. It is quite common for museums to rent audio-tape tour guides that viewers carry around with them as they tour the exhibits" (p. 210) . Furthermore, Rozier Rozier (2000), refers to audio tours as "perhaps the earliest form of 'augmented reality"' (p. 20).

Whereas audio guides typically provide factual information about the real surroundings, other possibilities exist. An artist that takes the idea of audio tours one step further is Willem de Ridder. In 1997, de Ridder realized an audio tour in the "Stedelijk Museum" in Amster-

dam that told visitors about the meaning of 'invisible' elements in the museum (*history and archive - Stedelijk Museum Amsterdam* n.d.). This shows that the virtual information can relate to the surroundings more freely. In fact, one could argue that Ridder's words also have the power to place imaginary virtual objects in a real environment and that they can create the experience of virtual objects existing in real space.[12]

[12] However, it can be argued that these objects are imaginary rather than virtual.

Google Glass

The concept of using a virtual layer of information to enhance our everyday lives has also been on the basis of the Google Glass project. Google Glass is essentially a head-mounted display in the shape of eyeglasses. A small display in one corner presents additional information (such as text and/or images) as an overlay on top of a user's view of the world.

The information displayed by Google Glass can be completely unrelated to a user's context (e.g., a random text message from a friend) but it can also relate to the user's real surroundings. For instance, the device can be used to translate text present in the real environment in real time, to overlay driving instructions onto a driver's view or to access relevant information in the kitchen (see figure 3.8).

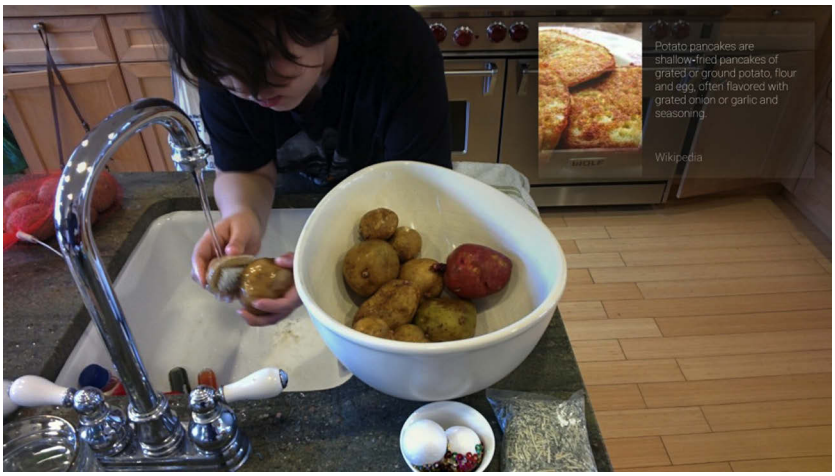

Figure 3.8: Google Glass can overlay information that relates to our real surroundings and context. This image is a screenshot taken with the device, illustrating the user's view. Image created by and courtesy of Ben Collins-Sussman.

The role of Google Glass in AR is controversial. As we know, 3D registration is commonly considered necessary. This view excludes all Google Glass applications from the realm of AR. However, the 2015 call for papers of the leading AR conference ISMAR (International Symposium on Mixed and Augmented Reality) argues that "[l]ightweight eyewear such as Google Glass can be used for augmenting and supporting our daily lives even without 3D registration of virtual objects". In line with this, some researchers consider systems like Google Glass in the context of augmented reality. For instance, Liberati and Nagataki (2015) consider Google Glass an AR device, and distinguish

among two types of current and future AR glasses: (1) AR glasses that inform the user about their surroundings and provide "informational text" to the user and (2) AR glasses that present additional objects, that are embedded in the real world and that potentially can interact with the real world as if they existed physically.[13]

If we apply Liberati and Nagataki (2015)'s distinction, Google Glass can act as an AR device, and falls in the first category of glasses, as it presents text (as well as other media) that informs us about our surroundings[14]. According to Liberati and Nagataki (2015), the information provided by such glasses modifies the objects they inform us about because the participant can change their attitude towards the objects based on the information.

We, too, believe that virtual information can modify (our perception of) real objects. Arguably, it can add to and affect our experience of the real world and in this sense become part of and augment the environment. However, we believe such augmentations are possible independently of how the virtual information is presented. In other words, information can augment our surroundings no matter whether it is, e.g., overlaid with AR glasses, displayed on a phone's screen or delivered by a recorded voice on headphones.[15,16] In our opinion, the question whether virtual content augments the real world (or vice versa) is not about the device we use, or the medium used to present such information. Instead, what matters is whether the presented content is experienced in relation to the real world. (This is likely the case when the two are inherently related on the content-level.) In line with this, the question whether Google Glass creates AR experiences depends on whether the presented information is perceived in relation to the real environment.

[13] In our opinion, the two categories are not exclusive. For instance, text can appear in the form of object-shaped letters that are integrated into the real environment and that seemingly interact with real objects.

[14] As we know, Google Glass can also present unrelated text. This is why we say that it can *act* as an AR device rather than that it *is* an AR device.

[15] In many ways, information defies the terms virtual and real. Arguably, information can have the same effects, no matter whether it is presented virtually or physically.

[16] In fact, we have to ask ourselves whether it actually matters whether the information is presented in a virtual form or, for instance, presented by a real person or on a physical information board. One can argue that information is never something physical, and always can affect and augment our experience of the world.

3.2.2 *Registration Without AR Experiences*

The previous examples have shown that spatial 3D registration is not the only link between the virtual and the real that allows us to experience virtual content as part of or in relation to the real world. Content-based relationships between the virtual information and the real environment, too, can facilitate the experience of an augmented environment. We thus believe there are different forms of augmentation aside from the apparent presence of virtual content in real space.

Another reason to look beyond registration is that registration alone might not always be sufficient in order to create AR experiences. This seems particularly relevant when it comes to the common goal of making virtual objects appear in real space. Here, many other relationships between the virtual and the real aside from spatial registration potentially contribute to the resulting experience. Among others, it can make a difference whether a virtual object appears to be affected by real light sources. For instance, Drettakis et al. (1997) claim: "Provid-

ing common illumination between the real and synthetic objects can be very beneficial, since the additional visual cues (shadows, interreflections etc.) are critical to seamless real-synthetic world integration" (p. 45). Sugano et al. (2003) go one step further and hypothesize that "[w]ithout shadows providing depth cues a virtual object may appear to float over a real surface even if it was rendered on the surface." (p. 76). In other words, registration alone might not suffice to create the desired effect. The subsequent experiment by Sugano et al. (2003) shows that presenting virtual objects with shadows as opposed to without shadows creates a stronger connection between virtual objects and the real world and increases the virtual object's presence in the world. (However, their research does not seem to support the idea that virtual objects appear completely detached from the real world due to the lack of shadows.)

In addition to optical interactions, a lack of other physical interactions and/or social interactions between real objects and virtual objects can potentially harm AR experiences and make virtual objects look "out of place" or appear as if they existed independently from the real world. For instance, Breen et al. (1996) point out: "For the new reality to be convincing, real and virtual objects must interact realistically" (p. 11). Likewise, S. Kim et al. (2011) write: "In order to make virtual objects move as if they coexisted with real objects, the virtual object should also obey the same physical laws as the real objects, and thus create natural motions while they interact with the real objects." (p. 25). Accordingly, for a virtual ball to appear as a believable part of real space, it might be necessary for it to bounce back when it hits a real wall. More than that—if we expect a realistic response, this movement might not be enough—the ball might also have to create a corresponding sound.

Furthermore, we can imagine that the presence of a virtual creature in the real environment is much more convincing if this creature seems to be able to perceive the environment and react to stimuli in the surroundings. For instance, a virtual creature might seem more present if it listens and responds to the sounds in the environment or if it sees and reacts to the participant when they are right in front of it.[17]

At the same time, the illusion of virtual elements being present in the space might be harmed if such interactions and perceptions are missing. For instance, it might disturb us if a virtual creature is not affected by real wind, if it is not reflected in real glossy surfaces or if it remains dry when it rains.

A first indication, that other factors aside from spatial registration indeed can affect the experienced presence of virtual objects in real space can be found in figure 3.9. In our opinion, the fact that the real cat does not seem to be aware of the virtual creature hurts the illusion of the virtual object actually being present in the space.

Unfortunately, a lack of empirical research makes it impossible to

[17] The idea of virtual creatures being more aware of their surroundings has been addressed by the developers of Pokémon GO with their AR+ update (Niantic, Inc., 2017). In this version, Pokémon seem to sense the player's movement. Consequently, players can scare virtual creatures away by approaching them too abruptly.

Figure 3.9: My cat shows no sign of awareness of the virtual beaver Sphero *(2011)*. According to our experience, this can harm the experience of Sphero being a part of the real environment. The picture is a screenshot showing the image displayed on the iPad. (The screenshot was taken by the author.)

conclude whether 3D registration is always sufficient to evoke AR experiences, i.e., to make participants experience virtual objects as part of or as related to the real environment. However, in our opinion, it is clear that other types of relationships also can facilitate and shape AR experiences. This should be reason enough to look beyond registration and consider relationships between the virtual and the real in general.

The notion that virtual objects should be able to sense and interact with the real world entails that we look beyond spatial registration and consider how the virtual and the real relate to one another on non-spatial levels. The idea that virtual content might have to react to non-visual aspects of the real world in order to appear as a believable part of the environment indicates that there is more to AR than what a participant sees. We will discuss this idea and in particular, the understanding of AR as a multimodal environment in section 3.3.

3.2.3 *Registration Versus Relationships*

In the preceding sections, we have argued that 3D registration between virtual content and the real world is only one of several ways to shape AR experiences. We believe that augmentation cannot only emerge from the *registration* of the virtual and the real but generally results from the *relationships* between the virtual and the real. In line with this, we believe that the spatial (and typically but not necessarily visual) presence and apparent existence of virtual content in the real environment is only one form in which the virtual can augment the real. Arguably, the virtual can also augment the real in different ways; e.g., by informing us about the surroundings. This, of course, raises one crucial question: If real-time registration by an interactive system in 3D is no defining factor, what then does define AR?

In our opinion, all AR scenarios have one characteristic in common: the virtual is experienced in relation to the real world and vice versa. Accordingly, we believe we are dealing with AR if one important criterion is met: the participant experiences a relationship between the virtual and the real. We thus propose that instead of defining AR in terms of registration between the virtual and the real on a technological level, to define it in terms of a relationship between the virtual and the real on an experiential level.[18]

So far, we have identified two key forms of AR. First, cases where a participant experiences the presence of virtual content in their otherwise real surroundings. Here, virtual content seemingly exists in real space, rather than, e.g., on a screen or in a separate virtual world. Second, environments where the virtual is experienced as pertinent to the environment on a content-level. The first form of AR is typically (but, as shown in subsection 3.1.3, not always) based on 3D registration of virtual content in real space. The latter form of AR, however, does not require registration. Rather, the virtual is likely experienced in relation to, as part of or as pertinent to the real environment because there is an inherent relationship between the virtual and the real in terms of content.

The question whether or not to define AR in terms of relationships between the virtual and the real remains a choice. In our opinion, AR comprises all cases where virtual content is experienced in relation to the real environment. We thus propose to broaden the view of AR and focus on the various possible relationships between the virtual and the real that facilitate such experiences. Presumably, there are many more relationships that still can be discovered. For instance, if we think about movies, a soundtrack can certainly become part of a scenery, although it is not *spatially* integrated into the movie. The virtual and the real might blend on such musical, non-spatial levels in AR as well.

Even when one disagrees with our view, it should be clear that spatial registration is not the only link between the virtual and the real that can shape AR experiences. Other relationships that play a role, e.g., include physical and social interaction between virtual and real objects.

[18] Of course, a mere link between the virtual and the real does not guarantee that the participant also experiences this relationship. What is more, a participant might experience relationships that have never been created or intended. For instance, a museum visitor might listen to a virtual museum guide, but associate the information with the wrong artwork. Consequently, they might not experience the intended relationship but experience another link instead. Similarly, one and the same scenario might be experienced as AR by one person but not by another. However, we believe it is safe to assume that by carefully considering and crafting relationships between the virtual and the real, we can shape AR experiences.

3.3 *From Visuals to Multimodal and Interactive Environments*

AR is commonly understood in terms of virtual *imagery* that is overlaid onto a user's or participant's *view* of the world (see chapter 2). Accordingly, AR is thought "to 'augment' the visual field of the user" (Caudell and Mizell, 1992, p. 660) or to provide a "composite view" (cf. *Augmented Reality* 2005). In this sense, much AR research is focused on what a user or participant *sees*.

However, if we approach AR from a participant's point of view, the resulting AR environments are not just something visual. Rather, they

are interactive, multimodal environments that potentially engage all of the participant's senses and that invite the participant to act in and interact with the space.

In our opinion, there are at least three reasons why we have to approach AR in terms of multimodal and interactive environments rather than focus on what a participant sees. First, the real world is not something visual but has multimodal and interactive qualities. Furthermore, virtual content, too, can exhibit non-visual qualities and allow for multimodal interaction. Finally, even purely visual virtual content that is superimposed onto our view can affect how we perceive non-visual stimuli. In the following sections, we develop these arguments in detail.

3.3.1 *The Multimodal and Interactive Real World*

If we approach AR from a participant's point of perspective, virtual content is experienced as part of or in relation to the otherwise real world. This world is not just something participants can *see*. Rather, it is a world that participants can perceive with all their senses, act in and interact with: we feel the ground beneath our feet, hear our footsteps, move over when a bike bell rings, we knock on doors and open them and engage in conversations with other people.

Although many AR systems focus on what a participant sees, non-visual qualities of the real environment often to play an important role in the overall resulting experience. For instance, in Caudell and Mizell's case of an AR system that helps assembly workers with virtual instructions (cf. section 1.1), it is crucial that the worker can touch, feel and physically interact with real objects: the worker might, for instance, drill a hole, connect wires, or place sticky fabric at the right spots. The ultimate goal of Caudell and Mizell's prototype is to support the worker with his actions in the world. Although the *virtual component* is strictly visual and intangible, the resulting augmented environment is more than what the user sees. Clearly, the system had little purpose, if the user could only see the augmented environment.

And Caudell and Mizell's project is hardly the only project where non-visual qualities matter. To mention just a few more examples: Participants need to touch and manipulate the world when they repair their printer with virtual instructions (Feiner, Macintyre, et al., 1993). A surgeon might listen to audible feedback of medical monitors, interact with colleagues, and of course, perform a surgery with the help of virtual indicators. Players of AR games like Pokémons GO walk through space, talk to our friends, and hopefully, hear a car bonk when they try to catch virtual Pokémons on real streets. Judging from these examples, AR is more than meets the eye.

If we understand AR in terms of the mixed virtual-real environments that a participant experiences rather than in terms of technolo-

gies, multimodality is the norm rather than the exception. In contrast to what widespread claims imply, AR not only "*might* [italics added] apply to all senses" or "*could* [italics added] be extended to include sound" (Azuma 1997, p. 361), but rather already applies to all our senses. AR is inherently multimodal, simply because it includes the multimodal real world. When dealing with augmented reality, we have to remind ourselves that AR not only "allows the user to see the real world, with virtual objects superimposed upon or composited with the real world." (Azuma, 1997, p. 3567) but that we can also hear, smell, touch, and taste this world.

Just like AR environments are multimodal because they entail the multimodal real environment, AR environments are also interactive simply because the real world allows for interaction. Whereas the fact that AR engages all our senses is often overlooked, the interactive qualities of AR are well known. For instance, Craig (2013) emphasizes that "[a]ugmented reality is interactive, so it doesn't make sense to watch it or listen to it" (p. 2) and argues that "that the way people engage with augmented reality is to experience it" (p. 1). Likewise, Hugues et al. (2011) point out that AR is not only something a participant can see but an environment that allows for action: "we define AR by its purpose, i.e. to enable someone to create sensory- motor and cognitive activities in a new space combining the real environment and a virtual environment" (p. 47). Furthermore, the fact that AR allows the participant to act in the world—and hence, choose their own perspective—has been key to AR's technological development. One of the most prominent topics in AR research is tracking techniques that make sure a virtual object's visual appearance matches the participant's current viewing perspective even when the participant changes their point of view (cf. Zhou et al., 2008).

Considering that participants experience multimodal environments that allow for interaction, we believe it makes sense to approach AR as a multimodal and interactive rather than solely visual phenomenon. However, the characteristics of the real world are not the only reason to do so. Another reason to think about AR this way is the fact that virtual content, too, can take multimodal and interactive forms.

3.3.2 *Multimodal and Interactive Virtual Objects*

How does it feel to touch a virtual object, to run one's fingers over it? Is there a chance that we can burn our hands when doing so? How does a virtual object taste or smell, what sound does it make if we shake it and how heavy is it, if we want to carry it around?

If we look at existing AR projects, tools and technologies, the answers to these questions can be disappointing. The majority of existing projects and devices allows us to view virtual content, rather than to experience it with all our senses. Sure, quite some virtual objects also

produce sounds. However, generally speaking, virtual objects in AR cannot be felt, have no smell, taste, weight, temperature or other physical properties that we know from real objects. To mention just a few examples: The *Sphero* (2011) project allows us to steer a virtual beaver through our living room, but we won't feel the beaver's fur if we try to pet it (instead we feel a robot ball). Likewise, the game Pokémon GO allows us to see virtual creatures in our everyday surroundings, but we cannot use dogs to sense their smell and chase them. Similarly, while virtual avatars can carry make-up, they typically cannot wear any perfume. Whereas touching physical artworks in a museum can get us in trouble, touching AR art (e.g., Veenhof, 2016) usually is safe but also boring: we will not feel anything if we try. Commonly, what it comes down to is that virtual objects can be *seen*—the typical virtual object is first and foremost a visual object.

While most AR projects make use of visual overlays that are superimposed onto a participant's view, there are exceptions that show that virtual content does not have to equal visual content. For instance, we have already thoroughly discussed Cardiff (1991)'s *Forrest Walk* that makes use of audio recordings. In the following, we will briefly point out a few more projects that illustrate that virtual content in AR can also take sonic, haptic, gustatory, olfactory and multimodal forms.

One of the works that use sound rather than visuals to convey the presence of virtual objects in real space is the SoundPacman game (Chatzidimitris et al., 2016). This game is an audio AR version of the traditional PacMan game. However, here all game elements are seemingly placed in the real, physical surroundings and the information about their position is provided solely by means of 3D sound. Using audio rather than visual cues clearly provides different possibilities. Among other things, the use of audio allows participants to perceive ghosts even if they are not in their direct line of sight and, for instance, positioned behind them.

An example that shows that we might perceive virtual content in real space haptically rather than visually has been realized by Bau and Poupyrev (2012). Their REVEL device injects electrical signals into a participant's body and thereby allows participants to feel virtual textures when running their fingers over real physical objects. As such, the system augments real physical objects with virtual tactile textures.

A project that aims at altering the taste rather than the tactile feel of real elements has been realized by Nakamura and Miyashita (2011). The authors propose a system that makes use of a fork or chopsticks connected to an electric circuit and thereby changes the taste of food. Similarly, they propose to change the taste of drinks by using two straws that are connected to an electric circuit. The experienced change in taste happens because the tongue is stimulated with electric current.[19]

Finally, also scents can be used in the AR context. For instance,

[19] This effect is nowadays known as "electric taste" and was discovered by Sulzer as early as 1752 (Bujas, 1971).

Lindeman and Noma (2007) suggest that under-nose displays, air-canon displays and scent emitters in the environment can be used to present the participant with a mix of computer-generated and real-world scents.[20] Also, Yamada et al. (2006) have proposed a wearable olfactory display that can present virtual odor sources in an outdoor environment. Their proposed setup takes the position of the virtual odor source as well as the position of the participant into account and varies the strength of the presented odor accordingly. This allows the system to simulate the spatial spread of the odor in the real environment. Although this project has been presented in the context of VR, it can be considered an AR project. This is because their system can simulate the existence of odor sources in an otherwise real environment.

[20] The term 'computer-generated' can be a bit misleading when it comes to scents. We thus propose the term 'computer-controlled' scents.

As the previous examples illustrate, AR can also work with non-visual virtual information. In addition, AR can make use of multimodal content that combines different types of sensory information. Researchers have, among others, used force feedback devices such as the Phantom in combination with HMDs in order to create viso-haptic virtual objects. For instance, Bianchi et al. (2006) have demonstrated the use of a Phantom device in an AR-based ping-pong game. In their setup, a virtual bat is attached to the haptic device and allows players to interact with a virtual ball. The player can not only see the virtual ball via a head-mounted display but also feel its impact on the simulated bat via the haptic device. Another example of an AR project that uses multimodal virtual content is the mobile (smartphone-based) AR game *GeoBoid* by Lindeman, G. Lee, et al. (2012). In their game, players are surrounded by flocks of virtual geometric creatures called GeoBoids. These creatures are represented both visually as well as by means of spatialized audio.

3.3.3 *Multisensory Perception and Cross-Modal Effects*

So far, we have argued that we have to treat AR from a multimodal perspective because both the real world and virtual content can engage all our senses. A third reason to treat AR from a multimodal perspective is the way our human perception works. When we perceive the world around us, our brain combines information from various sources. As Ernst and Bülthoff (2004, p. 162) point out:

> To perceive the external environment our brain uses multiple sources of sensory information derived from several different modalities, including vision, touch and audition. All these different sources of information have to be efficiently merged to form a coherent and robust percept.

For instance, when we sit at our desk, our arms resting on it, our fingers drumming on it, we can see, hear and feel the desk. These different sensory streams of information are integrated into our coherent perception of the desk. When information from different sensory modalities is combined, different sensory stimuli can interact with one

another. For instance, what we hear might influence our tactile experience and what we see can influence where we perceive a sound. Such influences, where information from one sense affects how we experience information from another sense are referred to as cross-modal effects and cross-modal interactions.

A popular example of cross-modal interaction is the "Parchment-skin illusion". According to Jousmäki and Hari (1998)'s findings, the sounds that accompany hand-rubbing can influence the tactile sensation of the skin. It was found that emphasizing high frequencies can make the skin feel rougher. Another popular cross-modal illusion is the McGurk effect (McGurk and MacDonald, 1976). In this illusion, what we see can affect what we hear. When being presented a video of a person saying "ga-ga", dubbed with the sound of a voice saying "ba-ba", participants in the study by McGurk and MacDonald (1976) reported hearing "da-da". This shows that different sensory stimuli can not only complement each other but also interact with each other to create a different experience than the sum of the individual experiences.

The fact that our perception is multi-sensory also plays a role in AR. Because interactions can occur between different sensory modalities, visual virtual information might affect our perception of real non-visual characteristics of the environment. In other words, even visual virtual overlays that are superimposed on a participant's view can affect what the participant perceives with other senses. More generally, virtual sensory information can interact with and affect how we perceive 'real' sensory information (information originating from the so-called real, physical world).

As it turns out, such cross-modal interactions between virtual and real stimuli are not only a theoretical consideration. Various AR projects have already utilized the phenomenon of cross-modal interaction and explicitly used visual virtual information to transform our experience of non-visual qualities of the world. For instance, Hirano et al. (2011) and Sano et al. (2013) use an HMD to display different computer-generated deformations on an object, when it is pushed down by a participant. Their experiments show that the perceived softness can be manipulated by means of visual virtual dents, without changing the actual material: The larger the visual dent caused by pushing the object, the softer seems the object. Other projects similarly show that virtual *visual* information can alter the perceived *temperature* (Ho et al., 2014), *texture* (Iesaki et al., 2008) and *center-of-gravity* (Omosako et al., 2012) of real objects. (These examples will be discussed in more detail in section 4.7).

As these examples show, visual virtual information can affect how we perceive non-visual qualities of the world. Even more possibilities arise if we present non-visual and multimodal virtual stimuli (cf. section 3.3.2). One of the projects that combine the idea of presenting

multimodal virtual information with the concept of cross-modal interaction is the MetaCookie project. Because this project supports our presented arguments for a multimodal perspective on several levels, we will discuss this project in more detail in the following.

The MetaCookie Project

The *MetaCookie+* headset[21] (Narumi, Nishizaka, et al., 2011a,b) aims at changing the flavor of a real plain cookie. The project is based on the idea that virtually changing the look and smell of a plain cookie might affect its perceived flavor. Consequently, the headset changes the visual appearance of a plain cookie and, for instance, makes it look like a chocolate, almond or cheese cookie (see figure 3.10). At the same time, it also features an olfactory display with scents that match the visual choices.

[21] A first versions of this system has surfaced under the name "Meta Cookie" (Narumi, Kajinami, et al., 2010b).

Figure 3.10: The MetaCookie+ headset. Reprinted from Narumi, Nishizaka, et al. (2011a). Reprinted under fair use.

In order to use the *MetaCookie+* system and experience the different tastes, the participant needs a real plain cookie with a special AR marker on it (see figure 3.10). (The marker makes it possible for the system to keep track of the cookies position).

When eating the cookie, the participant wears a custom head-mounted visual and olfactory display.[22] Before placing the cookie in their mouth, the participant can select a cookie of their liking from a list of options, including, for instance, chocolate, almond and cheese. After the participant has chosen their preferred cookie, an image of the selected cookie is integrated into his view at the position of the

[22] In the case of *Meta Cookie+*, the olfactory display consists of several air pumps, scented filters and a controller and is able to eject six types of scented air and fresh air.

real cookie, effectively making it look as if they were holding the cookie of their choice. In addition, the olfactory display dispenses a scent matching that of the selected cookie when the augmented cookie is within a range of 50 cm from the participant's nose. The strength of the scent linearly increases the closer the cookie moves to the participant's nose. When the participant is about the eat the cookie (and the cookie is in front of their mouth) the system produces the most intense version of the smell for about 30 seconds. This strong odor is produced to emulate retronasal olfaction (the stimulation of olfactory receptors via the mouth rather than nose, which typically results in a stronger sensation than stimulation via the nose). Because the smell is presented for about half a minute, the scent is assumed to be presented longer than it takes the participant to eat the entire cookie.

The authors evaluated their system (presumably very informally) with "a dozen people" in its initial version 2010. Its later incarnation was more systemically evaluated it with 15 participants (Narumi, Nishizaka, et al., 2011b) as well as 44 participants (Narumi, Nishizaka, et al., 2011a). Based on their initial trials (Narumi, Kajinami, et al., 2010b), the authors report that almost all participants perceived a change of the taste of the plain cookie. Similarly, the results from their later evaluation (Narumi, Nishizaka, et al., 2011a,b) suggest that Meta Cookie+ can change a perceived taste and allows users to experience different flavors, solely by changing the visual and olfactory information.

In our opinion, the project supports our argument for a multimodal view on AR in three ways. First of all, by augmenting the taste of a real cookie that one eats, it emphasizes that the real world is a multimodal world that we interact with and not a visual world that we look at. In other words, it emphasizes that the real component in AR entails more than what we see. Second, by displaying scents, it emphasizes that the additional information we present to participants is not limited to visual information. As such, it illustrates that the virtual component in AR, too, can be more than what we see. Finally, it builds on the concept of cross-modal interactions and thus shows that our senses do not work in isolation. This means, that we cannot simply treat what we see as independent of what we hear, smell, taste, or otherwise perceive. Naturally, this is not only true in the real world but also in AR.

3.3.4 *Visual Overlays Versus Multimodal Environments*

In the previous sections, we have presented various reasons to treat AR from a multimodal rather than vision-focused perspective.

The first reason is that AR experiences entail the real world and that this real world is a multimodal world. This means that multimodality in AR is the norm, not the exception.

The second reason is that virtual content can take non-visual and multimodal forms. In our opinion, considering these types of contents has no downside. Rather, working with other modalities opens up great opportunities: with sound, we can, for instance, also experience virtual objects that are hidden from our view. (We could, for instance, create the creepy feeling of being followed by presenting the participant with the sound of footsteps that seem to originate behind them and follow them around.) Furthermore, non-visual content can allow us to display invisible objects. This might make sense, if we, for instance, want to communicate the existence of ghosts in the environment (cf. Chatzidimitris et al., 2016).

The third reason for a multimodal perspective is that even visual additions that are integrated into a participant's view, can affect how the participant experiences non-visual aspects of the environment. Sight does not operate independently from our other senses; hence we should not treat it independently from it.

A final argument to treat AR from a multimodal perspective can be made for cases where AR technologies aim at making it appear as if virtual content existed in the real world or where it aims to imitate the real world.[23] In our opinion, multimodal and interactive properties of the real and the virtual world can play an important role in achieving these goals.

[23] As we have pointed out in section 3.2, we do not believe all AR projects have to create such an illusion. Because this argument only applies to cases where this illusion is desired, we present this argument as an additional reason to treat AR from a multimodal perspective rather than as a fundamental, general argument.

First of all, if virtual objects imitate real objects, we might expect them to display the same multimodal qualities that a real object would display. For instance, we would expect a virtual balloon to make a sound if it pops, and a virtual vase to make a sound if it breaks, and we might expect to feel something when we touch a virtual toy. Of course, our expectations are likely related to the specific object in question. However, a lack of multimodal qualities might harm the credibility of the virtual object and hurt the impression that the object is part of real space.

Furthermore, we believe that taking into account the multimodal qualities of the real world might make a virtual object's existence in the real world more believable. Imagine, for instance, a virtual pet that gets scared when there is a sudden sound in the surroundings, a virtual object that moves to the song playing on the radio, or a virtual character that puts on different clothes, according to the current outside temperature. We assume that if virtual elements react to the multimodal properties of the real world, such as its temperature and sounds, it might help convince us that virtual objects are actually in the same space as we are, and make the experience more entertaining and interesting.

At the same time, we hypothesize that the apparent presence of virtual objects in the real world may also be compromised when virtual content remains oblivious to the multimodal and interactive properties of its surroundings. For instance, we expect a virtual mouse to be

frightened (or at least react) when it hears a miaow.[24] Likewise, we would expect a virtual tree to get wet when it rains and expect virtual leaves to move if we feel that the wind is blowing.[25]

What is more, we might also expect that virtual content affects the real world and evokes multimodal responses: for instance, we would expect to hear sounds if a virtual ball bounces on a real wooden floor and expect a real window to break if the ball hits it.[26]

If, however, a virtual leaf does not move in real wind, a virtual mouse shows no reaction to a cat's miaowing, a virtual drum does not produce sounds when it is hit and if the floor remains silent when a virtual ball bounces on it, it might harm the impression that virtual content exists in the same space, even if this content is perfectly registered with the world visually.

To summarize, for objects to seemingly exist in the same space as real objects, it makes sense that they can interact with each other physically. Such interactions have visual and non-visual effects.

Ultimately, our various observations, assumptions and arguments boil down to one fundamental point: An AR environment is more than something a participant can see. In order to understand and advance AR, it does not suffice to study visual overlays that are integrated into a participant's view. Instead, we also have to study AR in terms of multimodal and interactive environments that a participant experiences.

[24] If we notice that the object is virtual, we also might expect the opposite and assume that the object cannot sense the world around it and interact with it. Here, we proceed under the assumption that the virtual object appears so present and real that it evokes the same kind of expectations like a real object. However, we are very much aware that an object's virtual nature might cause us to have different expectations instead.

[25] One might consider the question whether a leaf moves in the wind a minor detail. However, we mention such details because we expect they might considerably affect the AR experience, either consciously and unconsciously.

[26] Again, our expectations about this might be different if we can see that we are dealing with virtual objects. The topic of what we expect when we are aware that we are confronted with virtual objects is an interesting topic for future research.

3.4 *Synthesis, Discussion and Conclusion*

AR is commonly considered an interactive technology that overlays virtual imagery onto our view of the real world and that aligns this virtual content with a user's view of the real world in 3D and in real-time. In this chapter, we have argued that this image of AR is incomplete and we have proposed an alternative, more encompassing view of AR. This view departs from widespread understandings of AR in three complementary ways.

First, we do not view AR as a technology. Instead, we claim that AR technology enables augmented reality. We focus on the resulting augmented reality environments and experiences rather than on the technologies that enable them.

Second, we see AR as a result of the relationships between virtual content and the so-called real world. Whereas AR is generally assumed to involve the spatial integration of virtual content in (our view of) the real world, we believe that other types of relationships between the virtual and the real are possible. We hypothesize that different and new forms of relationships will enable different and novel forms of AR.

Third, we treat AR as a multimodal and interactive phenomenon and argue that AR engages all our senses and allows for action in and

interaction with the environment. Instead of focusing on what a user or participant sees, we focus on non-visual and multimodal aspects of both the real world and virtual content.

We can synthesize these views and propose to define AR in terms of interactive and multimodal environments where a participant experiences a relationship between virtual content and the real environment.

This definition allows us to distinguish AR from other environments. Most importantly, it allows us to distinguish AR from solely physical environments that do not contain any virtual information. Furthermore, the definition sets AR apart from scenarios where the virtual and the real merely coexist in the same space and where both are experienced as *independent* from each other. For instance, it does not include situations where a participant listens to an audiobook and experiences this story as unrelated to their actual environment. Likewise, it sets AR apart from entirely virtual environments. For instance, our definition does not include scenarios where participants are immersed in virtual worlds and where they experience virtual elements as independent from their actual, real environment. Finally, due to its focus on the *real* environment, our definition also sets AR apart from other mixed reality environments where a participant experiences a link between the virtual and the real. For instance, it does not include situations where the participant experiences real objects (e.g. a physical toy gun) in relation to an otherwise virtual environment (we will propose to refer to such environments as "augmented virtuality" below.)[27]

[27] Of course, the distinction between AR and other mixed reality environments is not clear-cut.

Although our perspective on AR deviates from prevailing ideas, our understanding of AR is not entirely new (see chapter 2). In particular, many researchers suggest that AR can engage all senses. To mention just a few examples: Azuma et al. (2001) point out that "AR can potentially apply to all senses, including hearing, touch, and smell" (p. 34). Furthermore, Craig (2013) points out "Augmented reality can appeal to many of our senses (although currently it is primarily a visual medium)" (p. 1-2). Lindeman and Noma (2007) explicitly explore the idea of multi-sensory AR and present a classification scheme that allows for visual, auditory, haptic, olfactory and gustatory forms of AR. However, these views generally assume that AR only will engage or address other senses, if non-visual virtual content is presented to the participant. Our view sets itself apart from such ideas because it considers AR as multimodal, even when virtual content using only one modality (e.g., only visual content) is presented to the participant. This is because we consider the multimodal real world as part of the experience. As a consequence, multimodality in AR is the norm rather than the exception.

We are not the first ones to focus on the relation between the virtual and the real rather than on registration. For instance, Manovich (2006) suggests that "a typical AR system adds information that is directly

related to the user's immediate physical space" (p. 225). Similarly, Klopfer and Squire (2008) define AR in terms of "situation[s] in which a real world context is dynamically overlaid with coherent location or context sensitive virtual information" (p. 205). While these claims suggest that the virtual has to relate to the real, they do not claim that the information has to be *registered* in real 3D space or aligned with real objects in 3D. The main contribution of our work in this context is that it provides a detailed rationale for deviating from commonly accepted focus on registration. Furthermore, our proposed definition differs from views such as put forward by Manovich (2006) and Klopfer and Squire (2008) with its focus on the participant's *experience* of those relationships.

While scarce, some existing definitions also focus on the participant's experiences. For instance, Spence and Youssef (2015) describe AR as "an experience of a physical, real-world environment whose elements have been augmented, or supplemented, by computer-generated sensory input" (p. 1). However, if we look beyond mere definitions, there are more views that emphasize the experiential qualities of AR. For instance, Craig (2013) focuses on the experience associated with AR and writes "the way people engage with augmented reality is to experience it" (p. 1). However, our view still differs from such existing views with respect to what constitutes an AR experience: We believe AR is characterized by the experienced relationship between the virtual and the real. Furthermore, we are not aware of any existing views that explicitly argue for a shift from a technology-focused definition of AR towards an experience-focused definition of AR.[28] While the individual points are not necessarily new in isolation, our contribution is unique and new in this combination.

In our definition, the participant's experience of the *real environment* plays a key role. Existing definitions of AR generally focus on the fact that a real environment is part of AR (e.g., Azuma et al. (2001)). Most notably, Milgram and Kishino (1994) describe AR as "all cases in which the display of an otherwise real environment is augmented by means of virtual (computer graphic) objects" (p. 1321) and places AR on the side of the real environment in their much-cited reality-virtuality continuum (see section 2.1). Our definition is similar to existing views because it also focuses on the real environment. However, unlike many views, our definition does not focus on the environments that are displayed by a system but on the environments that are perceived by the participant.

In our opinion, the proposed perspective on AR advances our understanding of AR on a fundamental and theoretical level. We hope that a better and broader understanding of AR will inspire new forms of AR. Our investigation has revealed various examples of interactive applications that defy prevailing definitions of AR but yet, augment our experience of our physical surroundings. This shows that narrow

[28] However, a similar shift has been postulated in the context of Virtual Reality by Steuer (1992). In his seminal paper, Steuer criticizes that the focus of virtual reality is technological, rather than experiential. Consequently, he focuses on the participant's sense of being in an environment and defines VR as "a real or simulated environment in which a perceiver experiences telepresence" (p. 7). While we propose a similar shift in perspective, our proposal differs fundamentally from Steuer's contribution because it addresses AR rather than VR and thus, focuses on a different kind of experience.

definitions not necessarily prevent practitioners to think outside of the box and to come up with different forms of (arguably) augmented reality. Yet, we expect that a better understanding of AR and of how to create (or facilitate) it, will inspire even more and new forms of AR.

When it comes to creating AR scenarios, our view of AR suggests that we have to consider and give form to the relationships between the virtual and the real. However, we also have to keep in mind that establishing a relationship between the virtual and the real not automatically ensures that a participant also experiences this relationship. What is more, a participant might experience relationships that have never been created or intended. For instance, a museum visitor might listen to a virtual museum guide, but associate the information with the wrong artwork. Similarly, the same scenario might be experienced as AR by one person but not by another. In our opinion, the question whether a scenario should be considered AR cannot be answered based on what a system does or displays. Instead, it remains a question of personal experience.

Our investigation has revealed two main forms of AR: First, cases where a participant experiences the presence of virtual content in the real environment. We propose calling this "presence-based AR". Second, cases where the participant experiences virtual content as related to or pertinent to their surroundings on the content-level. We suggest calling this "content-based AR". In future research, it would be desirable to explore if yet different forms of AR exist. For instance, can the virtual become part of the real world similarly to how a soundtrack becomes (a non-spatial) part of a movie? Furthermore, we would like to systematically explore what factors contribute to the experience of virtual content being part of the real space. We can imagine that next to registration, aspects such as the participants' imagination and an underlying narrative can play a major role in AR.

In line with our definition, our investigation has focused on situations where the participant experiences virtual content in relation to the so-called real world. However, all of our three main considerations—the focus on (1) experience, (2) multimodality and (3) relationships between the virtual and the real – can likewise be applied to the more general field of mixed reality (cf. section 2.1). If we generalize our definition, mixed reality can be seen as any environment in which the participant experiences a relationship between the virtual and the real. In line with this, *augmentation* can be seen as the result of the perceived relationships between the virtual and the real. Those specific mixed reality environments where the participant experiences real elements in relation to their otherwise *virtual* surroundings can be described as *augmented virtuality*.

In the future, it would be interesting to further explore our arguments in the more general context of mixed reality. However, it is also necessary to further investigate what our proposed view of AR entails,

and what forms AR can take if we apply our definition. This is the focus of this thesis. As a result, issues that play a role in other forms of mixed reality (i.e., augmented virtuality) fall out of the scope of this thesis. For instance, we will not discuss the experience of telepresence in virtual environments and the representation of real participants in virtual space in the form of an avatar.

In conclusion, we have presented new perspectives on AR, and arrived at an understanding that focuses on the participant's experience of the environment, the relationships between the virtual and the real and the multimodal and interactive qualities of the environment. More specifically, we have proposed that AR is characterized by the experience of virtual content in relation to an otherwise real, multimodal, interactive environment. We have already encountered several examples that fit and illustrate this broader definition, such as audio guides in a museum. However, many questions remain open: What else does augmented reality entail if we apply our definition? How can AR look, taste, smell, feel and sound like if we do not require registration? Are there yet other forms of AR, based on yet different relationships between the virtual and the real? In which ways can the virtual become part of and relate to the real environment? In the following chapters, we address these questions. We apply our proposed perspective and systematically explore the various forms AR can take.

Part II

What Forms Can Augmented Reality Take?

4 Relationships Between the Virtual and the Real

In the previous chapter, we have proposed that AR is characterized by the relationships between the virtual and the real. More specifically, we have argued that in order to experience AR, a participant has to experience a relationship between the virtual and the real. Simply put, we believe that the virtual and the real *augment* each other if the participant experiences a link between them. In line with this, we see augmentation as the result of the experienced relationships between the virtual and the real. This proposed view of augmentation does not necessitate a system that aligns virtual content with the real world interactively and in real-time and allows for new and different manifestations of AR. For instance, it encompasses scenarios where virtual content informs us about our real surroundings. In this chapter, we will build on this view of AR, explore possible relationships between the virtual and the real and investigate what AR is and can be if we approach AR from our proposed perspective.

As mentioned, the idea that relationships between the virtual and the real are pivotal for AR (and more generally, Mixed Reality) is not new. For instance, new media theorist Manovich (2006) notes: "In contrast [to a typical VR system], a typical AR system adds information that is directly related to the user's immediate physical space" (p. 225). According to MacIntyre (2002), the more general field of Mixed Reality (see section 2.1) is characterized by these relationships. He states that "[t]he relationships between the physical and virtual worlds is what makes Mixed Reality applications different from other interactive 3D applications" (p. 1). Looser, Grasset, Seichter, and Billinghurst (2006) refer to MacIntyre with their claim that "[c]reating content for Mixed Reality (MR) and specifically Augmented Reality (AR) applications requires the definition of the relationship between real world and virtual world" (p. 22). Hampshire, Seichter, Grasset, and Billinghurst (2006) make a similar reference to MacIntyre and state that "[d]esigning content for MR is driven by the need to define and fuse the relationship between entities in physical world and virtual world" (p. 409).

As these quotes show, the importance of relationships between the

virtual and the real for AR is also acknowledged by other researchers. However, existing AR research commonly reduces this topic to the *registration* of virtual objects with the real world in three dimensions and focuses on processes that make it look as if virtual objects existed in real space. For instance, existing research is very concerned with the tracking of the participant and the creation of correct occlusions between virtual and real objects (cf. Zhou et al., 2008).

In contrast, we believe that there is much more to AR than the apparent presence of virtual objects in real space. We expect that augmentation has many more facets and that relationships between the virtual and the real can be established on various different levels. For instance, a virtual museum guide might appear spatially present in the exhibition space and also inform us about our surroundings on the content-level. Likewise, a virtual bird might appear to sit on top a real tree branch and relate to its surroundings spatially, while at the same time also imitating the songs of real birds in the forest on a musical level. We believe that in such cases, the different relationships between the virtual and the real all contribute to and shape the resulting AR experience. What is more, we do not think virtual content needs to appear as if it existed in real space in order to augment this space—a relationship between the virtual and the real is enough.

The realization that AR is characterized by relationships between the virtual and the real rises several questions that have received little attention so far: What relationships between the virtual and the real are possible? How can the virtual relate to, and ultimately augment, the real world? What forms can augmentation take? What strategies are at our disposal to establish a relationship between the virtual and the real? And finally, what does AR entail, if we define AR in terms of relationships between the virtual and the real?

In this chapter, we address these questions. We apply our newfound definition of AR, explore different facets and forms of augmentation and identify various ways in which the virtual can relate to the real. Ultimately, our review reveals that there is much more to AR than the apparent presence of virtual objects in real space. For instance, we will see that virtual content can seemingly remove elements from the real world, transform the real world, or allow us to perceive aspects of our surroundings that typically are unperceivable to our senses.

In our investigation, we primarily focus on how virtual content relates to and affects the real environment in which it is presented.[1] We do this because typically, virtual content is added to our real existing environment as opposed to the other way around. By focusing on how the virtual relates to the real we do not mean to imply that the relationship is one-directional. In fact, we believe that typically, the virtual and the real relate to one another and augment one another.

The question of how the virtual relates to the real world serves as a basis for the structure of this chapter. The subsequent sections each

[1] Exceptions are section 4.1, where the virtual and the real exist independently, as well as section section 4.9 and section 4.10, which explicitly focus on how the virtual and the real interact with *each other*.

discuss one common relationship between the virtual and the real. In the following three sections, we discuss the *fundamental relationships* as well as the absence of a relationship between the virtual and the real:

- **(4.1) Coexistence: Independence of the Virtual and the Real**. Virtual content is presented in the real environment but seems to exist *independently* from it. The participant does not experience a relationship between the virtual and the real. According to our view of AR, coexistence is thus not enough to constitute AR.

- **(4.2) Presence: Spatial Relationships**. This section refers to spatial relationships between the virtual and the real. More specifically, it describes scenarios where virtual content seemingly exists *in* real space and at a certain position in the real environment, rather than, e.g., on a screen or in a separate virtual world.

- **(4.3) Information: Content-Based Relationships**. The virtual relates to the real content-wise. This is, e.g., the case when virtual content informs us about the real environment or when it tells a story about the real surroundings.

The subsequent five sections discuss relationships between the virtual and the real that potentially emerge from and build on these fundamental relationships. The question that we address on this second level is how the presence/presentation of virtual content affects its real surroundings. Based on the role of the virtual content in the real space, we distinguish between the following sub-forms of AR:

- **(4.4) Extended Reality: The Virtual Supplements the Real.** Here, virtual content acts as something additional that supplements the real world. As a consequence, the environments appear to contain *more* content.

- **(4.5) Diminished Reality: The Virtual Removes the Real.** In this case, there seems to exist *less* content in the surroundings.

- **(4.6) Altered Reality: The Virtual Transforms the Real.** In this instance, the virtual changes the apparent qualities of the real world. For instance, the virtual might alter the perceived size or shape, weight or texture of real objects. Here, the participant not necessarily perceives more or less information, but instead, perceives *different* information.

- **(4.7) Hybrid Reality: The Virtual Completes the Real.** Here, the virtual does not serve as "something extra" and *optional* in the otherwise real environment but rather *completes* a physical environment (or object) that would be considered incomplete without the virtual additions.

- **(4.8) Extended Perception: Translation-Based Relationships**. The virtual translates unperceivable aspects of the real world, such as radiation or ultrasound to information that we can perceive with our senses (e.g., sounds in our hearing spectrum, images or tactile stimuli). In other words, the virtual allows us to perceive real aspects of the environment in the context of this environment. We refer to this as extended perception.

The next two sections once more focus on scenarios where virtual objects seemingly exist and extend the real world. We notice that the presence of virtual objects in real space opens up possibilities for influences and interaction between the virtual and the real. The sections take our investigation one step further in the sense that we not only look at how the virtual content affects the real world but also at how the real world can affect the virtual in return. Furthermore, we emphasize the fact that virtual elements not only can appear to *exist* in the world but also can seem to *act* and *behave* in the real world. On this level, we distinguish among two main forms of relationships between the virtual and the real:

- **(4.9) Physical Relationships: The Virtual and the Real Affect Each Other.** This section discusses physical effects between the virtual and the real. Among other things, we discuss optical interactions, such as virtual and real objects casting shadows on each other and dynamic interactions, such as virtual objects being affected by the gravity and collisions between virtual and real objects.

- **(4.10) Behavioral Relationships: The Virtual and the Real Sense and React to Each Other.** In this section, we discuss influences and interactions between the virtual and the real that take place on a behavioral level. An example of such influences would be a virtual creature that is scared away by certain sounds in the real environment.

We conclude the chapter with two more sections. In these sections, we look beyond the previously discussed relationships as well as reflect on our findings in a broader context.

- **(4.11) More Relationships.** In this section, we emphasize that the collection of discussed relationships is not exhaustive. We briefly discuss other possibilities, such as temporal relationships between the virtual and the real and musical relationships between virtual and real instruments.

- **(4.12). Summary, General Discussion and Conclusion.** In this section, we summarize and reflect on our findings and discuss them on a general level and in the context of existing AR research.

Each section is heavily based on examples. In contrast to the previous chapter, the role of these examples is less argumentative and

more illustrative. This means that the examples showcase the different identified relationships. Together, the various examples also provide insights into the diversity of AR, which is an overall goal of this chapter and this thesis. For instance, we will see that AR projects have many different goals, make use of various different stimuli and technologies, are used in different application contexts and ultimately, can evoke a variety of experiences. Yet, the examples provided in each section also have an argumentative role: they prove that the identified relationship between the virtual and the real indeed exists and demonstrate its relevance in the field of AR. In this sense, they also support our choice to dedicate a category to the identified relationship.

In their totality, the various identified relationships between the virtual and the real form a topology. However, unlike in classical typologies, the identified types of relationships can surface in *combinations*. For instance, a virtual museum guide might visually appear as if they existed in the real environment *and* inform us about our real surroundings. Furthermore, some types of relationships can be considered subgroups of other types of relationships. An example is extended perception, where virtual stimuli are used to make unperceivable aspects of reality perceivable, and where this information naturally also informs us about the real world. Moreover, some relationships enable or build on other relationships. For instance, the presence of a virtual object in real space enables possibilities for physical interaction between the virtual object and its real surroundings. In order to emphasize that the different types are not exclusive, we will refer to the same examples in different sections. Furthermore, it is important that other types of relationships aside from the discussed ones are possible. As the identified types of relationships are neither *jointly exhaustive* nor *mutually exclusive*, we are not dealing with a classical typology. Rather, we present a hybrid, incomplete typology, as described by Bellamy and 6 (2012).

As the above overview shows, this chapter is rather comprehensive. It uses our definition of AR as a starting point and consequently, explores it by moving into many different directions. This results in a long and diverse chapter. The red line that holds the parts together is the notion that in AR, the virtual relates to the real. It is possible for the reader to follow this line in some directions while skipping others. In other words, the sections largely can be read and understood on their own. However, only together they provide an overview of the AR landscape and illustrate the diversity of what AR is and potentially can be. To the best of our knowledge, a comparably comprehensive overview of the different manifestations of AR has not yet been presented in AR research.

Throughout this chapter, we focus on relationships between virtual content and real content that appear in the *same* physical space. Relationships between virtual and real content that are not part of the same environment fall out of the scope of our investigation. (For instance,

we will not discuss the relationship between a virtual letter and the remote author of that virtual letter.) This is because, according to our definition developed in the preceding chapter, AR is concerned with the relationships that a participant experiences between something virtual and their real surroundings.

One aspect that we have to consider is that the participant typically also is a real part of this environment. In the AR research field, relationships between virtual content and a participant play an important role. As we know from the previous chapter, many interactive AR systems react to the participant's movement and display virtual content in a way that it matches the participant's perspective. Furthermore, several AR projects allow a participant to interact with the virtual content, and e.g., move virtual content (e.g., Billinghurst, Kato, and Poupyrev, 2008; Irawati et al., 2006). It should be emphasized that relationships between virtual content and the *participant* are not the primary focus of this chapter. Yet, we will consider relationships between the virtual and the participant in those cases where they play a prominent role. For instance, we discuss that virtual information can *inform a participant* about their surroundings.

Like in the previous chapter, we focus on conceptual and experiential aspects of AR and do not discuss technological issues. Whereas the previous chapter has focused on visually and sonically augmented reality (the two most common forms of AR), this chapter also considers other modalities. Consequently, many examples not only illustrate interesting relationships between the virtual and the real but at the same time reinforce our thesis-wide claim that AR is more than what meets the eye. Furthermore, while the previous chapter has focused on (a) virtual content that appears to exist in real space as well as on (b) virtual content that informs us about the real world, this chapter explores many more ways in which the virtual can relate to and augment the real world.

In order to distinguish between (1) the common understanding of AR in terms of systems that align virtual images and the real world in three dimensions interactively and in real-time, and (2) our newly proposed, broader understanding of AR in terms of relationships between the virtual and the real, we will refer to the former as "traditional AR" or as "registration-based AR" and to the latter as "AR in the broader sense" or "relationship-based AR".

4.1 Coexistence: Independence of the Virtual and the Real

In our everyday reality, virtual content is omnipresent: on advertisement screens, on the displays of mobile phones, tablets, smart watches, digital information boards, game consoles, radios, laptops and suchlike. Often, the information that reaches us through these channels has rather little to do with its physical surroundings. For instance,

the websites we skim while on the train do not concern the things we see when we look up or gaze out of the window. Likewise, the mails we read while waiting for our flight commonly have nothing to do with the airport we are at. Furthermore, computer games often take place in a virtual space that is independent from a player's real environment. Regularly, such games go as far as to separate us from the real world and temporarily take its place. In particular, Virtual Reality (VR) technologies aim at immersing participants in alternative, virtual spaces that typically have nothing to do with the player's immediate real surroundings (cf. Manovich, 2006).

As these examples illustrate, the fact that we engage with virtual content in our otherwise real, physical environment does not necessarily mean we experience a meaningful relationship between the two. Often, the virtual disregards its real surroundings and is experienced as an independent layer of information. In such cases, the virtual content and the real environment coexist, as opposed to relate to one another—they seem to exist in parallel, rather than integrate with each other.[2] Yet, one might argue that a relationship between such virtual and real elements exists. After all, virtual content is displayed or presented in the real environment. We refer to this basic and underlying link between virtual content and the world as *coexistence*.

[2] Some might object to the idea that the virtual *exists*. In this thesis, we treat the virtuality as a certain (simulated) form of existence. In our view, objects can exist both physically as well as virtually.

In our opinion, the mere coexistence of virtual and real content in the same environment is not enough to constitute AR. Instead, the virtual also has to *augment* the environment. In existing AR research, this augmentation is typically seen as a form of *supplementation* or *enhancement* of the real world by means of virtual content. For instance, Yuen et al. (2011) write "Augmented Reality (AR) is an emerging form of experience in which the Real World (RW) is enhanced by computer-generated content tied to specific locations and/or activities. " (p. 119). Similarly, Bederson (1995) states that "Augmented reality [...] uses computers to enhance the richness of the real world" (p. 210).

The fact that the virtual content is added to the real world is often seen as a factor that distinguishes AR from VR. For instance, Azuma (1997) compares AR to VR, and points out that in contrast to VR, "AR supplements reality, rather than completely replacing it" (p. 356). Likewise, Höllerer and Feiner (2004) point out that in contrast to virtual reality, AR "aims to supplement the real world, rather than creating an entirely artificial environment." (p. 221-222).

As we will see in the following sections, augmentation indeed often takes the form of virtual content that supplements and extends the real world. However, in addition, augmentation can also take other forms, and, for instance, transform or diminish the real world.

In our opinion, the fact that the virtual plays a role in the real world not only distinguishes AR from VR, but also distinguishes AR from environments where we experience virtual content as *independent* from the real world, rather than as *related* to or part of this world. We believe

the virtual *augments* the real environment if it is perceived as related to our real surroundings. In the following sections, we will explore the many ways in which the virtual can relate to, and ultimately add to, supplement, or *augment* the real world.

4.2 *Presence: Spatial Relationships*

AR involves the presentation of virtual content in real space. However, as we have shown in chapter 3, traditional registration-based AR applications go one step further than simply displaying or presenting virtual content. They also align virtual content with the real world in three dimensions and make it appear as if the virtual content existed in the physical environment, rather than on a display or in a separate second space. In such cases, the virtual is not only *presented* in a real environment but also appears to be *present* in this space. As mentioned in section 3.4, we propose to call this form of AR *presence-based AR*.

The presence of virtual content into the physical environment goes hand in hand with different *spatial relationships* between the virtual and the real. First of all, virtual content appears to exist *in* the real world and seemingly occupies real three-dimensional space. In addition, virtual content spatially relates to real objects in this space. For instance, a virtual object might appear to exist in front of, on top or next to real objects. (Technically speaking, they appear to share one coordinate system.)

The virtual content that appears to exist in the real environment can play various roles in this environment and take many forms. Most commonly, the virtual takes the form of virtual objects that appear to exist in real 3D space, alongside real objects. This is, for instance, the case in the first so-called augmented reality prototype by Caudell and Mizell (1992) (see figure 1.1). As discussed, their prototype was aimed at displaying virtual instructions about manufacturing processes in a way that they appeared in 3D space. In their paper, the authors sketch an example where a virtual arrow points at an exact location on a physical airplane fuselage to indicate the spot where a hole has to be drilled. In section 4.4, we will discuss such environments that appear to contain additional virtual elements or supplementary content in more detail. We propose to call this subform of AR *extended reality*.

In addition to supplementing the real world, virtual content can also *complete* the real environment. The difference is the following: When the virtual extends the real world, the real surroundings can still be considered "complete" without virtual additions. For instance, virtual ghosts play a crucial role in the AR game by Chatzidimitris et al. (2016). However, the virtual ghosts are not essential to the real streets—the real environment is also complete without them. In contrast, when the virtual completes the real, the environment is incomplete without the virtual component. The virtual is integral to the real environment

and thus completes rather than supplements the real. This happens in cases where the design of an environment or an object includes both a physical and a virtual component. In such cases, the real, physical component needs the virtual component. In other words, the virtual does not add "something extra" but completes the real.

An AR project where the real component deliberately leaves out certain characteristics to be filled in by the virtual is the augmented zebrafish by Gómez-Maureira, Teunisse, Schraffenberger, and Verbeek (2014) (see figure 4.1). With respect to the real component, this project consists of a physical, bigger-than-life zebrafish. On itself, this physical model appears rather incomplete: it is completely white; visual features of its skin such as colors and texture are missing. However, the zebrafish's skin is deliberately added virtually and projected onto the fish, which opens up possibilities that a solely physical model does not offer: The virtual projections not only add visual features but also allow the audience to interact with the object. If audience members step in front of the projector and move their shadow over the fish's surface, the shadow is filled by a second projector with additional information. For instance, their shadow will reveal an X-ray visualization and a basic anatomical schematic. In other words, the audience can look inside the fish and explore its anatomy by casting shadows on it. In section 4.5, we propose to call this form of AR *hybrid reality* and provide a more detailed discussion of cases where the presence of virtual content in real space completes rather than extends the real.

Furthermore, the spatial integration of virtual content in real space can be used to hide or seemingly remove or replace real elements from the real world. In this case, the participant experiences less rather than more content in their surroundings. This paradigm is also often referred to as "diminished reality" (e.g., Herling and Broll, 2010). The concept of diminished reality has, for instance, been explored by Mann and Fung (2002). The authors believe that diminished reality can be used to help avoid information overload. They introduce a system and algorithm that (among other things) is able to remove what they call *Real-world "spam"*, such as undesired advertisements from a user's visual perception of their surroundings (see figure 4.2). (The undesired 'spam' is replaced by different content). In line with existing research, we propose to call this form of AR that uses virtual additions to seemingly remove and replace real elements *diminished reality*. It is discussed in section 4.6.

In addition to adding and removing elements to and from the real world, virtual content that appears to exist in the world also can *transform* the real environment or real objects. For instance, Bandyopadhyay et al. (2001) have proposed a projection-based system that allows a user to transform real, physical (neutrally colored) 3D objects by virtually painting on them and by applying different virtual textures that can seemingly change their material properties (see figure 4.3). The

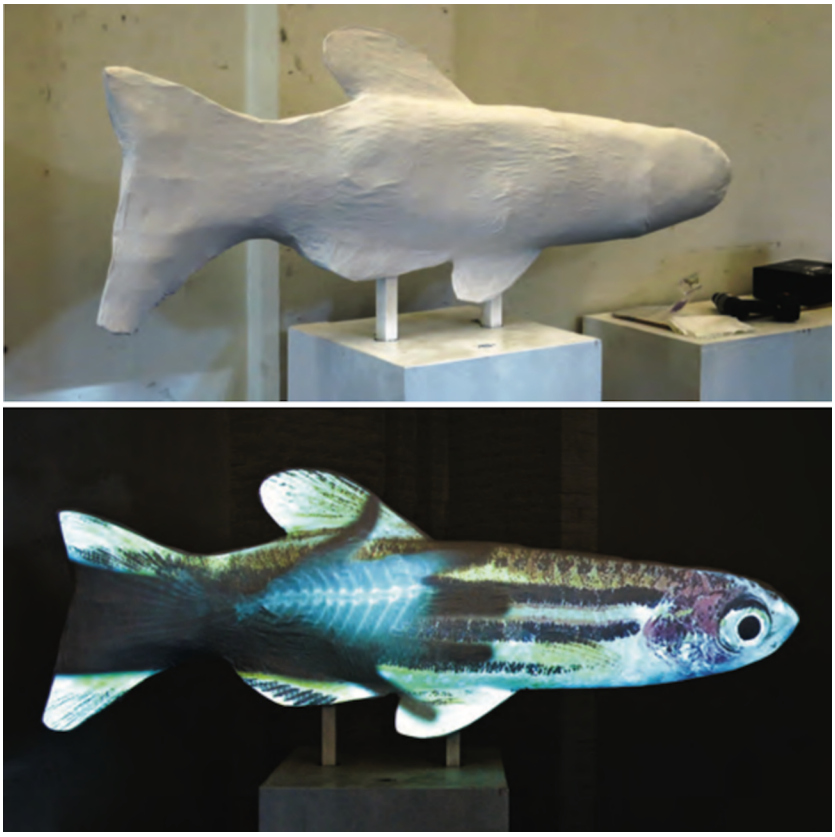

Figure 4.1: Virtual information completes a physical model of a zebrafish. Without the virtual component, the object is incomplete. Reprinted from M. A. Gómez-Maureira et al. (2014). "Illuminating Shadows: Introducing Shadow Interaction in Spatial Augmented Reality". In: *Creating the Difference: Proceedings of the Chi Sparks 2014 Conference*, pp. 11–18. Reprinted under fair use.

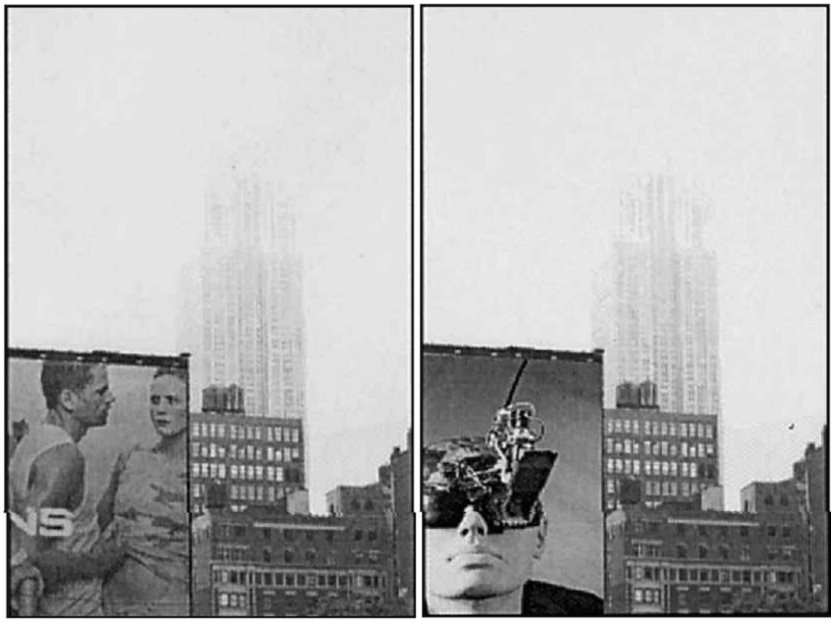

Figure 4.2: Virtual information removes advertisements on a billboard from the environment and replaces it with alternative content. Reprinted from S. Mann and J. Fung (2002). "EyeTap devices for augmented, deliberately diminished, or otherwise altered visual perception of rigid planar patches of real-world scenes". *Presence: Teleoperators and Virtual Environments*, 11(2), pp. 158–175. Reprinted under fair use.

concept of transforming the real world is popular in the context of projection mapping, where buildings can seemingly be transformed

by means of projections (see figure 4.3). We coin this form of AR *altered reality* and discuss the transformation of the real world by means of virtual content in depth in subsection 4.7.1.

Moreover, virtual objects in our real surroundings can be used to represent real but unperceivable aspects of the real world. For instance, virtual arrows could be shown to visualize the magnetic field, and virtual dust could be displayed to allow us to perceive air pollution. We call these instances of AR *extended perception* because they allow us to perceive more about the world. Extended perception will be discussed in section 4.8.

Figure 4.3: Virtual information can transform the real world. Here, artist Valbuena (2008) alters the appearance of the *The Hague City hall* with his dynamic installation *N 520437 E 041900 [the hague city hall*. Images from `http://www.pablovalbuena.com/selectedwork/n-520437-e-041900`. Reprinted under fair use.

The presence of virtual content cannot only extend, complement, transform or remove the real—it also opens up possibilities for (simulated) physical relationships between the two. For instance, if virtual objects appear in the real environment, they can seemingly be affected by gravity and appear to collide with real objects (Breen et al., 1996). Likewise, optical influences are possible. E.g., virtual objects can cast shadows on real objects and real objects can cast shadows on virtual objects (Madsen et al., 2006).[3] Physical influences and interactions will be discussed in section 4.9.

[3] We see these as physical influences because we choose to consider light as a particle as opposed to a wave. In line with this, we treat light-related influences as physical influences.

In addition, the presence of virtual objects in real space also opens up possibilities for *behavioral* relationships between the virtual and the real. For instance, in the AR version of the game Quake (Piekarski and Thomas, 2002), virtual monsters interact with the player on a behavioral level in the sense that they attack the player and that the player tries to shoot them. We will discuss behavioral relationships in more depth in section 4.10.

Although the examples above all deal with visual virtual content, it is important to note that the spatial presence of virtual content in real space is not limited to what we see. Rather, when the virtual is integrated into the real surroundings spatially, it becomes part of an environment we perceive with all our senses (see subsection 3.3.1). Furthermore, virtual content that spatially relates to the real world can take non-visual and multimodal forms. For instance, the Sound-Pacman game places virtual ghosts in the real environment by means

of spatialized sound (Chatzidimitris et al., 2016). Similarly, the Gravity Grabber by Minamizawa, Fukamachi, et al. (2007) allows us to feel virtual objects bouncing inside a real cube. Even smells, which typically are not perceived at an exact location in the surrounding space, might convey the presence of certain virtual elements in the environment. For instance, the mere smell of coffee might be used to create the illusion of real coffee being present in the environment. In the following sections, we pay close attention to the possibilities of augmenting the real world by means of non-visual content. We will discuss the above-mentioned examples in more depth as well as include a broad variety of other projects that illustrate the various possibilities of creating relationships between the real world and non-visual virtual content.

To summarize this section, virtual content can relate to the real world *spatially* in the sense that it appears to exist in this real space. We call this form of AR *presence-based AR*. In presence-based AR, virtual content appears *present* in the otherwise physical surroundings (rather than, e.g., on a screen or in a separate virtual world). The presence of virtual content in a real environment can affect the real world in many different ways. E.g., it can extend the real world as well as remove or transform real objects. The presence of virtual content in real space furthermore opens up possibilities for physical and behavioral influences and interactions between virtual and real content. The presence of virtual content in real space is often simulated visually, however it can also take non-visual and multimodal forms.

4.3 *Information: Content-Based Relationships*

As we have shown, the virtual can relate to the real by appearing spatially present in the real environment. Furthermore, the virtual can relate to the real on the *content-level* (see subsection 3.2.1). For instance, a virtual museum guide might inform us about a painting. In such cases, there is an intrinsic link between the additional virtual information and a participant's physical environment. In addition, the virtual content also relates to the participant in the sense that it informs them or tells them something about their surroundings. As mentioned, we believe that AR in the broader sense includes such scenarios where the virtual relates to its real surroundings content-wise. We have discussed this concept in subsection 3.2.1 and coined it *content-based AR*. In the following, we will revisit this topic and illustrate the prominent role that virtual information plays in the real world as well as in our everyday lives.

Virtual content that informs participants about their real surroundings is rather common in the western everyday world. Think, for instance, about digital information displays that tell us about the departure times of trains, about GPS devices, which help us navigate the space and about audio tour guides that inform us about exhibitions,

monuments or other points of interest.

The idea of informing participants about their immediate surrounding environment by means of virtual content is also often used in the context of traditional AR. An early example of an AR application that provides such information is the so-called "touring machine" prototype by Feiner, MacIntyre, et al. (1997). This system allows users to freely navigate a university campus. The users would receive information about the campus, both on a head-worn see-through display, as well as on a handheld opaque display. In their prototype, the head-worn display overlays the names of campus buildings over the participant's view of the actual buildings. In addition, the head-worn device shows different menu items. When selected, the handheld device will open documents that provide additional information about the university and the campus.

The mobile application *Layar* (2009), among other things, allows for similar experiences. The app can present site-specific content, such as information about nearby restaurants, metro stops and ATMs and other spatially related information, such as tweets that have been tweeted in the neighborhood.[4] This data is overlaid onto the real world using a mobile phone's screen and often includes images or icons that seem to float in the real 3D space, in front of the phone's lens. Aside from such imagery, the app presents text, as well as visually indicates the directions of the points of interest. In contrast to the "touring machine" prototype, this app makes use of user-generated content (theoretically, everyone can publish their own channels with additional information) and presents all information on only one screen. Also, Layar works globally as opposed to at one predetermined location. For instance, a user can receive information about their surroundings, no matter whether they open the app in Stuttgart (Germany) or in Leiden (the Netherlands).

[4] In addition, Layar also focuses on other scenarios, such as the augmentation of print content.

Aside from Layar, we can find many other phone-based mobile applications that present users with information that relates to the location where it is presented on the *content-level*. In order to inform the participant, this information not necessarily has to appear on top of or integrated into the real world. For instance, *Street Museum NL* (2013) dynamically displays old photographs that have been taken in the surrounding area on the smartphone screen. These images inform the user about the past and how the surroundings used to look a long time ago, even if they do not appear to exist in real 3D space or float over their view.

A dedicated device, which is built around the idea of enhancing and supplementing our everyday lives by means of additional virtual information is the Google Glass headset. As we have shown in section 3.2.1, this head-mounted display, in the shape of eyeglasses, presents additional information, such as text, images or videos as an overlay that appears on top of a user's view of the world. As mentioned, this infor-

mation can be completely unrelated, but also relate to a user's current context or location and e.g., present us with driving instructions.

Often, virtual information not only informs us about the world but also *instructs* participants about how to *act* in the world. Common examples are visual and/or sound-based driving instructions. In addition, the concept of guiding a person's actions in the world is also at the heart of several previously mentioned traditional AR applications. For instance, Caudell and Mizell (1992), who coined the term AR, originally saw AR as a means to guide workers in the manufacturing process. In line with this, they describe AR as "a technology [which] is used to 'augment' the visual field of the user with information necessary in the performance of the current task" (p. 660). Their proposed prototype, among other, uses a red line and descriptive text to illustrate which wire goes into which pin in a connector assembly task. Another previously mentioned example of a traditional application that informs the user and guides their actions in the real world is the AR system by Feiner, Macintyre, et al. (1993). This head-mounted display explains users how they can maintain and repair an office printer by means of line-based illustrations that appear to exist in real 3D space and that explain certain goals and actions.

As we have shown, virtual information is commonly used to inform us about points of interests and objects, such as monuments. However, it can also be used to inform us about *people* in our environment. For instance, the Recognizr concept/prototype by The Astonishing Tribe (Jonietz, 2010) intends to inform us about people in our surroundings. The underlying idea is that the software recognizes people who have opted in to the service using a face recognition algorithm and consequently displays their names as well as links to their profiles on social platforms when their face is viewed with a smartphone running the application.

Although the Recognizr concept was presented as early as 2010, the Recognizr app has not been realized in the meantime.[5] However, a similar concept was realized by Gradman (2010) in an art context. In contrast to Recognizr, Cloud Mirror is a *static* installation that takes the form of a digital mirror. This digital mirror temporarily merges the online identities of visitor's with their physical selves (Gradman, 2010). The installation identifies visitors based on their badges and consequently searches the Internet (facebook, twitter, flickr) for photographs of and facts ("dirt") about them. When visitors approach the digital mirror, the found data is, e.g., superimposed in an on-screen comic book-like thought bubble that follows the visitor's motion (see figure 4.4). (The virtual content thus relates to the human both spatially and content-wise).

[5] Their public facebook page displays a lost post from 9 September 2014, informing readers about the fact that their Kickstarter campaign has not been successful, promising to keep readers in the loop with their progress.

In addition to applications where virtual content informs us about physical and tangible elements in our surroundings (such as objects or people), we can also find applications where the virtual informs

Figure 4.4: In this digital mirror, virtual information about the person in front of the mirror is acquired and presented in a comic-like thought-bubble (Gradman, 2010). Photograph by Bryan Jones. Printed with permission.

us about something intangible. A well-known device that does this is a hand-held Geiger counter. This device informs us about our surroundings and produces audible clicks that correspond to the amount of radiation that is present at the current location. Another application that informs us about our intangible surroundings is the app *Shazam* (2008). This app listens to our environment and displays information about what songs or TV shows are currently playing.[6] In fact, the virtual can even inform us about things that do not exist at all. For instance, in 1997 de Ridder realized an audio tour in the *Stedelijk Museum* in Amsterdam that told visitors about the meaning of 'invisible' elements in the museum (*history and archive - Stedelijk Museum Amsterdam* n.d.).

[6] It is rather ambivalent whether music and televisions shows should be considered something real or something virtual. If we treat them as something virtual, this example shows that the virtual also can inform us about other virtual aspects of our surroundings. In any case, this example demonstrates that virtual content cannot only inform us about physical aspects of our reality, but also augment non-tangible aspects of our surroundings.

Whereas typically, virtual content is used to inform us about the real environment, the opposite is possible as well. For instance, in the Dutch seaside resort "Kijkduin" a physical sign describes the resort as the "Pokémon capital of the Netherlands", and thus informs visitors about the presence of the virtual Pokémon characters in the area (see figure 4.5.

As the various examples illustrate, content-based relationships between the virtual and the real are very common, both in the traditional field of AR, as well as in other areas. As we have shown, content-based augmentation can take many different shapes. One key form of content-based augmentation is augmentation by means of text, which can, for instance, be presented as a visual overlay, on a separate screen or in the form of a spoken text. However, the information can, for instance, also be conveyed by means of symbols (e.g., arrows) or guiding sounds. Furthermore, the virtual can relate to many different aspects of the real world. For instance, it can inform us about objects, places,

Figure 4.5: A physical information board informs visitors about the presence of virtual Pokémon in the environment. Photograph by ANP. Reprinted under fair use.

people and processes.

What roles can virtual information that relates to the real world content-wise play in our surroundings? Just like virtual objects that appear in space, virtual information presented on a separate screen or via speakers can supplement and extend the real environment. Hence, content-based AR also serves as a basis for what we call *extended reality* (which will be discussed in the following section).

Furthermore, in some cases, a real environment might be considered incomplete without additionally presented information about this environment. E.g., we can imagine an artwork where the descriptions provided by the audio guide are an integral part of the artwork, rather than supplementary information. In this sense, the virtual information can *complete* a real environment. Thus, just like presence-based AR, content-based AR can also serve as a basis for *hybrid reality* (see section 4.5).

We have suggested that presence-based AR can serve as a basis for *diminished reality* (section 4.6). It is difficult to imagine how content-based AR would allow us to seemingly remove content from the real world. We thus see no direct link to *diminished reality*. However, the additional virtual information that is presented in content-based AR might be able to distract us from some aspects of the real world. (Also, additional information might, e.g., take away our fear or discomfort in certain situations.)

Additional information that relates to our surroundings can change our experience of these surroundings (e.g., knowing more about an artwork can make us appreciate it more or see it differently). This means that content-based AR can also lead to what we call *altered reality* (cases, where the real is transformed by the virtual). However, this is not unique to AR (physically presented information can likewise

transform our experience of the real world). Because this phenomenon has mostly been explored in the context of presence-based AR, we will focus on examples where the presence of virtual content in real space transforms the real world when discussing *altered reality* in section 4.7.

Information that relates to our surroundings on the content-level can also be used to allows us to perceive more about the world. An example is the above mentioned Geiger counter, which translates the amount of radiation that is present at the current location into audible clicks. Although these clicks are only presented (rather than present) in the space, they translate aspects of the real world that we cannot perceive into virtual information that we can perceive. Hence, just like presence-based AR, content-based AR can be used for *extended perception*. More examples of *extended perception* will be discussed in section 4.8.

Finally, it is possible to imagine interaction between real content and virtual information that is solely presented (rather than present) in the real environment. For instance, a character on a digital advertisement board might speak to a by-passer. However, we believe the presence of virtual object in real space (and thus presence-based AR) opens up much more compelling and unique possibilities for interaction, as here both the virtual and the real seem to occupy the same space. This is why our investigation of physical relationships (section 4.9) and behavioral relationships (section 4.10) between the virtual and the real focuses on presence-based rather than content-based AR.

As we have shown, both content-based relationships and spatial relationships can serve as a basis for many subforms of AR. These subforms will be discussed in the following.

4.4 *Extended Reality: The Virtual Supplements the Real*

All forms of AR are characterized by a combination of the virtual content and the real world. This virtual content can play various different roles in the world. For instance, it can remove or transform real objects. However, most commonly, the virtual extends the real world. With this, we mean that the environment appears to contain *additional* virtual elements or *supplementary* content. We propose to call this subform of AR *extended reality*. It is important to not confuse this subform with AR in general. From a technological perspective, AR always presents additional virtual content to the participants. However, from a perceptual perspective, this additional content can play many different roles, such as supplement, diminish or transform reality. With *extended reality*, we refer to those cases where the virtual supplements the real and where the participant experiences additional virtual content in the environment.

The extension of the real world can take two main forms. First of all, the virtual can extend the real world by providing information

that relates to the environment on the content-level. This possibility has been discussed in depth in section 4.3. In such cases, the information extends the real, because it provides us with additional facts, instructions or stories. We can think of information that relates to our surroundings, such as an audio guide or museum app, as a *supplementary* layer of content—something extra or additional that becomes part of, shapes and extends the experience of the real world.

A second form in which the virtual can extend the real is in the form of additional virtual objects and elements that seemingly exist in the real space. As we know, creating the impression of virtual objects existing in the real world is one primary goal of existing AR research. Accordingly, we can find a huge variety of AR projects where virtual elements appear to exist in the real world and supplement the space.

In the following, we will provide a selection of examples that illustrate the many forms of how the virtual can extend the real. Because the addition of virtual elements to the real world plays such a prominent role in existing AR research, this section will be rather comprehensive. Also, because AR is very focused on making virtual objects appear in the real environment, many such examples will be included. Due to the length of this section, and because our senses work quite differently when it comes to perceiving virtual elements in space, we have decided to divide this section into several subsections: We first look at examples where visual elements extend the real environment. This form of AR is very common in the context of traditional AR. Subsequently, we explore approaches that have received less attention in the context of traditional AR research so far, and look at sonic, tactile, olfactory and gustatory extensions of the real world as well as at examples of multimodal additions.

4.4.1 Visual Additions

Examples of applications where additional virtual objects look like they existed in real space are very popular. They can, for instance, be found in the entertainment context, in manufacturing, in the medical domain, in education and in the art world.

As mentioned, the presence of virtual content in real space plays a fundamental role in the first so-called augmented reality prototype by Caudell and Mizell (1992), which displays virtual instructions about manufacturing processes in a way that they appeared in 3D space. Many others have followed Caudell and Mizell's example and created projects where virtual information appears in real space and is spatially aligned with physical objects. For instance, Feiner, Macintyre, et al. (1993) have presented an AR system that displays maintenance instructions for an office printer in real 3D space, spatially aligned with this office printer. Comparably, in the medical domain, research has focused on AR systems that display medical information in phys-

ical space, and more specifically, inside of the patient. For instance, the above-mentioned system by Bajura, Fuchs, et al. (1992) visualizes ultrasound echography data within the womb of a pregnant woman.

The fact that additional virtual content appears in real space opens up many possibilities for new forms of entertainment applications that make use of the player's real environment. For instance, AR games, such as Sphero (*Sphero* 2011) and ARQuake (Piekarski and Thomas, 2002; Thomas et al., 2000) present us with virtual game characters that move through the real environment.

For many projects, it is not only important that virtual content appears in the real environment, but also important that the virtual content appears in the same environment as the participant. Presumably, this is the case in the context of exposure treatment, where virtual fear stimuli can be displayed in the environment of the participant. For instance, Corbett-Davies, Dünser, and Clark (2012) have realized an AR project where virtual spiders appear in the real environment and even can be carried around and occluded by the user's hand.

Virtual content that is added to a real environment can allow people in this space to more effectively work together with *remote collaborators*. This is because unlike real content, virtual content can be modified both by people on site and remote colleagues. Such a collaborative AR scenario has been explored by Akman (2012). The author designed and implemented a multi-user system for crime scene investigation. Investigators are equipped with an AR headset, and can annotate the scene with virtual tags (e.g., to record the possible trajectory of a bullet). Both on-site team members and remote colleagues can subsequently see and modify these virtual annotations. Also, remote team-members can place additional virtual information in the scene.

Aside from extending in the environment of the participant, virtual content can also supplement *mediated* environments. For example, Scherrer et al. (2008) have created an augmented book that reveals additional 2D objects when this book is placed under a web-cam and viewed on the computer screen. These objects appear in the space that is depicted on the book's pages as well as seemingly float off the pages and enter the real environment that surrounds the viewer of the book.[7]

[7] It can be argued that such examples fall out of the scope of our definition of AR because the virtual content is not experienced in relation to the real world.

At times, virtual content is designed to extend or supplement *any* environment. In other words: sometimes, it does not matter in which specific environment virtual content appears. For instance, the Dutch super market chain Albert Heijn has published a series of stickers about dinosaurs, some of which make a virtual dinosaur appear above the card when the card is viewed through their smartphone application. In this case, where the card is viewed does not matter. The dinosaur appears as if it existed in the real environment, independently of where in the world, or in which context the card is scanned.

At other times, virtual content is designed to extend or supplement a specific real environment and only can be experienced in this space.

For instance, the artists Sander Veenhof and Mark Skwarek have created an additional virtual art exhibition in the famous MoMA (Museum of Modern Art) in New York City (without involving the museum itself) in 2010 (Veenhof, 2016, and personal communication)). Viewing the museum through the lens of their phones with the Layar application, visitors were able to see additional virtual artworks, as well as a virtual 7th floor alongside the actual physical artworks that were exhibited at that time. Judging from the video that shows the exhibition (Veenhof, 2010), the virtual artworks certainly became a crucial part of the museum experience.

Although technological questions fall out of the scope of this chapter, we would like to note that the virtual objects are typically displayed by means of head-mounted displays or hand-held displays. In addition, visual virtual content can be integrated into the world directly, e.g., by of projectors. This is typically referred to as spatially augmented reality (Raskar, Welch, and Fuchs, 1998) or spatial augmented reality (Bimber and Raskar, 2005). An example of such a spatial augmented reality project has been realized by Benko et al. (2014), who use three projectors in combination to allow two participants to see virtual content in the real environment, and, for instance, toss a virtual (projected) ball back and forth through the space between them (see section 3.1.2).

4.4.2 *Auditory Additions*

Aside from visual virtual elements, sounds can also extend and supplement the real environment. In the following, we will review examples that illustrate this point and briefly discuss the potential and unique opportunities that the addition of sound offers.

Like visuals, sounds are often used to convey the presence of virtual objects in real space. A project that uses audio sources for such purposes is the "Corona, an audio augmented reality experience" by Heller and Borchers (2011). In this project, the historic town hall of Aachen (Germany) was overlaid with a virtual audio space, representing an event from the 16th century. Virtual characters of people that attended the original event were placed at certain positions in the real space by means of spatialized sound. Another project, where sonic virtual content extends the real world is the SoundPacman game by Chatzidimitris et al. (2016) (also mentioned in chapter 3). This game makes use of 3D sound in order to give game elements a position in the real physical environment and to communicate their location to the player. Like in the original PacMan game, the player has to avoid being caught by the ghosts, and hence, has to monitor their spatial position.

In our opinion, these projects demonstrate an interesting quality of sound. Sound can be used convey a *spatial* presence of content in

the environment without implying a *tangible* or *material* presence of this content. This fits well with the example of ghosts (Chatzidimitris et al., 2016) and also with the idea of representing characters from the past (Heller and Borchers, 2011). We believe it makes sense for those characters to not appear as if they were present in the space in a material, tangible way.[8]

[8] Of course, sound is not the only medium that can create a spatial presence without implying a tangible and/or material presence. For instance, similar effects could be achieved with visually displayed semi-transparent virtual ghosts.

Whereas vision-focused projects typically focus on giving virtual *objects* a position in real space, sound-related projects often also focus on giving other types of virtual content (non-objects) a place in the real environment. For instance, the interactive sound installation Audio Space (2005) by designer and artist Theo Watson allows participants to hear audio messages that have been recorded and "left behind" by previous visitors in the same physical space. The audio messages are spatialized in 3D and seem to originate from the spot where they have been recorded. In addition, the participants can leave their own audio messages at any point within a room, simply by speaking into their microphone at the intended spot. (In later versions of this installation, sound effects were applied to the recorded messages, creating a more abstract sound environment.) This project showcases another quality of sound: it is relatively easy for participants to create virtual content in the form of sound and to add this content to the real world. (Arguably, it is currently much easier to record a spoken message than, for instance, to create a virtual object with a 3D modeling program.)[9]

[9] However, current technological developments, such as the integration of 3D camera's in smartphones undoubtedly make the creation of virtual 3D models much easier.

Another project that does not work with virtual *objects* is the LISTEN project (Eckel, 2001). This project includes the use of virtual *soundscapes* that, among other things, are used to create context-specific atmospheres. This project shows that sound not necessarily has to represent *objects* in space in order to extend the world.

If we compare the sonic examples to the previously discussed visual additions, it becomes clear that sonic additions provide us with possibilities that visual additions cannot offer us. One obvious point is that in contrast to vision, sound also allows us to hear what happens *behind* us. For instance, we can imagine a scenario in which virtual footsteps follow a participant around, only to stop and disappear when the participant stops walking and turns around. Naturally, such an experience that is based on what happens behind the participant is much more difficult to realize through visual additions.

4.4.3 Haptic Additions

In addition to projects that allow us to see or hear virtual objects, we can also identify projects that extend the real world with 'feelable' virtual objects. Although these projects also make it seem as if additional virtual content existed in real space, they often are not presented in an AR context and have received little attention in existing AR discourse. In the following, we will review some of these projects and place them

in the AR context.

An example of such a project that allows us to feel virtual objects in a real-world setting is the Gravity Grabber (mentioned in chapter 3) by Minamizawa, Fukamachi, et al. (2007). This wearable device consists of fingerpads that allow participants to perceive the ruffle of the water in a real glass, although they are actually holding an empty glass.[10,11]

Another project where the presence of something virtual is perceived tangibly is Sekiguchi et al. (2005)'s so-called Ubiquitous Haptic Device. When shaken, this little box conveys a feeling of a virtual object being inside the device. In contrast to the Gravity Grabber (Minamizawa, Fukamachi, et al., 2007), the tactile feedback is not simulated by a wearable device but by the box itself. Arguably, these projects qualify as AR and extend the real world, because they allow us to experience (and interact with) additional, simulated objects in the real world.

Furthermore, quite some research exists about providing tactile sensations when a user moves their hand through the air. For instance, Minamizawa, Kamuro, et al. (2008, e.g.,) propose a glove that a user can wear and that provides tactile feedback in order to convey the presence and spatial qualities of virtual objects. Another approach to haptic extended reality is the use ultrasound to provide mid-air haptic sensations. Hoshi, Takahashi, Nakatsuma, et al. (2009); Iwamoto et al. (2008) and Hoshi, Takahashi, Iwamoto, et al. (2010) have developed a tactile display that deploys airborne ultrasound and utilizes acoustic radiation pressure to create sensations that humans can perceive with their skin. Simply put, their display radiates ultrasound. When a user's hand interrupts this propagation of ultrasound (i.e., 'gets in the way'), a pressure field is caused on the surface of their hand. Because the pressure acts in the direction of the ultrasound propagation, the ultrasound "pushes" the hand and the user feels tactile sensations (Hoshi, Takahashi, Nakatsuma, et al., 2009). (The system can control the spatial distribution of the pressure using wave field synthesis.) What makes this approach special is that users can feel virtual objects, such as virtual raindrops or small creatures, on their hands without making any direct contact with a device. Hoshi, Takahashi, Nakatsuma, et al. (2009) combine this tactile display with a holographic visual display, which ultimately allows participants to both see and feel the virtual objects (see figure 4.6).

Before moving on, it should be noted that many of the reviewed techniques to make virtual objects tangible have not explicitly been explored in the context of AR yet. For instance, Minamizawa, Kamuro, et al. (2008) do not explicitly address whether they envision the tactile virtual objects in a virtual environment or in the context of the real world. However, we believe that techniques that allow us to display virtual objects in space can typically be used to extend the real environment and thus, used to create AR.

[10] The author of this thesis was able to experience this device in the context of a different application, where it allowed participants to feel virtual marbles moving in a transparent little empty box when shaking this box.

[11] The recent paper "Altered Touch: Miniature Haptic Display With Force, Thermal, and Tactile Feedback for Augmented Haptics" (Murakami et al., 2017) shows that the Gravity Grabber is now used in combination with a thermal display. The resulting system has been used to alter softness/hardness and hot/cold sensations in several augmented reality scenarios.

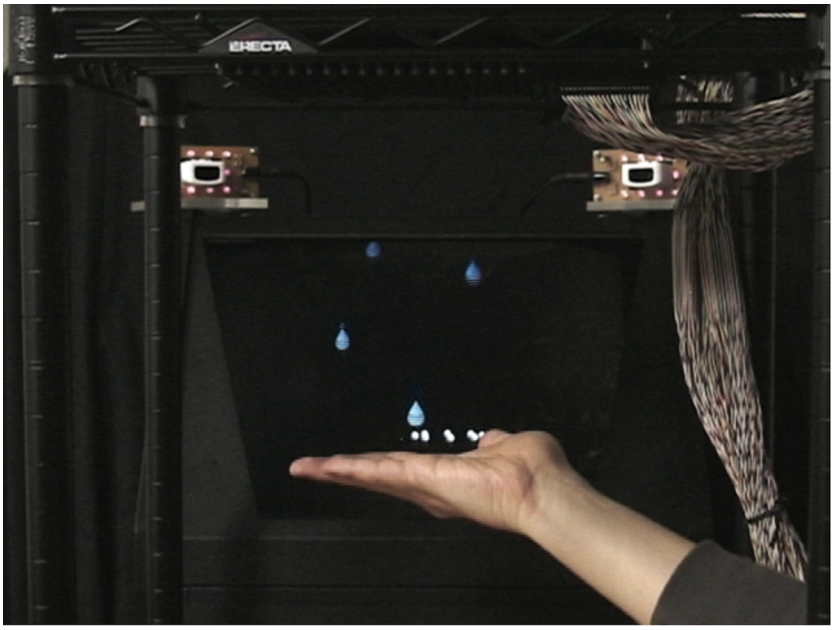

Figure 4.6: A combination of a tactile display and a holographic display allows participants to see and feel raindrops hit their palm. Reprinted from T. Hoshi, M. Takahashi, K. Nakatsuma, et al. (2009). "Touchable holography". In: *ACM SIGGRAPH 2009 Emerging Technologies*. ACM, p. 23. Reprinted under fair use.

4.4.4 *Olfactory and Gustatory Additions*

Aside from using sonic, tactile and visual stimuli, the real world can also be extended by means of olfactory or taste stimuli. However, our sense of smell and taste do not allow us to experience the same kind of spatial relationships between objects as our other primary senses do. For instance, we can see a virtual strawberry lying in front of a real banana, but we can presumably neither smell such relative positions in real space nor taste that the banana is lying behind the strawberry.[12]

Even if a smell does not convey us with an exact *location* of its source, it might nonetheless convince us of the *presence* of certain elements in the environment. For instance, if we look at the real world, the smell of a specific perfume might be enough for us to know that a certain colleague is in for work today and an unpleasant smell that follows us around might make us check our shoe soles for dog dirt or convince us that a baby's diapers have to be changed. Similarly, the taste of a meal might allow us to conclude about its ingredients, such as the presence of certain spices.

A question that arises is what exactly qualifies as virtual content when we are dealing with olfactory and gustatory information. Are we dealing with virtual strawberries if we can taste them in our yogurt, although the little pieces are made of pumpkin and artificial flavors? Are we surrounded by virtual flowers, if we smell them, but all we are actually dealing with is the new perfume of our colleague? As mentioned, in this thesis we consider stimuli as virtual if they have been synthesized or do not directly originate from their original source.

Regarding "virtual tastes", we can create taste experiences by stimu-

[12] Our senses of smell and taste work differently than our other senses. We can only perceive olfactory and gustatory information if our receptors are in direct contact with the molecules that contain this information (Köster, 2002). (In this sense, it is similar to touch, which also requires *direct* contact with tactile stimuli). In line with this, the sense of smell and the sense of taste are sometimes considered "near" senses (Köster, 2002). However, there is still some uncertainty about the spatial information that humans derive from olfactory cues. For instance, Köster (2002) claim that olfaction is "not involved in involved in spatial orientation" (p. 30). In contrast, Jacobs et al. (2015) have shown that humans can use a unique odor mixture to learn a location in a room and subsequently, navigate back to this location with only olfactory information guiding them, which suggests that humans can make use of olfaction in orientation.

lating the tongue with electric current. This effect is nowadays known as "electric taste" and was discovered by Sulzer as early as 1752 (Bujas, 1971). Reportedly, Sulzer touched two interconnected but different pieces of metal with his tongue, and experienced a ferro-sulphate-like taste, although the metals themselves were tasteless. Furthermore, presenting odors in the mouth can cause taste experiences (Lawless et al., 2005). For AR, what matters is if such taste experiences are experienced as related to the real (e.g., related to some real food).

When it comes to odors, these can be presented in real space by means of olfactory displays. One of the few projects that work with presenting smells at a certain position in real space is the "Projection-Based Olfactory Display with Nose Tracking" presented by Yanagida et al. (2004). This device is different from typical Olfactory displays in the sense that it does not focus on the synthesis of odors but on the spatiotemporal control of the odor. This means that unlike more common approaches, their prototype not simply diffuses odor in space but instead, projects scented air to the nose of people in the space. To do so, they track a participant's head/nose and use an air cannon aiming at the nose to transport/transfer clumps of scented air from the cannon to the user's nose. While the authors place their research in the context of VR, the actual proposed prototype and experiments simply "project" scents in the real environment. Because the participants experience virtual content as part of the space, this scenario can be interpreted as an olfactory example of AR. A challenge that comes with the presentation of virtual smells in the real space is that smells cannot easily be removed from the environment after they have been dispensed.

Existing AR research has paid little attention to the possibilities of using olfactory and gustatory information to supplement the real world. We suggest exploring this topic further in the future.

4.4.5 *Multimodal Additions*

In addition to using only one single modality to present additional and supplementary virtual content in real space, some projects also use combinations of different sensory stimuli. For instance, AR projects can make use of a combination of visuals and sound (such content is also referred to as *audiovisual* content). An example is the mobile AR game *GeoBoid* by Lindeman, G. Lee, et al. (2012). In this game, players are surrounded by flocks of bird-like virtual creatures called GeoBoids. These creatures are represented both visually as well as by means of spatialized audio using the player's phone. While there seem to be few projects that use sound as an integral (important) part of an audio-visual AR experience, sound is often used to accompany primary visual content. An example is the mobile game "Pokémon GO" (*Pokémon GO* 2016). Here, the visual creatures occasionally make

a sound and the movement of Pokéballs is accompanied by sound-effects.

In addition to AR applications that make use of audiovisual additions, we can also find several projects that allow participants to both see and feel virtual objects in real space. One early project that puts the idea of viso-haptic virtual objects in AR into practice, has been realized by Vallino and C. Brown (1999). Their augmented reality project displays virtual images in a live video stream of a real scene but also incorporates a Phantom force-feedback device that simulates the tactile characteristics of the object. This device has similarities with a small robot arm (cf. Vallino and C. Brown, 1999) with a thimble at the end, into which a user inserts their finger. It has motors driving each joint, which generate the force feedback needed to simulate the touch of virtual objects. Placing their finger in the device's thimble, the participant can feel the surface of the virtual object, experience its weight and dynamic forces, as well as move the object around within the real environment. (In their demonstrations, participants can, for instance, experience a virtual globe, spin it around its axis, feel the difference between water and land, or move a virtual cube around in real space with their finger.)

By now, this phantom-based approach has been pursued several times. For instance, Bianchi et al. (2006) have developed a similar system and realized an AR-based ping-pong game that allows players to play with a virtual ping-pong ball in the real environment and feel the impact of the virtual ball on a simulated bat. Later on, a two-player version of the same concept has been realized by Knoerlein et al. (2007).

4.4.6 Short Summary Extended Reality

To summarize, virtual content can extend and supplement an otherwise real environment. If we want to extend or supplement the real, we can build on both content-based and spatial relationships. In both cases, the participant has access to *more* content in the environment due to the virtual additions. Most commonly, AR extends the world by means of virtual objects that appear to exist in the real environment. There are many ways of conveying this presence of virtual elements in real space. Visual, sonic and tactile stimuli are particularly powerful to add virtual elements to our otherwise real surroundings, and they can be used to make them appear at specific locations in the environment. We propose to refer to forms of AR where participants experience additional, supplementary virtual content in their surroundings as *extended reality*.

4.5 *Hybrid Reality: The Virtual Completes the Real*

What role does the virtual play in the otherwise real environment? In the previous section, we have encountered examples where virtual content is designed to supplement the real world and serves as "something extra" and *optional* in the otherwise real environment. In such cases, the real surroundings can also be considered "complete" without the virtual additions. For instance, a museum is complete without a virtual museum guide, the streets are complete without virtual driving instructions or virtual Pokémon that appear on the sidewalk.[13] Because the real world is complete on its own, it can be experienced in two contexts: either independently, or in relation to the virtual additions (and hence, as part of an AR scenario). However, at times, the virtual not only supplements but rather *completes* an otherwise real environment (or a real object in the environment). In such cases, the physical environment (or object) is incomplete without the virtual additions, and the virtual is required. In line with this, the real is not intended to be experienced on its own—its sole purpose is to be experienced as a part of a mixed virtual-real scenario, and thus in the context of AR.

[13] Even if the virtual does not play an essential role in the otherwise real environment, it usually plays an integral role in the experience of the *augmented* environment.

Typically, such scenarios in which the virtual completes the real are achieved by not only designing virtual additions for an existing real world but by designing a mixed environment or object that consists of both a real component and a virtual component from the very start. In such cases, the virtual can fill in aspects that are missing in the real world, and vice versa—the virtual and real complete (and in this way augment) one another.

The idea of creating hybrid objects is often applied in the field of augmented prototyping. Like the above-mentioned augmented zebrafish project, augmented prototyping makes use of digital imagery that is projected onto physical models, resulting in partially virtual, partially real prototypes (see, e.g., Verlinden et al., 2003). A setup for such hybrid models has, for instance, been proposed by Raskar, Welch, and Chen (1999). Their research explores the use of light projectors to augment physical models with virtual properties. For instance, they use ceiling-mounted projectors to extend physical objects from wood, brick, and cardboard on a tabletop with virtual textures and colors.

In the context of hybrid reality, it is important to note that the virtual not only completes the real but that the virtual and the real complete each other. In projection-based setups, the virtual usually completes the real *visually*, whereas the real completes the virtual *physically*. However, other possibilities exist. For instance, a karaoke version of a song deliberately leaves out certain elements of a song, which then have to be filled in live by a participant. Ideally, the real singing of the participant mixes in with and becomes part of the played music.

As the discussed examples show, the virtual and the real can com-

plement and complete each other in different ways. For instance, the virtual can complete the real on a musical level, or visually. Similarly, the real can complete the virtual musically, or physically.

If we look at the entirety of reviewed examples, we can identify two main approaches to creating AR: First of all, we can take the real world as it is, and aim at creating virtual content that relates to this world. Furthermore, we can give shape to virtual content *and* the real world. This approach, too, allows us to establish relationships between the virtual and the real. When desired, it allows us to make sure the two complete one another. Considering that AR environments and experiences are characterized by the relationships of the virtual and the real, we believe that designing both the virtual component and the real component with respect to each other offers many possibilities for creating and shaping AR experiences.

In order to be able to easily refer to environments and objects where the virtual completes the real, we propose the terms *hybrid object*, *hybrid environment* and more generally, *hybrid reality* to denote such scenarios.[14] We see hybrid objects and environments as a subgroup of AR. Hybrid objects and environments are intended to be experienced in their hybrid form—neither the virtual nor the real makes sense on its own. (This sets hybrid objects and environments apart from many other augmented objects and environments that also can be experienced without visual additions.)

[14] In existing AR research, there is no clear, agreed upon definition of what constitutes a hybrid environment or object and the term "hybrid" is only used occasionally. For instance, Lok (2004) use it to refer to virtual environments that contain virtual representations of real objects (or in other words, incorporate real objects into virtual environments). In contrast, Raskar, Welch, and Fuchs (1998) speak of a "hybrid environment" to refer to AR environments that are build with a combination of different technologies, such as a combination of projectors as well as see-through head-mounted displays.

4.6 *Diminished Reality: The Virtual Removes the Real*

As we have seen, virtual content often supplements and augments the real world in the sense that there is *more* content in the environment. However, we can not only use virtual information to add content to the world—it can also be used to hide or seemingly remove real elements from the world.

The process of removing real content from our perceived environment is also referred to as "diminished reality". Diminished reality is sometimes seen as its own field of research (e.g., Herling and Broll, 2010). In fact, we could argue that it forms a "counterpart" to augmented reality, as it is focused on removing rather than adding something to the world. Yet, diminished reality is also considered a subset of AR (e.g., Azuma et al., 2001).

In this chapter and throughout this thesis, we treat diminished reality as a as a form of AR. We believe this makes sense because diminished reality applications also present us with virtual information that relates to the real world. Just like the creation of additional objects in the perceived environment, the deliberate removal of real objects from our perception of the world is realized through the addition of virtual content. More than that, the addition of virtual content and the removal of real content from the perceived environment often go hand

in hand. If, for instance, a virtual chair appears to stand in front of a real desk, parts of this real desk will be hidden from our view. In this sense, adding virtual information to our perception of the world on the one hand, and removing real information from our perception on the other hand, can be considered two sides of the same underlying process.

Whereas many AR projects are focused on adding virtual elements to the world, AR research and development has also explicitly focused on how to remove real elements from the world. One of the key questions is how to fill the space of the removed object. Many different approaches have been proposed to make it seem as if a real objected did not exist. For instance, Herling and Broll (2010), have presented a system that can remove arbitrary real objects from a live video stream of the environment by filling the resulting empty space using an image completion and synthesis algorithm. Simply put, their algorithm removes the area in which the undesired object is located and uses information in the remaining parts of the video image to fill up this area.

Zokai et al. (2003), too, have been working on removing real objects from (a view) of the real word. However, unlike Herling and Broll (2010), they use images from different viewpoints in order to determine what lies behind the removed object. Consequently, their approach replaces the real-world object with an appropriate background image.

A yet different approach to removing real content is found in the art context. Instead of simply *removing* elements from the world, the artist Julian Oliver has worked with the principle of *replacing* real content with different, arguably more desirable, virtual content. His mobile augmented reality project called *The Artvertiser* removes advertisements in the city and replaces them by art. (In this way, the project is quite similar to the previously mentioned work by Mann and Fung (2002) that likewise can replace advertisements.)

Just like the general field of AR, diminished reality is very focused on vision. In other words, real objects are commonly removed from our *view* of the world. However, the idea of removing aspects from a person's experience is not unique to the field of visually augmented reality. For instance, the same idea has quite a tradition in the audio context.[15] Here, active noise control systems are used to reduce undesired real sounds from a user's perception. This is achieved by playing back additional sounds that are specifically designed to cancel out unwanted sounds (Leitch and Tokhi, 1987).

[15] The fact that similar concepts have been applied in the audio domain for a long time has also been pointed out by Herling and Broll (2010).

The idea of presenting additional information in order to not make us notice existing aspects of the real world is also a common everyday strategy when it comes to unwanted smells or tastes. Unlike with sound, we cannot simply dispose a smell or taste signal that cancels out existing tastes or smells. However, additional smells or tastes can

mask, overpower or subdue existing smells and tastes. For instance, many people use deodorant to cover up (and ideally prevent) body odor.

A project that approaches the idea of removing taste and smell differently is the "Straw-like User Interface" by Hashimoto et al. (2006). The project explores *removing taste and smell* from the drinking experience by solely simulating the tactile sensation of drinking at the mouth and lip. In their own words, they hope to allow participants to "experience a new sensation by extracting the drinking sensation from that of taste and smell, and in doing this present a comfortable and exciting sensation to the lips and mouth" (p. 2). While the interface consciously does not provide taste and smell sensations, it simulates and combines three aspects of the drinking experience: (1) the pressure change in the mouth (normally caused by foods blocking the straw), (2) vibrations at the lips and (3) sounds.

Of course, the "Straw-like User Interface" does not actually remove something real from a real experience. Rather, it only simulates parts of a real experience. However, by only simulating *some* properties and leaving out others, they indirectly simulate the removal of those properties that have not been simulated. In this sense, many AR projects might allow us to explore the removal or absence of real aspects from objects. For instance, we might be able to see a virtual teapot, but not be able to feel anything when we touch it. Likewise, we might see a spider walking over our hand, but not feel it on our skin (Corbett-Davies, Dünser, and Clark, 2012). We assume, such partial simulations might not only allow us to experience the *presence* of an object but might also allow us to experience the *absence* of some of its characteristics or aspects, such as the absence of tactile qualities. However, this remains speculative. A question that could be researched in the future is how partial simulations are experienced. For instance, it would be interesting to know whether and under which conditions we experience a solely visual simulation of a teapot as an *intangible* teapot. Similarly, it would be interesting to further research the experience of *partial* removals. For instance, what do we experience when we happen to touch an object with our hands that has been removed from our view by means of diminished reality technologies—do we experience the object as being *invisible*?

To summarize, virtual content can be used to add elements to the world, but also can be used to remove real elements (or aspects of real elements) from the world. In the context of traditional AR, the focus lies on removing real objects from our view. However, we can also use additional sonic, olfactory or gustatory information to mask real sounds, smells or tastes. This is quite different from AR in the traditional sense. However, we believe that in cases where virtual stimuli (e.g., synthesized stimuli) are used to seemingly remove real stimuli, we can speak of AR in the broader sense. After all, we are dealing

with additional virtual content that relates to its real environment.[16]

Whereas both traditional AR and AR in the broader sense can seemingly remove and replace some aspects of the real world, AR projects never replace the real surrounding world *entirely*. This sets AR apart from Virtual Reality (VR), where participants experience a completely synthetic environment, rather than a partially real, partially virtual environment (cf., e.g., Milgram and Kishino, 1994).

[16] When it comes to sound, one can argue that it actually falls within the scope of traditional AR, as the virtual and real sound waves have to be properly aligned with each other interactively and in real-time, so that the canceling effect is achieved.

4.7 *Altered Reality: The Virtual Transforms the Real*

The presentation or presence of virtual information in an environment always changes or transforms the environment. For instance, an environment is not the same when it contains virtual ghosts (Chatzidimitris et al., 2016), virtual spiders (Corbett-Davies, Dünser, and Clark, 2012) or virtual voices (Watson, 2005). Similarly, the world appears differently, if real objects are hidden from our view or undesired sounds are removed from our sonic environment. However, whereas many AR projects focus on adding or removing information, some projects explicitly aim at *transforming* the environment. In particular, many projects focus on transforming real-world objects. In the following, we will have a look at such cases where the virtual *transforms* the real. We propose to call this *altered reality*. Altered reality scenarios are very common and take many different forms. In the next sections, we explore how visual, tactile, sonic, olfactory and gustatory qualities of the real world can be transformed by means of virtual additions. Subsequently, we explore projects where the virtual seemingly transforms other aspects of the real world, such as the room temperature. Finally, we take a closer look at the transformation of multimodal perception and the phenomenon of cross-modal interaction, where information that stimulates one sense transforms our perception of information that stimulates another sense.

4.7.1 *Transformations in Visual Perception*

Transforming how real objects look is especially popular in the context of projection mapping and so-called spatial augmented reality. In projection mapping, light is used to project virtual content directly onto the real world. Spatial augmented reality more generally refers to all cases where virtual content is directly integrated into an environment (rather than, e.g., overlaid onto a participant's view)—including scenarios where projected light is used to alter the appearance of physical objects (Raskar, Welch, and Fuchs, 1998). [17]

Often this method of projecting virtual content onto the real world directly is used to seemingly transform the underlying real objects. An artist who works with this method is Pablo Valbuena. For instance, his video-projection on the city hall in The Hague called "N

[17] The terms "projection mapping" and "spatial augmented reality" are often used interchangeably. However, strictly speaking, the term spatial augmented reality is broader. It is not limited to the use of video projection technologies, but also includes other forms of embedding virtual content in the real world directly, such as the use of flat panel displays (cf. Raskar, Welch, and Fuchs, 1998).

520437 E 041900 [the hague city hall]" has followed this principle and has transformed the physical building into a large dynamic sculpture (Valbuena, 2008). Through the virtual projections, the city hall has gained virtual and dynamic properties, such as moving walls, or temporary convexities and indents.

Figure 4.7: A comparison between traditional augmented reality (left) and stylized AR (right) as implemented by Fischer et al. (2005). In both images, the teapot is a virtual object, while the cup and the hand are real. However, the stylized version uses an image filter and non-photorealistic rendering. Reprinted from J. Fischer et al. (2005). "Stylized augmented reality for improved immersion". In: *Proceedings IEEE Virtual Reality 2005*. IEEE, pp. 195–202. Reprinted under fair use.

The same concept also plays a role in the previously discussed context of augmented physical models and prototypes (see, e.g., Raskar, Welch, and Chen (1999) and Verlinden et al. (2003)). Using projections, physical models can quickly and cheaply be transformed and give us an impression on how an object would look with different types of colors or different textures.

In addition to projection-mapping, there are other means to alter how the real world looks. For instance, Fischer et al. (2005) have proposed to transform a participant's view of the real world with a painterly image filter in the context of video see-through AR.[18] More specifically, they suggest applying the same stylization to (1) the participant's view of the real world as well as (2) the virtual additions. Reportedly, this makes the virtual elements and the real world look very similar, and ultimately, makes it look as if virtual objects were an actual part of the real environment.[19,20]

[18] Video see-through AR captures the real world with (a) camera(s), combines the live video images with virtual imagery and present the result to the participant via a video display.

[19] In this project, the transformation of the environment does not seem to be the ultimate goal in itself. Rather, the transformation serves the purpose of making virtual objects mix in with the real environment.

[20] It is debatable if this transformed version of the real environment should be referred to as a *real* environment.

4.7.2 *Transformations in Auditory Perception*

The concept of changing qualities of the real world is also quite popular in the audio domain. For instance, mobile apps like *RjDj* (discontinued, see RjDjme (2008) for a video) and more recently, *The app formerly known as H _ _ r* (2016) and *Inception - The App* (2016) focus on transforming a user's real sonic environment. These apps use sound-input from a user's phone and apply filters and delays to transform the sonic environment of the user.[21]

The idea of "remixing" the sonic environment, which underlies these applications, is not new and has previously been explored in the art context. For instance, the artist Akitsugu Maebayashi has worked with similar concepts with his sound work *Sonic Interface* from 1999. The project makes use of a laptop, headphones and microphones and

[21] These applications are implementations of so-called "Reactive music" (Bauer and Waldner, 2013; Bondo et al., 2010; RjDj, n.d.). Reactive music reacts to the listener and his environment in real-time, e.g., by using the data from a phone's camera, microphone, accelerometer, touch-screen and GPS as input. Unlike traditional music, reactive music is distributed in the form of software that produces the actual music. A platform that provided the possibility for sharing and experiencing reactive music is the discontinued "RjDj" application (RjDj, n.d.).

uses delays, overlapping repetitions and distortions in order to recompose ambient sounds in urban space (Maebayashi, 1999; Unstable Media, n.d.). Judging from the description of the work found on the website of the Unstable Media (n.d.), the resulting soundscapes break the usual synchronicity between what one hears and what one sees.

In addition to transforming the general sonic environment, we can also change the sound of a specific object in the environment. This approach is common in a musical context, where musicians often use audio effects, such as vocoders, filters and delays to change the sound of their instruments. However, we can also imagine changing the sound of everyday objects in a similar way. For instance, the closing sound of a car door might be altered by sensing the original sound with microphones, emphasizing certain frequencies and playing the result back by embedded speakers. Likewise, we might transform the physical clicking sound of a clock by means of audio effects. This might, for instance, allow us to transmit additional information about the current time with the ticking sound.

In our opinion, the idea of physically embedding speakers into real objects to make them sound a certain way can be considered a sonic form of so-called spatial augmented reality. As mentioned, the concept of spatial augmented reality refers to cases where virtual content is embedded in the real world directly. Typically, the concept is discussed in a visual context, and used to describe cases where virtual content is embedded into the environment by means of projectors or flat panel displays (cf. Raskar, Welch, and Fuchs, 1998). However, we can apply the same concept to sound, and augment the real world by embedding *sonic* virtual content in the real world directly, e.g., by means of loudspeakers that are placed in the environment or embedded inside physical objects.[22]

While the discussed projects differ from traditional, registration-based AR applications on a technical level, they share important conceptual and experiential qualities: Judging from our own experience with current apps such as *The app formerly known as H _ _ r* (2016), the virtual sounds are perceived in the context of the real world, as linked to real events, and as related to our surroundings. We hence believe that such scenarios fall within our definition of AR.

[22] Ultimately, this would suggest that we can see a teddy bear that emits a pre-recorded grumble sound when it is shaken as an augmented object.

4.7.3 *Transformations in Haptic Perception*

In addition to changing how the real world looks and sounds, we can also find also various projects that focus on changing how the real world feels. This idea has a long tradition in AR research. For instance, in his seminal review of AR, Azuma (1997) suggested the idea of augmenting the feel of a real desk, "perhaps making it feel rough in certain spots" (p. 361) by means of gloves with embedded effectors. If we look at the current AR research landscape, such tactile

transformations have become possible—even without gloves. A tactile technology that enables feeling virtual textures on real surfaces is the previously mentioned REVEL device (Bau and Poupyrev, 2012). This device injects electrical signals into a participant's body and thereby allows them to feel virtual textures when running their hands over real objects and surfaces.

Another project that focuses on changing how an object feels when we touch it, and more specifically, on altering how warm or cold it feels, has been conducted by Ho et al. (2014). With their study, the researchers address the common belief that the color blue evokes cold feelings whereas the color red evokes warm feelings. Their study is based on several experiments in which participants touch an object with their hand and subsequently judge whether the object felt warm or not. The effect of color on temperature judgments was investigated by manipulating either the color of the object or the color of the participant's hand. (The color of the object was altered physically, whereas the hand color was changed by projecting either blue or red light onto the hand. However, we assume that similar results can be obtained when an object's color is altered virtually.) In contrast to the common belief, their results indicate that blue objects are more likely to be assessed as warm than red objects of the same temperature. A red object, relative to a blue object, was found to raise the lowest temperature required for an object to be judged as warm by about by about 0.5°C. Similarly, a blue hand, relative to a red hand, was found to raise the lowest temperature for the object to be experienced as warm by about 0.5°C. As the researchers elaborate, "this change [in the lowest warm temperature] is sufficient to induce a clearly perceptible change in the perceived temperature of an object in contact" (p. 2).[23] Although this project shows that color can alter temperature judgments in an experimental setting, more research is needed to explore whether virtual colors can be used to transform our temperature experience of real objects in our everyday world.

[23] As Ho et al. (2014) propose, the fact that the result seems to contradict the common belief can be explained by the hypothesis that the color can modulate the *expected* temperature of the object. In line with this, the researchers interpret their findings in terms of "Anti-Bayesian" integration, which suggests that our brain integrates the felt temperature with those prior expectations in a way that emphasizes the difference between them.

Other projects likewise explore the possibilities of visually altering how an object feels but focus on transforming other qualities of the object, such as its softness. As mentioned in section 3.3, Hirano et al. (2011) and Sano et al. (2013) use an HMD to display different computer-generated deformations on an object, when it is pushed down by a participant. Their experiments show that the perceived softness can be manipulated by means of virtual dents, without changing the actual material. The larger the dent caused by pushing the object appeared, the softer seemed the object. Similarly, the *softAR* project by Punpongsanon et al. (2015) manipulates how soft a physical object feels when a user is pushing it. Here, this is achieved by means of spatial AR: a projection changes the surface appearance and alters how deformed the object looks as well as changes the color of the finger of the user. According to the authors, the augmented object can feel

significantly softer than it actually is.

In addition to studies that focus on the perceived temperature and softness of an object, there is also research that focuses on the texture and material of objects. With their research, Iesaki et al. (2008) address the question of how we tactually experience an object when we touch one kind of material while it looks as if we were touching another type of material. Their study uses an HDM to change the visual appearance of physical objects created from geometrical data using rapid prototyping techniques. If viewed through the HMD, a plastic physical object might, for instance, look as if an object were made of wood, cloth, leather, stone or steel. In their experiment, participants were presented with pairs of such visually augmented objects and subsequently identified which of two objects felt rougher. Reportedly, although the compared objects had the same actual roughness, participants felt a difference between them. Hence, the authors conclude that tactual experiences can be deliberately altered by means of visual stimulation. However, they point out that such an influence of visual stimulation on the tactual experience was only perceived when the roughness of the virtual texture and the tactile texture of the physical prototype was almost the same.

Furthermore, Omosako et al. (2012) have created a similar study, but with a focus on changing the perceived center-of-gravity of an object by changing its visual appearance. In order to evaluate whether the perceived center-of-gravity can be affected by superimposing virtual objects, they conducted two experiments. In their first experiment, they superimposed virtual cases of different sizes and aspect ratios onto an actual physical plastic case. Subsequently, participants reported where they perceived the center-of-gravity of the object. In their second experiment, the same virtual object was repeatedly superimposed onto the plastic case, which was filled with different weights. Again, participants reported the location of the perceived center-of-gravity. Based on the results, the authors confirmed that changing the visual appearance of an object indeed can change the perceived center-of-gravity of the object.

Other projects that change our haptic experience of real objects are the previously discussed example of the Gravity Grabber by Minamizawa, Fukamachi, et al. (2007), which can make an empty glass feel as if it were filled with water, as well as Sekiguchi et al. (2005)'s Ubiquitous Haptic Device, which makes it feel as if a box contained a small virtual object. However, these projects focus on communicating the presence of additional elements inside of a real object. This means that here, the haptic transformation is not the goal in itself.

Whereas some projects transform the feel of distinct objects, other projects transform the environment more generally. An example is the Gilded Gait system by Takeuchi (2010), which seemingly changes the environment's ground. This system comes in the form of insoles

that can be placed in existing shoes. The insoles are equipped with embedded actuators that can provide vibrotactile feedback. When the person wearing the insoles makes a step, the insoles simulate different ground textures, such as soft ground or a bumpy ground.

In addition to changing the tactile quality of physical objects, there is quite some interest in changing the tactile qualities of graphical user interface (GUI) elements on touchscreens. For instance, Poupyrev and Maruyama (2003) have proposed a system that can be used to augment and transform the feel of interface elements such as buttons, scroll bars and menus on small touchscreens. For instance, touching a button results in a click under the user's finger. One can argue that such transformations change the way *virtual* (on-screen) objects feel. On the other hand, one can argue that they change the feel of a real touch-screen. In any case, when touching these augmented interface elements, virtual and real tactile stimuli mix in with each other, transforming the original real tactile experience.[24]

In our opinion, projects that change how the real world feels by means of virtual stimuli should be considered part of AR in the broader sense, as virtual content is experienced in relation to (and as part of) the real world.

[24] This approach has, e.g., been pursued by Apple with their so-called 'Taptic Engine' that allows users to feel force feedback when interacting with their iPhone.

4.7.4 *Transformations in Olfactory Perception*

In addition to changing our visual, tactile and sonic environment, there also exist possibilities of changing the olfactory qualities of the environment. Typical means to change these properties are air fresheners, which come in a broad variety of scents. We could, for instance, argue that the "Hawaiian Tropical Sunset" air freshener by Air Wick adds a 'virtual' hint of Hawaii to otherwise non-Hawaiian environments. Consequently, one could go as far and consider environments where virtual scents (scents that are synthesized or that do not originate from their original source) change the olfactory characteristics of the real environment AR.[25]

[25] We are aware that few people would agree to such a broad view of AR. One could more strictly define what counts as virtual to exclude such examples.

4.7.5 *Transformations in Gustatory Perception*

In addition to changing how the real world looks, sounds, feels and smells, we can find various attempts at changing the taste and flavor of real food or drinks. In fact, changing the flavor of foods and drinks by means of food additives is extremely common in our everyday lives. Many food additives are artificial and, for instance, simulate the taste of certain real ingredients. For instance, artificial sweeteners simulate the taste of sugar (and consequently, also can be used to *replace* the ingredient). We could argue that here, additional virtual (synthetic) flavors are integrated with real foods and transform the taste experience on a gustatory level, similarly to how virtual projections can mix in with real objects visually. If we follow this argument, foods with

additives can be seen as a form of AR.[26]

Next to the use of food additives, we also can find various attempts at changing the flavor of foods without changing underlying chemical composition. For instance, Nakamura and Miyashita (2011) approach this by stimulating the tongue with electric current. As mentioned above, the resulting sensation is called electric taste and was discovered by Sulzer in 1752 (Bujas, 1971). Nakamura and Miyashita (2011) built on this phenomenon, and propose a system that changes the taste of drinks using two straws that are connected to an electric circuit. Furthermore, they propose a system that changes the taste of food, which makes use of a fork or chopsticks connected to an electric circuit. Based on preliminary experimentation, the authors conclude that it is possible to distinguish tastes using different voltages. However, their ultimate goal is not only to create different taste experiences, but to increase the sensitivity of the taste organ, and allow participants to taste subtle differences they normally cannot perceive. Furthermore, they aim at making previously tasteless aspects of the environment, such as atmospheric CO_2 concentration perceivable. (Projects that intend to allow us to perceive unperceivable aspects of reality are discussed in section 4.8.)

Although the underlying perceptual principles and the technological implementations between the "Augmented Gustation" project by Nakamura and Miyashita (2011) and the above-discussed tactile feedback technology REVEL (Bau and Poupyrev, 2012) certainly differ, the use of electric current to change a food's taste is conceptually similar to the idea of changing the tactile feeling of real objects by injecting an electrical signal into the user's body. We thus might consider "electric taste" as an augmented reality gustatory technology, just like the revel device is considered "an augmented reality (AR) tactile technology" (Bau and Poupyrev, 2012).

Another project that aims at changing flavor without changing the underlying chemical composition has been realized by Narumi, Sato, et al. (2010). The authors approach this by changing how the drink *looks*. In their experiments, the researchers succeed in creating various different taste experiences of the same drink, simply by virtually changing the drink's color. (This change of color is achieved by placing the fluid into a little bag, and then placing it in a glass filled with white-colored water. The color of the water surrounding the actual drink could be altered with an embedded LED that also was placed in the water.)[27]

Finally, the previously mentioned *MetaCookie* headset (Narumi, Nishizaka, et al., 2011b) (see section 3.3.3), aims at changing the flavor of a real plain cookie by changing the visual appearance of a neutral cookie (e.g. making it look like a chocolate flavored, almond or cheese cookie) and by presenting the user with the matching olfactory information. This reportedly can alter the taste of the cookie. As the

[26] Again, we expect that few people would agree to such an encompassing view of AR. As mentioned, the definition of what counts as virtual could be changed to create a more narrow notion of AR.

[27] Of course, this project not only alters the taste but also alters the visual appearance of the drink. Because more than one modality is transformed, it can be considered in the context of "multimodal transformations". Furthermore, because virtual information from one sense (the color) influences how we experience real information that we perceive through another sense (taste), the project demonstrates what we call "cross-modal" transformations. We will discuss multimodal and cross-modal transformations in more detail in subsection 4.7.7.

MetaCookie project (Narumi, Nishizaka, et al., 2011b) demonstrates, there is an intersection between traditional AR and food experiences.[28] However, we believe there is much more to the field of "gustatory AR". In our opinion, all changes of the taste of real food or drinks by means of virtual stimuli can be considered a form of AR in the broader sense.

[28] In their paper, "When AR Meets Food: A Structural Overview of the Research Space on Multi-Facets of Food", Wei et al. (2012), review how AR technologies have been applied to different aspects of food.

4.7.6 *More Transformations*

The projects that we have discussed in this section so far show that virtual information can change how the real world looks, sounds, feels, tastes and smells. However, there is more to our experience of the world than visual, auditory, tactile, olfactory and gustatory qualities. For instance, we experience the temperature of our surroundings and gravitational forces. What is more, we also experience the passage of time, even though we do not have a specific sensory organ to do so. So far, these kinds of experiences have received little attention in the context of AR. Yet, existing research indicates that virtual stimuli can also target other senses, and as a result seemingly transform even more aspects of the real world.

For instance, informal self-experimentation by Ruhl (2013) has revealed possibilities for transforming the experienced resistance of our surrounding space by means of galvanic vestibular stimulation (GVS). GVS refers to the electric stimulation of the human vestibular system, which plays a key role in our perception of balance. In his experiments, the author used a self-built head-mounted (bilateral bipolar) GVS device in combination with an accelerometer, which was used to measure the orientation of the device (and likewise, the orientation of the wearer of the device). The author went on to explore everyday activities wearing this device, and conducted little experiments, such as using different stimulation intensities based on the researcher's own orientation (Ruhl and Lamers, 2011). This, for instance, did allow him to counteract or amplify his angular movement.[29] As he reports, this revealed potential for AR applications based on vestibular stimulation: When the device was counteracting his movement, it felt to the author as if he "was moving through a liquid or a thick syrup-like medium" (Ruhl, 2013, p. 27). He furthermore reports: "The GVS device counteracted all my movements, so it took more effort to move around" (p. 27). In contrast, when the device amplified his movement, the author reports: "it felt like my resistance was really low since the device backed up every movement I made" (p. 27). Based on his experiences with the device, the author concludes that GVS might be used to simulate "the suggestion of being in a different medium than air" (Ruhl and Lamers, 2011, p. 2). Of course, this self-experimentation is only one of the very first steps towards GVS for AR. However, the more general field of using GVS for altered experiences seems to ad-

[29] In the article from 2013, the author speaks of the device counteracting his balance and counteracting his movement. However, based on the overall description, we interpret this to mean that the device counteracts his angular movement.

vance quickly. For instance, in 2016 the company Samsung revealed an experimental GVS-based headset, which intends to make users experience movement in VR environments (see, e.g., Newsroom (2016) and Engadget (2016)).

Other research efforts that go beyond what we can see, hear, touch, smell or taste, focus on the experienced *temperature* of an environment. Several studies have investigated whether colored light affects our perception of temperature (see, e.g., Van Hoof et al., 2010)—unfortunately with different outcomes.[30] It is commonly assumed that environments with dominant wavelengths toward the red end of the visual spectrum feel warmer and that environments with wavelengths predominantly toward the blue feel colder, which is also called the "hue-heat" hypothesis (Bennett and Rey, 1972). For instance, Winzen et al. (2014), who studied the influence of colored light on the perceived room temperature in an aircraft cabin, found that the temperature in the cabin was experienced differently under different lighting conditions. With yellow lighting, the room temperature was experienced to be warmer than with blue lightning. (Interestingly, the air quality was experienced as being higher in blue light.) Similarly, Fanger et al. (1977) found that participants in their study preferred a slightly lower (0.4 °C) temperature when exposed to extreme red light as opposed to during exposure to extreme blue light. However, the authors concluded that this effect is "so small that it has hardly any practical significance." (p. 11). In contrast, an earlier study by Berry (1961) did not reveal an effect of colored illumination on thermal comfort.

[30] We will discuss cases like this, where virtual information from one sense influences how we experience real information that we perceive through another sense in more detail in subsection 4.7.7.

Of course, colored light in itself is hardly something virtual. Yet, we believe it makes sense to have a look at the effect of colored light onto a participant's temperature perception of the environment. This is because we could use AR technologies to introduce virtual light sources that seemingly change the color of the environment. This could, e.g., happen on an individual level, supporting individual temperature preferences. However, judging from existing studies, additional research is needed to see whether virtually changing the color of a person's surroundings might allow us to affect the perceived temperature of the environment.

In addition to studies that aim at altering the perceived qualities of the air around us (e.g., by making it feel more syrup-like, or making it feel colder/warmer), there are studies that focus on altering our experience of time. Strictly speaking, it is debatable whether we should treat time as a characteristic of the real world that we perceive (e.g., Schäfer et al. (2013) present it as a characteristic of our mental representations of objects and events instead). In this sense, it is debatable whether changing how we perceive time by means of virtual stimuli would fall within the scope of AR. However, let us assume that we experience the passage of time similarly to how we experience sensations of the external world. This raises the question whether virtual

stimuli can alter our time experience, just like they can transform real-world sensations. We believe this might be the case, and actually quite commonly occur in our everyday lives. For instance, we assume that seeing a virtual clock (see figure 4.8) can affect our experience of time passing.[31] One might, e.g., look at the current time and as a consequence, feel like time flies by or as if time stands still. Similarly, one might experience the passage of time differently, if the clock showed another time instead or if there were not clock available at all. We believe that a virtual clock (and possibly, real clocks as well) can be seen as form AR in the broader sense, as they provide an additional layer of information that is typically experienced in relation to the real world.

[31] As discussed earlier, the terms virtual and real are somewhat inappropriate when talking about facts, knowledge or information such as information about the current time. In the end, it likely does not matter whether the time is presented by a virtual clock or by real clock, as the current time itself is neither something real or something virtual.

Figure 4.8: A virtual clock informs customers of the Gusto Esporessobar in Winterswijk about the time. Image by Hanna Schrraffenberger.

In our opinion, it would be exciting to further research how far and in which ways virtual stimuli indeed can be used as a means to alter

our experience of time in a deliberate way. For instance, it would be interesting to know whether specific virtual stimuli can make it seem as if time went slower, faster or maybe even move in reverse.

4.7.7 *Transformations in Multimodal Perception and Cross-Modal Transformations*

As we have shown, there are many ways to change the perceived quality of real objects. Among other, there are projects where visual information changes how an object looks, where sonic information changes how the environment sounds and where tactile information changes how the real world feels. However, we have also seen other approaches to virtually transforming the real. In particular, we have seen many examples where *visual* information changes how an object *feels*, and, for instance, alters the perceived temperature (Ho et al., 2014), texture (Iesaki et al., 2008), softness (Hirano et al., 2011; Sano et al., 2013), or center-of-gravity (Omosako et al., 2012). Furthermore, we have seen examples where the color of a drink alters its taste (Narumi, Sato, et al., 2010) as well as project where a combination of smell and visual overlays is used to alter the taste of a real cookie (Narumi, Kajinami, et al., 2010a; Narumi, Nishizaka, et al., 2011b). There are two interesting aspects of these projects that we have not discussed yet.

First of all, these projects seemingly transform more than one type of sensory stimulus. For instance, the *MetaCookie* project alters what we see, changes the smell of the cookie (or adds a smell) and ultimately, also changes how the cookie tastes. Likewise, other projects change the visual appearance of an object, but by doing so, also affect tactile qualities, such as roughness or softness. As such, these and similar examples can be understood in the context of *multimodal transformation*: multiple sensory modalities of a real object are transformed.

Secondly, the projects all build on our brain's capability of integrating different sensory stimuli. More specifically, in all above-summarized examples, virtual information from one sense influences how we experience real information that we perceive through another sense. For instance, visual information affects what we feel or taste. Such influences where information from one sense affects how we experience information from another sense are also referred to as cross-modal effects and cross-modal interactions.

Cross-modal effects are usually studied in the context of multimodal perception and multi-sensory integration. Multi-sensory integration is concerned with how information from our different senses is combined into one coherent, seamless experience of the world.[32] Cross-modal interactions can occur between different types of real stimuli (e.g., between a real visual stimulus and a real auditory stimulus). However, as we have seen, they can also occur between virtual and real stimuli. As the reviewed projects show, it is possible to make

[32] A comprehensive overview of research in the field of multimodal perception and cross-modal effects is provided by Bertelson and De Gelder (2004).

use of this fact and deliberately utilize *cross-modal* relationships between different types of virtual and real stimuli in order to create and shape AR experiences.

4.7.8 Short Summary Altered Reality

To briefly summarize this section, virtual content can be used to alter how the real world in general, as well as real objects in particular, appear (look, feel, smell, taste, sound) to us. Simply put, the virtual can transform the real. We call this form of AR *altered reality*.

4.8 Extended Perception: Translation-Based Relationships

When we think about the real world, we typically think about things we can see, touch, hear, smell or taste and more generally, the things we can perceive. At the same time, we unconsciously exclude aspects of our reality that we cannot perceive, such as ultrasound and magnetism. Fortunately, there are devices that help us to overcome some of those sensory limitations and that allow us to perceive things about the environment we normally cannot perceive. In this thesis, we refer to this process as *perceptualization*. The term perceptualization was coined by the Media Technology Master program for a course in the program's curriculum. As mentioned on the course website, perceptualization is a generalization of the terms "visualization" and "sonification" that applies to all human senses. It "describes the translation of signals and information to modalities that appeal to any of the human senses" (Media Technology MSc Programme - Leiden University, n.d.). In this thesis, we adopt the same understanding and usage of the term. It is important to note that perceptualization can occur in many contexts. For instance, perceptualization can also be used for the exploration and communication of datasets and translate values into something we can, e.g., feel or hear. Here, we look at perceptualization in a real-time and real-world context. In other words, unperceivable signals from the real world are translated into signals we can perceive with our senses—the real is translated into something virtual.

Devices that perceptualize unperceivable signals are rather common in our everyday world and have existed for a long time. A well-known example is a hand-held Geiger counter, which translates the amount of radiation that is present at the current location into audible clicks. Another common device that translates information we cannot perceive to stimuli we can perceive are night vision goggles, which allow a person to see in the dark.

The idea of perceptualization relates to AR in the sense that virtual stimuli can be used to represent real but unperceivable aspects of the real world. For instance, virtual imagery can visualize the magnetic field, and virtual soundscapes can allow us to perceive air pollution.

In fact, one can argue that perceptualization always transforms real information that we cannot perceive into virtual (synthetic, artificial, generated) information that we can perceive. As such, perceptualization falls in the scope of AR in the broader sense. Because real-time and real-world perceptualization projects are concerned with what we can perceive, rather than with extending the environment, we refer to this group of projects under the umbrella *extended perception*.

The idea of translating signals we cannot perceive to signals that we can perceive is closely linked to fields of sensory substitution and sensory augmentation. Sensory substitution refers to cases where one human sense (e.g., touch) is used to acquire information that is normally acquired by a different sense (e.g. vision) (Kaczmarek, 1995). Sensory substitution systems often aim at allowing people to perceive information they cannot perceive due to an impairment. For instance, sensory substitution systems have been proposed to allow blind people to see via their ears or via their skin receptors (see, e.g., Bach-y-Rita and Kercel, 2003). An example is "the vOIce" by P. B. L. Meijer (n.d.). This device for the blind translates a live camera view into sound. The images are scanned from left to right. Pixels higher in the image are mapped to higher frequencies (the pixel's position on the y-axis determines the pitch) and brighter pixels are louder (the brightness is mapped to volume). A study by Auvray et al. (2007) shows that blindfolded participants could use the device to localize and point at a target and to recognize objects and discriminate between objects of the same category. [33]

[33] See Ward and P. Meijer (2010) for a description about the visual experiences of two blind users who have been using the vOIce over a period of years.

Like sensory substitution devices, sensory augmentation devices translate information that cannot be perceived into stimuli that can be perceived. However, they aim at allowing us to perceive information humans in general cannot perceive due to the way our senses work. They aim at extending our sensory abilities so that we can perceive additional and 'new' aspects of the environment. In other words, they hope to provide us with an additional sense and new sensory experiences.

An example of a sensory augmentation device is the vibrotactile magnetic compass belt called *feelSpace* (Nagel et al., 2005). The belt is worn on the waist and indicates the direction of magnetic north with vibrations. Reportedly, none of the participants in Nagel et al.'s study experienced a local magnetic field. However, two participants, after wearing the belt for a longer period of time, experienced the input from the belt as a property of the environment rather than as mere tactile stimulation. Similarly, a follow-up study by Kaspar et al. (2014) concludes that the feelSpace device led to subjective changes in space perception and enabled the use of new navigation strategies. (In this study, eight out of nine belt wearing participants agreed with the statement that they are developing "a new sense of spatial perception with the belt/with training" after using the belt for several weeks. For

instance, one participant describes the following: "Often I do not perceive the vibration any more. It is rather a direct feeling of knowledge – not even really a perception. It does not feel like any other sense" (belt wearing participant 3, p. 54).

Another project that aims at augmenting our perception is the previously mentioned 'Augmented Gustation' project by Nakamura and Miyashita (2011). As discussed, their project stimulates the tongue with electric current. The authors hope that this will allow participants to perceive tasteless properties of the real environment, such as CO_2 concentration, as well as increase the sensitivity of the taste organ so that humans can distinguish among tastes that they normally cannot discern.

Perceptualization also has interesting overlaps with the field of traditional AR. Established AR technologies and concepts, such as the visual integration of virtual objects into our view can be used to translate what is hidden from our senses into something we can perceive. For example, AR applications allow for a form of virtual X-ray vision and make it possible to see hidden or occluded objects (see, e.g., Bane and Hollerer, 2004). Furthermore, the mobile AR platform Layar has been used to visualize the air quality in the Dutch city Leiden in the context of the MIMAQ (Mobile Individual Measurements of Air Quality) project (iReport, 2010). More specifically, virtual clouds were used to represent the air quality/pollution. Judging from the image that can be found of this project online (see iReport (2010)), these virtual clouds (more or less) appeared to float in the real space, when the environment was viewed through the application.

Interestingly, the general field of AR is often seen as a form of augmented perception. For example, Normand, Servières, and Moreau (2012) point out: "Reality can not be increased but its perceptions can. We will however keep the term 'Augmented Reality' even if we understand it as an 'increased perception of reality' " (p. 1). Similarly, Ross (2005) refers to AR as that "what should be called augmented perception of time and space" (p. 32). Furthermore, the widespread survey of AR by Azuma (1997) states that AR enhances a user's perception of and interaction with the real world. In contrast to these views, we treat augmented perception as a subset of AR that is explicitly focused on allowing humans to perceive more about their surroundings. We see extended perception as a form of AR because—no matter whether we are dealing with a mobile app that displays virtual clouds (iReport, 2010), a Geiger counter that presents us with audible clicks, night vision goggles or other sensory augmentation systems such as the compass belt (Kaspar et al., 2014; Nagel et al., 2005)— the additionally provided information relates to the surrounding environment.

To summarize, we can translate unperceivable but real aspects of the environment into virtual but perceivable information. In such cases, the link between the virtual and the real is a *mapping* or *translation* from

something humans cannot perceive to something we can perceive. As such, the virtual can augment our perception of the real world. This augmentation always also informs us about the real environment. In this sense, perceptualization always goes hand in hand with content-based relationships cf. section 4.3). In addition, the information can also appear present in the environment and relate to the surroundings spatially (cf. section 4.2).

4.9 *Physical Relationships: The Virtual and the Real Affect Each Other*

Real objects have physical qualities such as a mass and temperature, and consequently are affected by physical laws such as gravity. In contrast, virtual objects do have virtual/simulated qualities and do not have to follow physical laws. The fact that we can see virtual objects in space does not necessarily mean that they appear to exist in a physical or material form or that they adhere to physical laws. Consider, for instance, the previously discussed project by Feiner, Macintyre, et al. (1993), which presents line-based illustrations that help with the maintenance of an office printer. Aside from their color, these lines do not seem to have any physical (material) properties. They appear as if it existed in 3D space, however, unlike physical objects, they are not affected by gravity or cast shadows. Judging from their appearance, we would not expect them to offer any resistance when we try to touch them. Simply put, they appear to be present spatially, but not in a physical or material form. Similarly, we can easily imagine virtual ghosts that do not obey to physical laws and that move through walls and hover over ground. As these examples suggest, virtual content can appear to be part of and present in real space without displaying traditional *physical* qualities. Yet, more commonly than not, virtual objects also simulate some physical qualities and seem to relate to the real world physically.

This (simulated) physical relationship between the virtual can take many forms. For instance, virtual and real objects can affect each other on an optical or acoustic level (e.g., casting shadows or causing resonances). Dynamic (movement-related) effects are also possible, for instance, if virtual and real objects collide. The effects furthermore can have different *directions*. For one, the real world can affect the virtual content physically. Furthermore, the virtual content can seem to physically affect the real world. What is more, the two can influence one another and *interact*. In this section, we will explore such (simulated) physical relationships between the virtual and the real. Unlike in previous sections, we explore both how the virtual affects the real as well as focus on how the real affects the virtual.

4.9.1 *The Real World Affects Virtual Content*

There are many ways in which the real world can seemingly affect a virtual object physically. First of all, there is quite some research that focuses on *optical effects*, such as occlusions, reflections and refractions. For instance, Madsen et al. (2006) present a method for taking the illumination of the real world into account when rendering virtual objects. As a result, light changes in the real environment affect the appearance of virtual objects, making sure they are shaded realistically as well as that they cast fitting shadows. Furthermore, Kán and Kaufmann (2012) focus on rendering and displaying realistic reflections and refraction of the real world in virtual objects. They demonstrate their rendering system with a virtual glass, that shows correct refractions/reflections of surrounding elements such as a person's hand and physical colored cubes that stand next to the virtual glass (see figure 4.9).

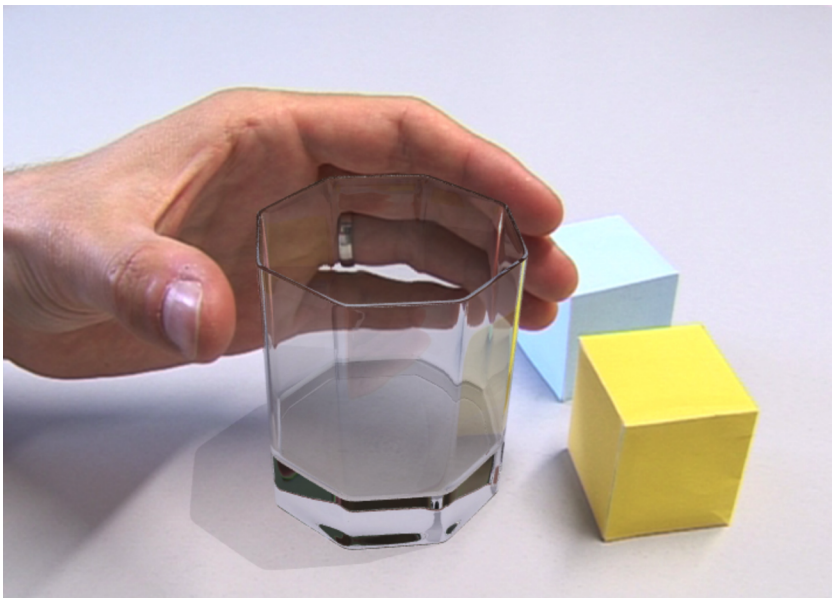

Figure 4.9: The real world affects the appearance of the virtual glass. For instance, we can see the person's hand refracted in the glass. Reprinted from P. Kán and H. Kaufmann (2012). "High-quality reflections, refractions, and caustics in augmented reality and their contribution to visual coherence". In: *International Symposium on Mixed and Augmented Reality (ISMAR 2012)*. IEEE, pp. 99–108. Reprinted under fair use.

Similarly, Pessoa et al. (2010) also propose a photorealistic rendering technique that focuses on the effects of the real environment on the appearance of virtual objects. Their demonstrations include, for instance, a virtual vase that appears to be illuminated by the real environment. (This is achieved by virtual light sources that manually were positioned to mimic the position of the most prominent lights in the real environment.) Furthermore, they, e.g., show a teapot reflecting surrounding physical objects as well as color bleeding effects where light from real surfaces appears to color virtual objects.

In the previous examples, the surrounding real world has an effect on the *optical* appearance of virtual objects. However, influences of the real world on virtual content are not limited to the visual domain. We can easily imagine virtual objects that either seem to be or that actu-

ally are affected by the acoustic properties of their real surroundings. If for instance, the sound of a virtual object is played back in real space by means of loud-speakers, it will naturally be affected by the properties surrounding space. It will, e.g., sound different when played back in a church as opposed to on the streets. A similar effect can be simulated by means of audio effects (e.g., reverb) even if the sound is played back by means of headphones rather than loudspeakers. In other words, we can make it sound as if a virtual sound were reflected in and affected by the surrounding physical structures. Of course, just like there are many possible visual effects, acoustic effects are not restricted to reverb. For instance, virtual object could start to resonate due to a real sound that occurs at their resonant frequency. However, so far, such acoustic influences have received rather little attention in AR research.[34]

[34] One exception is the research by Lindeman and Noma (2007). The authors argue that computer-generated stimuli generally should undergo the same transformations as real-world stimuli. They state that virtual characters should receive "the same lighting effects (light position and intensity) as objects in the real world" but emphasize that this holds for all senses and point out that "the voice of a virtual character should also be influenced by environmental objects, such as occluders or reflectors" (p. 175).

Aside from these influences that affect a virtual objects' (visual and non-visual) appearance, the real world can also affect a virtual object's movement and/or position. In other words, there are also possibilities for *dynamic* influences and interactions.[35] A common real-world force that often seems to affect virtual objects is gravity. Many virtual objects seem to have a physical mass and seem to be affected by gravitational forces of the real world. At least, this interpretation seems natural, given that virtual objects often appear to stand, lie or move on real objects rather than float around in space freely. To mention just few examples, virtual Pokémon game characters appear to sit on the real pavement, virtual spiders clamber over real obstacles and can be carried by participant's hand's (Corbett-Davies, Dünser, and Clark, 2012) and virtual architectural models appear to stand on top of real tables (Broll et al., 2004)—all of which implies that the virtual objects are affected by gravity.[36]

[35] Breen et al. (1996) distinguish between visual and physical forms of interaction. In contrast, we summarize both forms in the context of physical interactions.

[36] Strictly speaking the virtual architectural models by Broll et al. (2004) are actually affected by gravity in the sense that they are linked to real, physical placeholder objects that are placed on the table top.

An early research project that focuses on gravity, kinematic constraints and collisions between virtual and real objects has been realized by Breen et al. (1996). In their paper, they present AR techniques to automatically move virtual objects downwards in the real environment until they collide with real objects in the environment. As the authors mention, this process can be viewed as "simulating virtual 'gravity"' (p. 11). Chae and Ko (2008) similarly simulate gravity, but take things a step further with the use of a dedicated physics engine that applies physical attributes such as weight, gravity, friction, elasticity and force. The virtual ball in their demonstration not only falls downwards until it collides with a real object but also bounces off this object. Furthermore, in their setup, the angle of the floor determines the resulting motion of the virtual object.

In addition to gravity, other types of physical forces can affect the movement of virtual objects. This happens, for instance, in the case of van Velthoven's (2011) interactive installation and car racing game *Room Racers*. Unlike traditional computer games, which are displayed

on a screen, this game takes place in real space. Virtual cars are projected onto the player's floor. Real objects, such as shoes, keys and toys are placed on the ground and define the racing course. During the game, players can steer the cars around the track with traditional game consoles. The physical objects act as barriers that cannot be crossed or passed through by the virtual cars. Furthermore, S. Kim et al. (2011) provide another example of how real objects can affect the movement of virtual ones. In their setup, the collision between a real ping-pong racket and virtual spheres (essentially virtual balls) and boxes results in what they call "feasible responses" (p. 26): the objects seem to bounce off the racket in a plausible way. Although the authors do not discuss this explicitly, it appears that one can use the real racket to play with the virtual objects similarly to how one would play with real objects. However, while the racket affects the movement of the virtual objects, the collision does not affect the movement of the physical racket in return (and no impact will be felt by the participant holding the racket). This raises the question whether virtual content also can affect the real world.

4.9.2 *Virtual Content Affects the Real World*

Aside from projects and situations where the real affects the virtual, we can also find cases where the virtual affects the real. As the virtual often has no way of actually affecting the real world, these effects often are simulated.

Like in the previous section, optical effects play an important role when it comes to influences between the virtual and the real. Typically research into illumination in AR (see above) not only discusses how the real world affects virtual objects but at the same time also addresses how virtual objects can influence the real world. For instance, virtual objects can seemingly affect the appearance of the real world by casting virtual shadows on the real world. An example is the system by Madsen et al. (2006), which not only realizes realistic lighting of virtual objects (including shadows that real objects cast on virtual objects) but also makes sure virtual objects cast shadows onto the real environment. Similarly, Sugano et al. (2003) explore what effect shadows of virtual objects have on AR. Based on experiments, the authors conclude that shadows increase the presence of virtual objects as they provide a stronger link between the virtual object and the real world. Of course, optical effects are not restricted to shadows. For instance, we might also expect to see reflections of virtual objects in real objects. We can, e.g., easily imagine scenarios where a virtual character should appear in a real mirror. However, while the reflection of the real world in real objects is commonly addressed, we can find little research dedicated to the reflection of virtual objects in the real world.[37]

As one might expect, effects of virtual objects on the real world are

[37] One rather specific exception is the research by Bimber, Encamacao, et al. (2000). This work addresses the reflection of *stereoscopically projected* virtual scenes in a mirror and explores the idea of using a mirror as a means to look at and interact with the virtual information from otherwise difficult-to-reach positions.

limited to the visual domain. We can, for instance, imagine a virtual singer, whose voice causes a real object to resonate. This can either happen virtually (by simulating the resonance) or actually (if the virtual sound is played back by a loudspeaker and thus causes the resonance).

As we have seen above, real objects can collide with and thereby affect the movement of virtual objects. The opposite—virtual objects affecting the movement of real objects—is much more difficult to realize. This is because virtual cannot directly apply forces to real objects. So far, little research has been invested in realizing such physical effects. One of the few projects that address this challenge has been realized Kang and Woo (2011). In their *ARMate* project, they extend a physical toy cart with electronics so that a virtual character can push and pull the cart.

Other situations in which virtual objects can affect physical objects arise when the virtual object has a physical counterpart. This is, e.g., the case with the virtual toy beaver Sphero, which is physically represented by a robot ball. If the beaver/the robot ball collides with another physical object, such as a football, this collision will naturally have some sort of effect. However, aside from the discussed examples, it remains rather unclear in what ways and to what extent virtual objects can (appear to) affect real objects physically. We will explore this question in more depth in the following chapter.

4.9.3 *Interaction Between the Virtual and the Real*

So far, we have discussed the possibilities of virtual content affecting the real world and the real world affecting virtual elements. If we combine these possibilities, it is easy to imagine scenarios in which the virtual and the real affect each other, or in other words, *interact*. For instance, we can envision the collision of a virtual ball and a real ball, that would cause both balls to change their path. However, as we have seen, it is rather difficult for virtual objects to affect real objects. As a consequence, there are only few examples of projects in which the virtual and the real influence one another physically.

An artwork which demonstrates that real and virtual elements in an environment can physically interact with each other is *Radioscape* by Edwin van der Heide (2012; 2000-). This art installation makes use of several radio transmitters that are distributed over a part of a city, each transmitting one layer of a meta-composition. By navigating through the city with a custom developed receiver, a listener can pick up several signals at a time. The volume of the single layers depends on one's distance to the corresponding transmitters. Due to the chosen wavelength, buildings become conductors and resonators for the transmitted signals. The physical environment is excited by and responds to the transmitted radio waves. As such, they influence the

waves in return. Ultimately, this causes the physical environment to affect what one hears.

Just like we wonder in what forms and to what degree the virtual can (seemingly) affect the real physically, we wonder to what extent and in what ways interactions between virtual and real objects are possible to produce in AR. One the one hand, it would be interesting to know if we can reproduce real-world interactions such as collisions. On the other hand, it might be even more interesting to explore whether other and new types of interactions might be possible—after all, virtual objects do not have to adhere to the same laws as real objects. We will address this question in more depth in the following chapter.

4.9.4 Short Summary Physical Relationships

The presence of virtual content in real space opens up possibilities for interactions between the virtual and the real. We have shown that there are many ways in which the real world can seemingly affect a virtual object physically. In contrast, it is more difficult for the virtual to affect the real world. Yet, such influences can be realized and simulated. The virtual and the real can also influence one another and interact. This possibility will be explored further in chapter 5.

4.10 Behavioral Relationships: The Virtual and the Real Sense and React to Each Other

If we look at the real world, physical interaction between the elements in the real world is only one of various forms of interaction that occurs. For instance, people also interact on a behavioral level. Imagine e.g., people talking to each other or reacting to each other's movement to avoid collisions on a crowded street. Furthermore, animals react to one another on non-physical levels. An example would be dogs barking at each other or chasing one another in a park. What is more, interactive objects also react to and interact with the environment. Consider, e.g. interactive doors that sense the area in front of them and automatically open when people approach them.

Of course, people and doors differ considerably. Yet, the described actions interactions have something fundamental in common: All of them are based on information sensed in the surrounding environment. In the case of people and animals, information is obtained by means of senses. In the case of objects, the information about the environment is acquired by means of sensors. In both cases, the sensed information ultimately prompts some sort of response or action. (Typically, this response might elicit a change in the environment in return, resulting in a chain of cause and effect or in other words: interaction).

Although virtual objects have no real senses, they can nonetheless

sense the world by means of sensors and act in this world based on what they sense. Hence, behavioral relationships can also be established between virtual objects and the environment. For instance, a virtual trainer and a real runner might race against each other on the running track and react to each other's movement. Likewise, a virtual animal might react to alluring sounds and a virtual car might avoid colliding with real objects. In this section, we discuss such relationships between the virtual and the real that are based on either sensory input or sensor input under the term "behavioral relationships". This term is chosen because here, virtual and real objects not only (seem to) exist in space, but also exhibit some kind of behavior that relates to the environment.

Just like physical relationships, behavioral relationships between the virtual and the real can take different forms. First of all, the real can sense the virtual and change its behavior based on the sensed information. Secondly, virtual objects can sense the real world around them and act according to the acquired information. Finally, the virtual and the real can sense each other, react to each other and ultimately, react to each other's reactions - resulting in interaction on a behavioral level.

Cases where the real senses the virtual and changes its behavior based on the sensed information are quite common: Participants typically react in some way to the virtual content they perceive in the world. For instance, participants often see a virtual object and consequently move around in the space to have a look at the virtual element from different perspectives. One can argue that here, the real world (a real participant) reacts to the presence of virtual objects on a behavioral level. However, with the exception of participants, the real world seldomly reacts to virtual additions on a behavioral level. For instance, real doors typically do not open for virtual creatures (although this could be realized on a technological level) and pedestrians typically walk right through virtual elements (such as virtual Pokémon), simply because these elements are not part of their perception of the world.

Just like real elements rarely react to virtual elements in an augmented space, virtual elements only occasionally sense and react to real elements in the space. It is often apparent that virtual animals or creatures are not able to sense their immediate surroundings. An example are the previously discussed virtual Pokémon . These virtual creatures appear to exist in front of and face the player but at the same time, have literally little sense about what is going on around them. They seem rather oblivious to their surroundings. This also reflects in the limited ways we can interact with virtual creatures: We cannot scare them with sudden noises or lure them closer with the smell of real food. Judging from personal experience, the virtual creatures are not affected by humans making faces at them. Considering their rather apathetic attitude towards sounds, smells, or even visual occurrences

in their surroundings, it can quickly become apparent that essentially, they cannot see, hear, smell or otherwise sense the world around them.

Yet, the idea of virtual objects sensing the environment is not new. Many virtual objects and characters exhibit some kind of geometric awareness of their surroundings. For instance, in the AR version of the game Quake (Piekarski and Thomas, 2002), virtual monsters appear to walk around the real campus. Although registration issues cause monsters to seemingly walk through walls or appear out of nothing (Piekarski and Thomas, 2002), the fact that they walk around in the environment presumably causes the impression that they can sense the surroundings to some (at least geometrical) degree. Furthermore, as illustrate, during the game play, virtual monsters attack both each other and the player. We assume this creates the impression that the virtual monsters, in fact, can see the player as well as each other. In other words, the virtual monsters seem aware both of virtual as well as real elements in the environment. (Unfortunately, it remains unclear if the monsters also can sense bystanders and whether they are aware of the actions of the player that they should be able to 'see' from their perspective.) Arguably, the battle between the real player and the virtual monsters can be seen as a form of behavioral interaction between the virtual and the real. We believe such interactions can be taken to the next level by also incorporating interactions between virtual objects and the general surroundings. For instance, in the ARQuake game, the virtual monsters could be able to avoid collisions with real people in the environment or recognize real doors to seemingly enter and hide in real buildings.

The idea of virtual objects being aware of the topography of the environment as well as of a participant's position in this space also comes back in other AR games. For instance, the sound-based AR version of PacMan (Chatzidimitris et al., 2016) makes use of ghosts that chase the player. The ghosts move through the actual streets and try to catch the player. (Three of the ghosts move randomly whereas one of them actually takes the player's position into account). The ghosts are clearly aware of the streets, as their movement through the space always follows existing real-world paths. In this sense, the behavior of the ghosts relates to the surrounding environment on a behavioral level. Presumably, the chasing dynamic between the participant and the ghosts is also experienced as a form of behavioral interaction between the participant and the ghost. With respect to this, it would be interesting to know if the player actually feels like the ghosts can sense them in the space.

Projects where virtual objects also sense non-visual and non-spatial information about their surroundings are sparse. One example is the mobile AR game *GeoBoid* by Lindeman, G. Lee, et al. (2012). In their game, players are surrounded by flocks of virtual geometric creatures called GeoBoids. These creatures are represented both visually as well

as by means of spatialized audio. Players move towards a swarm of GeoBoids by running to their location in the real world. They can capture individual creatures by pointing the device at them and swiping over the screen of their mobile device. However, players can also scare the flock by whistling at a certain pitch and for a certain duration. In other words, the birds seem to be able to listen to their surroundings and act according to what they hear.

The idea of virtual elements sensing and acting in the world relates the field of AR to that of *Intelligent Agents* as well as to the field of *Sentient Computing*. *Intelligent agents* have been defined by Russell et al. (1995) as "anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors" (p. 31). Arguably, if virtual objects sense their surrounding environment and (seemingly) act in and upon this environment, they can be considered virtual intelligent agents. The concept of sensing the environment also comes back in the context of *sentient computing*. Sentient Computing refers to the concept of making applications "more responsive and useful by observing and reacting to the physical world" (Hopper, 1999, p. 1). As Addlesee et al. (2001) explain, sentient computing systems can adapt their behavior based on a model of the surroundings that they create using sensor data. Ultimately, virtual objects that sense the world and react to their surroundings would rely on some sort of system that maintains such a model of the surroundings. So far, AR systems primarily sense the world with respect to geometry, as this is often needed to register them in three-dimensions and to make it seem as if a virtual object existed in the real world. Furthermore, as we have seen, some AR systems take the illumination of the real world into account as well. However, if we want virtual objects to not only *exist* but also *act* and *behave* in the world, the environment has to be sensed and interpreted on additional levels. For instance, for a virtual mouse to react to sounds, it makes sense to use microphones to sense the sound in the surrounding space. If the mouse also should be afraid of real cats, the AR system also has to detect cats and determine whether the cat can sense it from its own perspective.

A project that includes the idea of virtual objects sensing the environment (albeit still to a small extent) is the previously mentioned *ARMate* by Kang and Woo (2011). Here, the virtual character that can push and pull a toy cart also has what the authors call "synthetic vision". This means it can autonomously perceive virtual and real elements in its view. Furthermore, as mentioned in subsection 3.2.2, the idea of virtual creatures being more aware of and reacting to their surroundings has been addressed in AR+ update of the game Pokémon GO (Niantic, Inc., 2017). In this version, Pokémon seemingly sense the player's movement. Because of this, players can scare the virtual creatures away with sudden movements.

In additions to projects that simply incorporate such behavioral re-

lationships between the virtual and the real, there is some research that addresses these possibilities more explicitly. For instance, Barakonyi et al. (2004) have developed a framework called *AR Puppet* that combines the concepts of AR, sentient computing and intelligent agents (more specifically, autonomous and animated agents). The framework builds on the idea that real-world objects such as printers, digital instruments and interactive robots can be both queried for status information and controlled with commands. This opens up possibilities for virtual characters to sense and affect such objects. The authors introduce an example application of a virtual LEGO repairman. This virtual repairman guides the assembly of a real, physical LEGO robot and, e.g., illustrates how to mount the next pieces onto the robot. Although the authors do not describe this, we can easily imagine this repairman to also actually steer/drive the real robot around in the physical environment. In a similar way, virtual characters could play the physical (but digital) piano or cause an automatic door to open when they approach it. As such, the work of Barakonyi et al. (2004) can serve as an important inspiration for behavioral interactions in AR.

In addition, Gelenbe et al. (2005) explicitly address the idea of introducing virtual autonomous agents in AR environments. The authors approach this from the context of training simulations (such as medical or military training) where it is important that simulated entities act autonomously and realistically. The authors point out that "[t]he behavior of injected artificial entities can be as important as their appearance in a visual simulation" (p. 260). In line with this, they address questions such as how virtual objects can be designed to exhibit intelligent behavior in AR settings and propose an agent model that operates under the assumption that the virtual agents perform "outdoor" missions in an environment that only contains little obstacles and enemies. Although rather specific, their research shows ways in which AI (artificial intelligence) research and in particular work on multi-agent systems can inform and potentially advance the field of AR.

Given that—with the exception of the real participants—real elements in the world rarely sense virtual additions there is currently little ground for behavioral interactions between the virtual and the real (aside from interactions between the participant and the virtual content). We see possibilities for advancing AR in this area. For instance, we can easily imagine scenarios where virtual birds sing along and interact with real birds, where virtual characters interact and play with a real automatically closing door or where virtual and real toys interact with one another.

To summarize this section, the real world can relate to the virtual world on a behavioral level. At the same time, virtual content can relate to the real world on a behavioral level. Furthermore, the virtual and the real can sense each other, react to each other and react to

each other's reactions. We refer to this chain of action and response as behavioral interaction. Currently, virtual elements commonly seem aware of the geometry or topology of their surroundings. We believe that there is plenty of room to extend their 'senses' and further explore behavioral relationships in AR. For instance, we see much potential in making virtual objects react to multimodal properties of the real world. We believe strengthening the relationships between the virtual and the real on a behavioral level will have two main benefits. First of all, we believe that making virtual objects react to the multimodal properties of the real world can help to convince us that they are part of this world. Imagine, for instance, a virtual pet that gets scared when there is a sudden sound in the surroundings, a virtual object that dances to the song playing on the radio, or a virtual character that puts on different clothes, according to the current temperature. Presumably, such relationships will strengthen and contribute to the illusion of virtual objects existing in and being a part of the otherwise real environment. We expect that, if virtual content matches the multimodal properties of the real world, the virtual might blend in with the real world more seamlessly, ultimately enabling more holistic experiences.

Second, behavioral relationships provide many possibilities to entertain and engage participants. If, for instance, a virtual creature senses the world, a participant might lure it closer with certain sounds, change their appearance by placing them in a colder environment or by turning on the heat, or affecting their behavior by putting on a different song or by shedding light on them with a torch.

4.11 *More Relationships*

In the previous sections, we have discussed various relationships between the virtual and the real. Although we believe we have identified those links between the virtual and the real that are fundamental to AR, the presented typology is certainly not *exhaustive*. In this section, we want to emphasize the fact that more relationships exist and briefly discuss some of those relationships, although in less detail.

One relationship we have only mentioned in passing is a musical relationship between the virtual and the real. An example of which would be the relationships between the sounds of a virtual piano that plays along with real instruments. The fact that we have not discussed a relationship does not mean that it cannot play a role in AR. For instance, apps like the above-mentioned RjDj (n.d.) might provide us with virtual sounds that relate to the sounds of the real surroundings harmonically.

Another type of relationship that has not been addressed in detail is a temporal relationship between the virtual and the real. Typically, information about our surroundings informs us about the *"here and now"*. Aside from telling us more about the current characteristics

of our surroundings, the virtual can also inform us about the past and future of the surroundings. In such applications, *temporal* relationships between the virtual and the real play a key role. Examples in which temporal relationship play an important role are, e.g., the previously mentioned "Street Museum" apps (*Museum of London: Streetmuseum* 2014; *Street Museum NL* 2013). As discussed above, these mobile apps display images of the past on the location where they have originally been taken. Of course, this concept is not limited to images. For instance, also sounds can be played back where they were recorded earlier. Likewise, AR applications might show 3D models in real space that suggest how the area will look in the future.

Also, it should not go unmentioned that the virtual and the real can be related on a *narrative* level. For instance, virtual objects might be experienced as part of our environment, simply because a story relates them to the environment. According to the author's experience, this happens in the running application "Zombies, Run!". This app uses narrative to connect the virtual audio story with the player's reality. It presents the player/runner with the sounds of "Zombies" that, according to the story, chase the runner. The sounds of the zombies are not spatialized, and from a perceptual point of view, it is quite obvious that the Zombies are not really present in the same space as the player. Yet, the narrative tells the runner that this is the case, and thus establishes a link between the runner, the surroundings and the Zombies.[38] Based on the personal experience of the author, the Zombies are not actually *perceived* in the surrounding environment, but nonetheless *imagined* in the space.

[38] However, as the running game also describes an environment that typically differs from a runner's actual environment, it remains questionable, whether the Zombies are experienced as a part of the otherwise real environment, or whether the player is transported into another, virtual environment instead. In this sense, the lines between AR and VR blur.

Another type of relationship that might support the impression of virtual content being part of the real environment is *similarity* between virtual content and its real surroundings. The audio artist Janet Cardiff, who creates walks where virtual pre-recorded soundscapes mix in with the actual sounds of the environment (see chapter 3) has emphasized that similarity/imitation is important for the virtual soundscape to mix in with the sounds of the real world. On her website, she explains: "The virtual recorded soundscape has to mimic the real physical one in order to create a new world as a seamless combination of the two." (Cardiff, n.d.). Of course, imitation is not limited to the sonic domain. For instance, the previously discussed project by Fischer et al. (2005) makes use of visual and stylistic similarities between the virtual and the real by applying the same stylization to the participant's view of the real world as well as to the virtual additions that are included in this view. As mentioned, the authors suggest that this process this makes the virtual elements and the real world look very similar, and ultimately, makes it look as if virtual objects were an actual part of the real environment.

The relationships we discover always depend on the chosen perspective. For instance, we might speak of environments where the

virtual *enhances* the real if we were interested in the *quality* of the resulting environment. In our exploration, we have approached AR from a conceptual and experience-focused perspective.

As mentioned, the presented overview is not exhaustive. We expect that many more relationships between the virtual and the real can be discussed, especially if one discusses the relationships on a more granular and detailed level or shifts the perspective—for instance, by approaching AR from a technological perspective or by focusing on the relationships between a *participant* and the virtual content.

4.12 Summary, General Discussion and Conclusion

In augmented reality, virtual and real content are combined in our so-called real world. However, simply presenting or displaying virtual content in the real world arguably is not enough to create AR: in AR environments, the virtual relates to the real world in which it is presented.

Our investigation has shown that the virtual can relate to—and ultimately augment—its real surroundings in many ways. On a fundamental level, virtual content can relate to the real world spatially and content-wise. Furthermore, it can translate unperceivable but real aspects into to a perceivable but virtual form. If the virtual relates to the real on such a fundamental level, it can play various different roles in the real world. First and foremost, it can extend the real and provide additional content to the participant. We suggest summarizing these scenarios under the term *extended reality*. Furthermore, it can hide or seemingly remove real objects from the perception of the participant. This is already known under the term *diminished reality*. In addition, the virtual can transforms the real environment or real objects in the environment. We propose the term *altered reality* to describe this subform of AR. In cases where the real environment is incomplete without the virtual elements, the virtual can furthermore complete the real environment. Our proposed term to single out this form of AR is *hybrid reality*. Furthermore, the presence of virtual objects in real space also opens up possibilities for *physical* as well as *behavioral* relationships between the virtual and the real. Here, it is important that virtual content not only can appear in but also potentially *act* in the real world. It furthermore is important that the virtual not only relates to the real world, but that the real world also relates to and potentially affects the virtual.

As emphasized, many more relationships could be discussed. However, we believe that we have identified the most prominent links between the virtual and the real as well as brought attention to less commonly considered relationships that likewise can shape AR experiences.

It is important to note that the discussed relationships are not mu-

tually exclusive. For instance, virtual information can both appear to exist in the real world *and* inform us about our surroundings.

As the previous chapter has shown, AR is often defined in terms of interactive systems that align virtual content with the real world in 3D and in real-time. This understanding of AR is linked to the desire of making it seem as if virtual objects existed in the real world. This chapter reaffirms our belief that there are many other factors aside from spatial registration that can contribute to the impression of virtual objects being part of the real environment. For instance, whether an object appears present in the real environment, likely also depends on whether this object physically interacts with the real objects, whether it appears to sense and react to the real environment on a behavioral level and whether it relates to the real scene on a content-level. It would be interesting to investigate what factors influence whether we experience virtual content as part of real space systematically with experiments in the future.

At the beginning of this chapter, we have asked ourselves what AR entails if we define AR in terms of relationships between the virtual and the real. Our investigation has revealed that this understanding of AR describes an extremely diverse field. Our definition, for instance, encompasses projects that make use of a variety of different technologies and stimuli as well as projects that focus a wide range of different experiences. Whereas some might question the need for such an encompassing view on AR, this broad picture of AR aligns well with the overall goal of this thesis to address "AR in the broadest sense". To the best of our knowledge, no other equally broad, diverse and comprehensive overview of the different forms of AR exists.

We believe that in order to work and communicate in such a complex field, we have to be able to clearly identify and single out specific forms of AR. Our proposed typology can help with this. In our opinion, it makes sense to distinguish between presence-based AR and content-based AR (however, both can be combined). Furthermore, it often can be helpful to further specify the role of the virtual content in the real world. For this, the distinction between *extended reality*, *diminished reality*, *hybrid reality* and *altered reality* can prove to be helpful. Of course, the proposed typology can be extended as needed.

In this chapter, we have encountered a variety of strategies that are used to augment the real environment. As expected, relating virtual and real content spatially or content-wise are prominent fundamental approaches to AR. However, the design of AR experiences does not have to stop on this level. Designers and developers can build on spatial and content-based relationships, and for instance, include narrative elements, utilize cross-modal effects or simulate interactions between virtual and real elements. In any case, the creation of AR experiences not only requires the development of interesting virtual content but also necessitates the design and establishment of relation-

ships between the virtual and the real. In line with this, we want to encourage AR developers, artists and designers to compose their own, novel, and possibly unique relationships between virtual content and the environment.

As discussed, we understand augmentation as a result of the *perceived* relationships between the virtual and the real. Accordingly, we believe what ultimately matters is whether participants perceive a link between the two. Our investigation builds on the premise that the links between the virtual and the real are *experienced* by potential participants. However, in practice we have to be more careful: Establishing a relationship between the virtual and the real not automatically ensures that a participant also experiences this relationship. What is more, a participant might experience relationships that have never been created or intended. For instance, a museum visitor might listen to a virtual museum guide, but associate the information with the wrong artwork. Consequently, they might not experience the intended relationship but experience another link instead. Similarly, one and the same scenario might be experienced as AR by one person but not by another.

Unfortunately, the question of whether we experience a relationship between the virtual and the real (or two things) is difficult to rise during an AR experience. This is because the question alone causes us to think about the virtual in relation to real, and thus establishes a link between them. Furthermore, while participants have to experience a relationship, they not necessarily have to be conscious of this fact. For instance, a person might not be aware that a drinks' taste is affected by its virtual color (cf. Narumi, Sato, et al., 2010). Yet, if their drink tastes differently from how it "normally" would taste, the person experiences the influence of the virtual color on the real drink. Likewise, a participant might not even be aware that a certain object in their environment is virtual, and hence, not consciously experience any relationship between something virtual and the real.

In our opinion, the challenges with making sure participants experience the desired relationships between the virtual and the real, should not stop us from thinking about and designing those relationships. We believe it is safe to assume that participants will be much more likely to experience the desired link between the virtual and the real if this link has been deliberately designed.

In this chapter, we have placed AR in a broader context. This has revealed that many of the underlying concepts that play a role in AR also are at play in areas that usually are not considered AR. For instance, ideas such as removing real stimuli from a participant's perception are common in AR, but also have been a popular research topic in the audio engineering context. Furthermore, ideas such as changing the properties of real objects through virtual additions are certainly not exclusive to the field of AR. For instance, one can argue that foods

that contain artificial flavors also present us with a combination of the virtual and the real.

The chapter furthermore has revealed interesting ties between AR and other research fields. For instance, the idea of virtual objects sensing the environment and acting in it relates AR to the fields of AI and Sentient Computing. Similarly, some of the encountered approaches to transforming the qualities of real objects show the relevance of perception research to the AR research field.

Finally, our review has shed light on several topics that seem to have received surprisingly little attention in existing AR research so far. One such topic is multimodal perception. Considering that AR is often concerned with blending virtual and real stimuli to create one seamless experience, it would make sense to explicitly explore how different stimuli are integrated on a perceptual level. We thus suggest considering multimodal integration of virtual and real stimuli in future research. Given that AR commonly deals with transformations of the real world, it would be particularly interesting to further explore the role cross-modal effects can play in AR in order to facilitate such transformations. Furthermore, we have gained the impression that relatively little attention has been devoted to realizing new, non-realistic interactions between the virtual and the real. We expect that AR allows us to create new forms of virtual content that does not appear to adhere to physical laws and consequently, allows us to realize new forms of interactions between the virtual and the real. We will explore this topic further in the following chapters.

Hitting imaginary walls, pulling virtual strings

What augmented reality can learn from urban dance

A few weeks ago my colleagues convinced me to join their weekly Hip Hop fitness exercise at the university Sports center. Moving my limbs in the rhythm of well-known radio hits turned out to be more difficult than I had anticipated. After all, I had been running to similar music on a regular basis. A particularly difficult move required us to turn 360 degrees while at the same time imitating a windmill with our arms. In order to help us get the movement right, our instructor gave us a simple but effective hint: "imagine two walls, one in front of you and one behind you. You can only move between them, your arms should not hit the walls." To be honest, this tip didn't help me at first. Rather, I was distracted—those invisible walls reminded me of my research into augmented reality (AR) and the presence of virtual objects in real space. These walls we had to avoid were solely a product of our imagination. Nonetheless, our movements acknowledged their presence. The walls were, in a most basic and fundamental way, becoming part of and augmenting our surroundings... could we call this a form of imagination-based AR? Could it be that dance and AR had more in common than I thought?

Only minutes later this suspicion got confirmed. By now, our hands were connected to our feet with imaginary strings. In order to move our feet, we had to pull the strings. To my surprise, when our teacher illustrated the movement, it appeared as if those strings indeed existed. Although I knew that they were merely imaginary, and even though I could not see the strings, some part of me was fooled into believing that they were actually there. Given the teacher's movement, her hands and feet simply had to be connected by a thin, invisible rope! There was no digital technology required, I was not wearing a headset, nor was I staring at a screen: a relatively simple movement was sufficient in order to convey the presence of virtual objects (or, to be precise, virtual strings) in real space. It might not have looked like AR, but watching these invisible ropes certainly *felt* a lot like AR!

Over the next days, aching muscles reminded me to investigate this phenomenon further. Luckily, I already knew where to start. In 2013, I had attended a presentation about illusion-based dance by Diego Maranan at the Creativity and Cognition conference in Sydney (see Maranan, Schiphorst, Bartram, and Hwang, 2013). During his talk, Maranan not only illustrated technological metaphors used in the urban dance styles 'liquid', 'digitz' and 'finger tutting', but at the same time mesmerized the audience with movements that made us doubt whether his hands were constrained by the same kind of bones we had. Among the videos that were shown, one dancer had left a lasting impression: Albert Hwang, a master in making three-dimensional boxes appear in real space—solely by running his hands through thin air. A quick look at his YouTube channel (Hwang, 2006) decided the matter; I had to find out how dancers created the illusion that imaginary objects existed in space, I wanted to know how much illusion-based dance styles and augmented reality had in common and I definitely had to master some of those movements myself.

Dance AR?

Compared to learning the basics of liquid dancing, my theoretical considerations were rather simple. AR and illusion-based dance styles have one central aspect in common: both create the impression that virtual objects actually exist in our real, physical environment. If we understand augmented reality as a concept of combining and relating the virtual and the real (see chapter 3) rather than a collection of technologies, it is not far-fetched to think of these dance-illusions as a time- and movement-based form of augmented reality. What is more, the traditional, technology-focused field of AR can learn quite a few things from urban dance!

So how does urban dance approach the virtual and how do their methods inform the general field of AR?

No technology required!

First of all, dance teaches us that there are alternative means to display virtual objects in space besides AR technology. AR most commonly uses smartphone screens, heavy headsets or other kinds of visual displays that overlay the real world with virtual elements. In illusion-based dance, imaginary objects are revealed to the audience through a dancer's body movement. The dancer can, for instance, run his or her hands over the shape of an imaginary object in order to make it appear as if the object is actually present (Hwang, 2012). Illusion-based dance reminds us that AR is not restricted to digital mediums and that we do not have to resort to computer technology in order to make virtual objects appear in real space. Lamers (2013) has discussed the Pepper's Ghost as an instance of pre-digital AR. In this regard, dance-illusions can serve as yet another compelling example of AR that remains in the physical domain.

Realism, really?

AR should be more like reality and virtual objects should both look and behave like real, physical objects! At least, this is the impression I get from much existing AR research. Scientists and developers strive for photorealism, they struggle with occlusion and investigate how virtual objects can cause reflections and cast shadows just like real objects do (see, e.g. Agusanto, Li, Chuangui, and Sing, 2003; Gibson and Chalmers, 2003; Kanbara and Yokoya, 2004). Likewise, it is said that virtual objects should behave and interact with the world like real objects (S. Kim, Kim, and Lee, 2011). If we are to believe existing research, a virtual ball is supposed to drop and bounce on the floor, just like a real ball would. There is certainly nothing wrong with that. However, illusion-based dance shows us that another approach is possible. Dance shines when it comes to expressing simple geometrical shapes and structures, such as rectangular boxes or walls. In some respect, these 'dance-objects' could not differ more from real objects. First of all, dance-objects do not adhere to our physical laws; they commonly float in space, right before the dancer. At the same time, the way a dancer moves them about in space implies that they do, however, have a certain mass—the mass just does not cause them to fall down. And of course, unlike real objects, these imaginary objects are essentially invisible and certainly do not occlude what's placed behind them. More than that, they often appear out of nothing just to disappear in thin air a few seconds later. Fascinatingly, it does not bother us that these imaginary objects are not really present, don't look like real objects and do not behave like anything we know from the physical world—the objects are believable and convincing nonetheless!

What you see isn't what you get

I expect multimodal AR to become one of the more interesting topics in the future. However, I do not think that a multimodal or richer sensory experience is always better. In their paper on illusion-based dance styles, Maranan et al. (2013) make an interesting observation: when dancers let imaginary boxes appear in space through their movement, the viewer can interpret this in two different ways. Either there is no box in space and the dancer is moving in a very complicated way or

there is a box in space that guides the movement of the dancer's hand. While watching, our eyes tell us that there is no box but our body (or our embodied cognition) tells us that there is. Maranan et al. propose that it is "this moment of embodied/cognitive dissonance [that] makes the movement compelling" (p. 173). I believe that AR can benefit from a similar dissonance: looking at a breakfast cereal box through our phone's screen, we see the virtual dinosaur eating our cereal, but we cannot touch it. Our eyes tell us "it is there" while our body and mind tells us that it isn't. I do not claim that all AR benefits from such a dissonance. But I am convinced that it can actually add to—rather than subtract from—the overall AR experience.

The power of movement

Ultimately, AR can learn from illusion-based dance that movement is a powerful means to express the presence and properties of virtual content. By moving virtual objects through space, AR can communicate properties that it could hardly convey otherwise. If a virtual leaf moves through space in a certain way, its movement shows us that there is wind. If a virtual ball rolls over a real floor, it tells us something about its weight and resistance. Furthermore, using movement, we are able to create the impression of yet other—invisible—objects being present in space. How would you display an invisible wall with AR technology? Dance gives the answer: by having something bump against it, by movement! And there are more possibilities: if a virtual object looks heavy but moves through space weightlessly, we might be able to discern a change in gravity. By rewinding their movements, good dancers are almost able to fool me into believing that time goes backwards. Maybe AR technology can evoke a feeling of time moving differently by rewinding the movement of objects or by varying their speed. I hope future AR will explore what can be expressed by simply moving virtual objects through real space.

Future AR is not reality, it is our imagination

Let us return to the imaginary walls that were occupying the university's dance studio some weeks ago. I am not sure whether these walls can be called AR. But I am sure that a dancer will not be able to create the illusion of a virtual wall in space without imagining the wall first.

In the future, AR will surely overcome many technical challenges. However, the future of augmented reality is not only about what is or will be possible technically. It is also about what we can imagine and how our imagination works. One of AR's unique powers is that it can be different from our real, unaugmented reality. But how can virtual objects differ from real objects without losing their believability? How can augmented reality differ from reality? Studying related arts such as dance, mime or magic helps us find answers and think outside of our imaginary, invisible and virtual boxes.

5 *From Imitative to Imaginative Realities: Influences and Interactions Between the Virtual and the Real*

AR allows us to experience virtual objects in our otherwise real environment. These virtual objects can not only passively *exist* in otherwise real the world, but they can also *act* in and *interact* with this world. For instance, a virtual ball can seemingly collide with a real wall, and a virtual toy can sense and react to its owner (see chapter 4). In this chapter, we take up this idea of virtual-real interactions and examine how the virtual and the real can influence each other in augmented reality. We explore whether and how real objects can affect virtual objects and vice versa.

Our exploration is driven by our own curiosity and imagination. We envision scenarios where a virtual ball bounces on a real sidewalk, where real wind moves virtual leaves, where real doors open for virtual objects (see figure 5.1) and where virtual objects get wet when it rains. Furthermore, we wonder, whether virtual and real objects can interact in novel ways, allowing us to experience influences that cannot exist in a purely physical world. Ideally, AR would allow us to both imitate the real world, as well as realize new imaginative realities that go beyond physical laws and allow the virtual and real to behave and interact according to our own ideas.

Of course, the suggested ideas of imitating the real world on the one hand and creating new realities, on the other hand, are not new. The strive for realism as well as the creation of new and imaginative forms of realities can, for instance, be witnessed in the context of literature, gaming, photography and painting. When it comes to computer-generated virtual content, both directions can be traced back to Sutherland's (1965) vision of an 'ultimate display'—a room in which a computer controls the existence of matter. In the paper that describes his vision, Sutherland (1965) suggests: "A chair displayed in such a room would be good enough to sit in. Handcuffs displayed in such a room would be confining, and a bullet displayed in such a room would be fatal" (p. 2). With this, he describes computer-controlled objects that interact with the real world just like their real counterparts. At the

Figure 5.1: Real doors can open for virtual objects. Image © Hanna Schraffenberger and Edwin van der Heide.

same time, Sutherland also implies possibilities for realizing different types of behaviors and creating imaginative environments. He emphasizes that such an ultimate display "could literally be the Wonderland into which Alice walked" (p. 2). The idea of moving beyond the simulation of physical laws also comes back in his comments on computer displays in general. In this context, Sutherland (1965) explicitly argues that "[t]here is no reason why the objects displayed by a computer have to follow the ordinary rules of physical reality with which we are familiar" (p.2).

As Sutherland's paper shows, the ideas of mimicking the real world as well as creating new types of realities have a long history. However, augmented reality is no ultimate display, and AR technology cannot control the existence of matter. This raises the question of whether and to what degree both visions can actually be realized in the context of AR. In fact, there are reasons to doubt the feasibility of either idea when it comes to interactions between virtual objects and the real world.

With respect to imitating real-world interactions, one faces the challenge that many virtual objects cannot directly apply forces to real objects (cf. S. Kim et al., 2011). Usually, the real world can affect virtual elements, but virtual objects cannot affect the real world in return. In the context of Sutherlands examples, this means that we can make virtual bullets fly through a real environment, but that these bullets won't have any effect when they hit someone or something real.[1] If the real world does not seem to be affected at all by the behavior and actions of virtual objects, this might seem unbelievable.

[1] Arguably, in the context of bullets this can be considered an advantage rather than a problem.

When it comes to creating new and imaginative forms of actions and reactions, believability is an important issue. It is not clear what interactions and influences between the virtual will be perceived as credible and meaningful. Technologically, there is nothing keeping us from

having virtual raindrops 'fall' upwards, from turning virtual frogs into princes when they are kissed, or making virtual objects 'teleport' to an entirely different position when they collide with real elements. A question that arises is whether behaviors that defy physical laws are plausible in a real-world context. As we will see, some researchers seem to believe that for virtual objects to appear as if they were part of the real world, they also have to behave like real objects and stick to the rules of that world. The question arises whether virtual objects might "have to follow the ordinary rules of physical reality" (see Sutherland, 1965, p.2) after all when they appear to exist in the context of our "physical reality". Personally, we do not expect this to be the case and hope to dispute the claim that virtual objects (always) have to behave like physical objects.

In this chapter, we take up these different lines of thought about interaction between the virtual and the real. In particular, we address the following three considerations: First, the virtual is free from physical laws. Hence new forms of influences between virtual content and the real world can be realized. Second, virtual content cannot directly apply forces to real objects. As a consequence, interactions that we know from the physical world might not be possible. Third, not everything that is technically possible is necessarily also credible. For instance, in order to appear as a believable part of the physical environment, virtual objects might have to adhere to the same laws as real objects.

These three considerations inspire us to ask the following questions: What types of interaction between the virtual and the real are both possible and credible? Can the virtual and the real interact like physical objects? Can they interact in new but believable ways? We are interested in both problems that arise when virtual and real objects seemingly exist in the same space, as well as in possibilities that emerge from such an AR setting. In particular, we are interested in new forms of interactions that are unique to AR and that could neither exist in a solely physical nor in an entirely virtual world.

In order to answer the presented questions, we follow both a theoretical and a practical approach. We review existing research and AR works, conduct our own initial series of practical experiments as well as reflect upon these experiments.

The topic of interaction between the virtual and the real has emerged as a central theme in in the previous chapter. Because we want every chapter to be able to stand on its own, we will revisit topics and examples discussed in the previous chapter, and in particular section 4.9 and section 4.10. However, we will move far beyond the previously discussed material and primarily address the topic from new perspectives. For instance, we have made a distinction between physical interaction on the one hand and behavioral interaction on the other hand in chapter 4. In this chapter, we choose a different point of view. We focus on interactions that mimic real-world interactions,

as well as explore imaginative forms of interactions, that do not exist in reality, but that nonetheless appear believable. In line with this, we distinguish between (1) *imitative* interactions that could actually occur between physical elements in the real world and (2) *imaginative* interactions that cannot exist in a purely physical world, but that are perceived as credible or convincing nonetheless.

The main goal of this chapter is to answer two key questions: (1) whether virtual objects can interact with physical objects in a realistic manner as well as (2) whether they can interact in imaginative but believable ways. We first search for answers to these questions in existing AR research. This theoretical exploration is presented in section 5.1. Subsequently, we take a more practical approach to interaction between the virtual and the real and address the questions with a series of small exploratory experiments. This practical exploration is presented in section 5.2. Finally, we present a general discussion and conclusion (section 5.3). We reflect on our findings and conclude that virtual and real objects can believably simulate real-world influences as well as influence each other in imaginative ways that have no equivalent in the physical world.

As mentioned, we are particularly interested in imaginative but yet believable forms of interaction between the virtual and the real. In this study, the question whether the interaction is believable was evaluated from the author's subjective point of view. Furthermore, the question was addressed in the context of an 'ordinary everyday environment'. This is important because the believability of an object's behavior likely depends on the situation and context in which the behavior takes place. For instance, different forms of behavior will be accepted as believable in the context of a game than in the context of a working environment. (This is likely true both for real and for virtual objects.)

Our interest in believable forms of *interaction* between the virtual and the real entails an interest in the *behavior* of virtual objects. However, our key interest is virtual behavior in relation to the real world, rather than virtual behavior as such. The more general question of when the behavior of virtual objects is believable falls out of the scope of this thesis. We are focusing on the interaction between the virtual and the real because this issue is specific to the field of AR.

This chapter addresses two issues that are often approached independently from each other in existing AR research: First, the interaction between a participant (user) and virtual content. Second, the interaction between virtual objects and other physical objects in their surroundings. We address both of these topics, but propose a view that consolidates the two: We see the participant as part of the augmented environment. Accordingly, we see interaction between virtual content and the real environment as a broader, more general field that also encompasses the interaction between a participant and the virtual

objects.

The idea of interaction between a real environment and virtual content entails that there is some kind of mutual influence between the virtual and the real world. Given our interest in the participant's experience (rather than technological aspects), we are focusing on scenarios where the participant either witnesses interaction between virtual objects and the real environment and/or interacts with virtual objects her/himself. If we look at existing AR projects, participants often interact with virtual content on a technological level: Many AR systems react to the movement of the participant and consequently, present virtual content that depends on the participant's point of view (see subsection 3.1.2).[2] For instance, a virtual cup might look different, depending on whether a participant looks at it from above or from the side. If the participant reacts to what they see, and e.g., move to see an object from yet a different angle, one could speak of a mutual influence (and thus interaction) between the virtual and the real. However, in our opinion, this does not mean that the participant also *experiences* some interaction with the virtual content. Arguably, simply looking at an object from different points of view is not experienced as interacting with the object since the object itself does not react to the actions of the participant. Similarly, the fact that an object looks different from different angles does not make it feel like the object is affected by us. Accordingly, such scenarios fall out of the scope of our exploration. Instead, we focus on scenarios where virtual objects actually appear to be affected by the real world and vice versa. This is, for instance, the case when a virtual object changes its size, color, shape or position as a response to colliding with a real object. In our review of existing AR literature, we will make different views on what constitutes interaction explicit. However, unfortunately, it is not always clear how other authors define interaction.

[2] In fact, AR is often defined in terms of such systems (e.g., Azuma, 1997).

Our exploration is focused on underlying ideas and conceptual possibilities rather than issues of implementation. Yet, we will at times mention different technological approaches that facilitate interactions between the virtual and the real. This is because sometimes, conceptual ideas and technological solutions are closely interlinked. Furthermore, we want to support future research and development that intends to implement the underlying ideas.

Given that we are interested in conceptual rather than technological possibilities, our practical explorations use basic technological implementations. We generally work with cheap and readily available office hardware rather than dedicated AR devices. In our opinion, this is sufficient to experience (basic/fundamental) interactions between the virtual and the real and hence, we see no need to work with different materials instead. So far, this practical exploration is solely based on our own experiences with the AR scenarios. It does not yet include any empirical research with participants. However, it can serve as the

first step towards such empirical studies, as it identifies possible forms of interaction and scenarios that could be studied with participants in the future.

In addition to the AR research field, many other disciplines are also interested in interactions between the virtual and the real. For instance, research into conversational agents has been very concerned with creating virtual humans that converse with and react to real human input just like a human (see, e.g., Cassell et al., 2000). In this chapter, we primarily focus on issues that are unique to AR and that arise from the fact that virtual and real objects seemingly exist *in the same physical space*. Topics that are a primary concern in other research areas fall out of the scope of our investigation.

5.1 Theoretical Exploration

The presence of virtual content in an otherwise real environment opens up possibilities for influences between this environment and the virtual content. However, it is still unclear what forms these interactions can take. In order to get a first idea about what reactions and interactions between the virtual and the real are possible and credible, we will have a look at existing AR projects and review opinions on how interaction between the virtual and the real can, should or could look like.

In the following, we will first address imitative interactions and subsequently explore imaginative interactions. However, this clear distinction between the two is somewhat misleading. Rather than as two distinct groups, the two forms of interaction can be seen as a continuum. Often, projects will mimic reality in some form, while deviating from it in other ways. We have placed ideas in one of the two categories based on the concept we want to emphasize and illustrate—this might not always be the most prominent feature of a certain project.

5.1.1 Imitative Interactions

Can virtual objects interact with the real world in the same manner as real objects? According to some researches, realistic interaction between the virtual and the real are not simply a possibility but rather, a *necessity* for successful AR.

For instance, Breen et al. (1996) point out: "For the new reality to be convincing, real and virtual objects must interact realistically" (p. 11). Effects that, according to the authors, need to be considered in AR include occlusions, shadows, reflections, refractions, color bleeding, kinematic constraints, collisions as well as responses to collisions and external forces. The authors not only assert that such real-world effects and influences have to be implemented in AR, but also propose techniques that approach some of the issues. In particular, they present

techniques for realizing occlusions between virtual and real objects as well as for placing dynamic virtual objects on top of static real objects. They do this by 'simulating gravity' and detecting collisions between virtual and real objects. Essentially, virtual objects are moved downwards in the real space until they collide with a real object. As a result of this process, virtual objects are placed on real objects. For instance, a virtual lamp might appear to stand on a real desk.

The Physical Artifact

Breen et al. (1996) are not alone with their view that realistic interactions between the virtual and the real are necessary. For instance, S. Kim et al. (2011) write: "In order to make virtual objects move as if they coexisted with real objects, the virtual object should also obey the same physical laws as the real objects, and thus create natural motions while they interact with the real objects." (p. 25).[3] The authors not only argue for such realistic interactions but also identify challenges that arise when attempting to implement them. More specifically, they illustrate that problems can arise due to the inability of virtual objects to affect real objects. They argue that when a virtual and real object collide, both objects should be affected by this collision. However, as AR systems typically only can control the movement of the virtual object, a real object will usually appear unaffected by a collision with a virtual object. The authors believe that such interactions "may contradict the physical cognition of humans" and argue that it "diminishes the sense of realism of AR". S. Kim et al. (2011) coin this phenomenon "physical artifact" and continue to explore when these artifacts occur, demonstrate instances of the problem and also present ideas about how the problem can be avoided.

[3] It is not clear whether the authors believe this to be generally true, or only assume this to be the case when a participant perceives the virtual object as a real object. As we will see later, they give an example where a real paper cup remains unaffected when it is hit by a virtual ball. In this context, they write "If a viewer perceives the virtual ball as a real one, this physically incorrect response will contradict the physical intuition of the viewer, and thus may harm the immersiveness of the viewer considerably" (p. 27). In line with this, it might make a difference whether the virtual object is perceived as a real or as a virtual object.

The practical exploration of these physical artifacts by S. Kim et al. (2011) includes several interesting examples of influences between virtual objects and the real world. For instance, they present an example where virtual boxes and spheres fall down, collide with a real table tennis racket and, according to the authors, show plausible responses. Furthermore, they demonstrate an example of the "physical artifact". A virtual ball falls down, bounces off a physical slanted plane, and collides with a real paper cup. This collision causes the virtual ball to move into a different direction. However, the real cup remains unaffected. From a technological point of view, this is not surprising, as the virtual ball does not actually apply any force to the cup. However, from a perceptual perspective, things might appear differently. As the authors explain "[i]f a viewer perceives the virtual ball as a real one, this physically incorrect response will contradict the physical intuition of the viewer, and thus may harm the immersiveness of the viewer considerably" (p. 27).

The authors also propose a solution to avoid such collisions. Their idea is to change the parameters used in the physical simulation in

a way that maintains the realism and at the same time, avoids the collision. They demonstrate this by adapting the previous example. Due to a small adjustment in one parameter, the virtual ball bounces off the real plane in a slightly different (but still credible) angle, thereby avoiding the collision with the cup.

Imitating Optical Interactions

In line with the belief that realistic interactions are necessary, we can find a wide variety of projects that attempt to realize such realistic interactions (cf. section 4.9). With respect to this, a lot of research seems to focus on realizing realistic *optical effects* between virtual objects and the real world. To mention just a few examples: Many researchers work on methods that allow an AR system to take the illumination of the real world into account when rendering virtual objects (e.g., Madsen et al. (2006) and Kanbara and Yokoya (2004)). This makes it possible for real light sources to affect the appearance (e.g., shading) of virtual objects. Furthermore, the information about the illumination of the real world can be used to make virtual objects cast realistic shadows onto the real world and affect the appearance of the real world in return. In addition, AR research focuses on realistic caustics, reflections and refractions. For instance, Kán and Kaufmann (2012) demonstrate a rendering system that is capable of these optical effects. Their demonstration displays a virtual glass that casts a virtual shadow onto the real world as well as features correct refractions of surrounding elements, such as a person's hand and physical colored cubes that stand next to the virtual glass (see figure 4.9). Similarly, Pessoa et al. (2010) propose a rendering technique that focuses on the effects of the real environment on the appearance of virtual objects. Their demonstrations include, for instance, a virtual vase that appears to be illuminated by the real environment, a teapot reflecting surrounding physical objects as well as color bleeding effects where light from real surfaces appears to color virtual objects.

While visual effects get a lot of attention, very little research addresses similar issues with respect to other modalities. One of the few exceptions is the work by Lindeman and Noma (2007). The authors point out:

> In order to attain a truly merged experience, the two [real-world and computer-generated] stimuli should undergo similar transformations, so that, for example, a virtual character receives the same lighting effects (light position and intensity) as objects in the real world. In fact, this applies to all sensory modalities; the voice of a virtual character should also be influenced by environmental objects, such as occluders or reflectors. (p. 175).

Imitating Dynamic Interactions: The Real Affects the Virtual

In addition to research that focuses on realistic optical interactions, research has also pursued realistic *dynamic* influences and interactions. Here, the main focus is on making virtual objects move as if they were indeed affected by the real world. This often happens with respect to gravity. Often, virtual objects appear to have a physical mass and seemingly are affected by gravitational forces. This happens, for instance, in the above-discussed exploration of the physical artifact by S. Kim et al. (2011). As illustrated, it is difficult for virtual objects to physically affect real objects. However, real objects can easily affect virtual objects. In the exploration by S. Kim et al. (2011), a physical slanted plane and a real paper cup affect the trajectory and movement of the virtual ball. Similar examples have been presented by Chae and Ko (2008), who simulate gravity, and take attributes such as weight, gravity, friction, elasticity and force into account when determining the movement of virtual objects. They demonstrate this with a virtual ball that falls downwards, collides with a real object, and bounces off this object. Another similar example of a virtual ball that bounces on a real table is provided Valentini and Pezzuti (2010).

The idea of simulating real-world interactions between real and virtual objects often comes back in the context of AR games. For instance, in the AR version of Air Hockey by Ohshima et al. (1998), hitting a virtual puck with a real mallet, causes the puck to change direction—presumably in the same way as a real puck. Furthermore, Namee et al. (2010) propose an engine for creating plausible physical interactions between virtual and real objects in the context of AR games. To demonstrate this engine, the authors present two proof-of-concept AR games. The first is a table-top racing game where virtual cars interact with both virtual and real objects. For instance, they can crash into a real object or drive over a real ramp. In their second game, the player has to move virtual crates around the environment with a small real robotic forklift. The real forklift can, e.g., raise and lower crates with its fork, push around and carry crates or crash through multiple crates. All of these interactions have an equivalent in a solely physical world.

As discussed above, simulations of dynamic real-world interactions are often incomplete due to the fact that many virtual objects cannot affect the real world. Essentially most projects only simulate the influence of the real world on the virtual objects. For instance, to the best of our knowledge, the collision between the virtual puck and the real mallet in "AR²Hockey" (Ohshima et al., 1998) only affects the puck and has no effect on the physical mallet at all.

Imitating Dynamic Interactions: The Virtual Affects the Real

Whereas many projects explore how the real world can affect virtual objects, rather few projects focus on the way virtual objects can affect the real world. One of the few exceptions is the project called *Kobito - Virtual Brownies-* by Aoki et al. (2005). Here, virtual creatures (so-called Kobitos) move a real tea caddy. A similar project has later on been realized by Kang and Woo (2011). In their project, a virtual character is able to interact with a physical toy cart. For instance, it can push and pull the cart. Furthermore, participants can interact with the virtual object through interaction with the physical object. E,g., they can move the cart, which can cause the virtual character to fall down. Both projects extend real objects with electronics (e.g., motors) to allow the virtual characters to move the real objects.

A different approach to allowing virtual objects to affect real objects is found in the table-top game called *IncreTable* by Leitner et al. (2008). In this game, both virtual and real items can be arranged on the table to solve puzzles. Among the available objects are, e.g., virtual and real domino stones. In order to facilitate interaction between the virtual and real dominos, the authors implemented so-called portals (see figure 5.2). These special physical interfaces can both push a real domino stone when it is hit by a virtual one as well as detect a falling real domino stone to push a virtual one.

Figure 5.2: Virtual and real domino stones can interact with the use of so-called portals. Reprinted from J. Leitner et al. (2008). "IncreTable, a mixed reality tabletop game experience". In: *Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology*. ACM, pp. 9–16. Reprinted under fair use.

Another project where the virtual affects the real is the artwork "Beyond Pages" by Masaki Fujihata (see Kunst und Medientechnologie Karlsruhe, n.d.; MediaArtTube, 2008). The work consists of a real room, that contains, a real desk, chair and lamp. On the desk, there is a virtual book and stylus that allows the visitor to interact with the virtual book. On one of the pages, a virtual light switch is depicted.

If the visitor switches the virtual switch, the real lamp in the room can be turned on/off. Here, a physical interface from the real world is replaced with a virtual interface.

As these projects show, the virtual can have an effect on the real world. This effect can be *simulated*, as in the case where virtual objects cast virtual shadows onto real objects (see, e.g., figure 4.9). However, the outcome of the interaction can also be real. For instance, the above-discussed virtual tea caddy and toy cart *actually* move in the real world and the real domino stones *actually* fall when 'hit' by virtual dominos.

Utilizing Real-World Interactions

Another approach to interactions between the virtual and the real is not to mimic them, but to make use of actual interactions in the physical domain. A simple example of this concept would be playing back the voice of a virtual character via speakers. If this happens, the characteristics of the surroundings will naturally affect the voice. For instance, if the sound of a virtual creature is played back on a speaker in a big church, it will sound different than if it is played back outside without any need to simulate this effect. We hence can utilize natural interactions that occur in the physical domain.

An example of a project that makes use of interactions that naturally occur in the physical domain is the installation *Radioscape* by Edwin van der Heide (2000-). The installation consists of several radio transmitters that are distributed over a part of a city. Each transmitter broadcasts one layer of a meta-composition. Listeners can pick up several signals at a time with a custom developed receiver. The volume of each of the single layers depends on the listeners' distances from the corresponding transmitters. Due to the chosen wavelength, buildings become conductors and resonators for the transmitted signals. The physical environment is excited by and responds to the transmitted radio waves, ultimately affecting the virtual content and influencing what one hears. Although this interaction happens in the physical domain, we can argue that the transmitted virtual content interacts with the physical landscape.

A completely different approach that also utilizes real-world interactions is found in the commercial product *Sphero* (2011). Sphero is a robot ball that—when viewed with the corresponding smartphone app—is turned into a virtual beaver (cf. J. Carroll and Polo, 2013). Because the virtual ball is affected by the real world (it can, e.g., not pass through real walls, and is affected by gravity) the virtual beaver is also affected by the real world accordingly.

Participant-Focused Interaction

The idea of mimicking real-world interactions also comes back in the specific case where a *participant* interacts with virtual content. For

instance, Craig (2013) claims that AR not only allows us to interact with virtual content but also, that we can interact with it *in the same way* as we interact with physical objects:

> In brief, the core essence of an augmented reality experience is that you, the participant, engage in an activity in the same physical world that you engage with whether augmented reality is involved or not, but augmented reality adds digital information to the world that you can interact with in the same manner that you interact with the physical world. (p. 2)

It has to be noted, though, that Craig's definition of interaction is rather broad, and appears to include looking at virtual content from different perspectives. For instance, he writes:

> [...] a person can sense the [digital] information and make changes to that information if desired. The level of interactivity can range from simply changing the physical perspective (e.g., seeing it from a different point of view) to manipulating and even creating new information. (p. 16)

Furthermore, Craig (2013) also points out the possibility to interact with virtual content in new and additional ways that have no equivalent in a physical world. For instance, unlike a real house, a virtual house in a vacant lot could be moved around or viewed in different colors. (Interactions that are impossible in a physical world will be discussed in subsection 5.2.2.)

An example that shows that real-world interactions can indeed be imitated, is the above-mentioned AR version of AIR hockey "AR²Hockey" by Ohshima et al. (1998). Here, two players play air hockey using a real mallet to hit the virtual puck that moves over a real table (cf. Azuma et al., 2001). This means, that the game simulates the interaction between a puck and mallet that we know from the real world. Similar ideas have been used in other contexts. For instance, the AR version of the game Quake (Piekarski and Thomas, 2002) allows participants to virtually shoot at monsters by means of a physical toy gun—essentially also copying an interaction that we find in the real world.

Another project that allows a participant to interact with virtual content in the same way we interact with real content has been realized by Corbett-Davies, Dünser, and Clark (2012) and Corbett-Davies, Dünser, Green, et al. (2013). In contrast to the above-mentioned projects, it does not make use of a physical interface but allows participants to interact with virtual spiders with their bare hands. Participants can, for instance, pick spiders up and carry them around.[4]

The idea that AR can allow for similar interactions as the real world is also taken up by Bau and Poupyrev (2012). Similar to Craig (2013), the authors state that "[t]he fundamental premise of AR is to enable us to interact with virtual objects immediately and directly, seeing, feeling and manipulating them just as we do with physical objects." (p. 89:1).

[4] The underlying idea that people with a fear of spiders can interact with these virtual spiders, while they could never interact with real spiders in the same way, suggests that interaction with virtual objects in some way differs from interacting with actual spiders even when it is extremely realistic. Presumably, the fact of knowing that something is virtual will change how the interaction is experienced, and thus creates some difference.

However, Bau and Poupyrev (2012) emphasize that this goal is often *not* reached and consequently propose a means to change this. Their tactile technology REVEL provides virtual tactile feedback that can extend real, physical objects by means of virtual textures that are felt when touching the object. Ultimately, such a tactile augmentation of real objects leads to an object with both a virtual and a real component that a user can not only see but also touch and interact with physically.

Like Bau and Poupyrev (2012), several other researchers emphasize common limitations when it comes to interacting with virtual objects in AR, and consequently, propose a way to change things. For instance, Vallino and C. Brown (1999), point out that "Augmented reality systems have been interactive only to the extent that the user could move about the workspace and be a passive viewer of the visually augmented scene" (p. 199).[5] To change this, the authors then propose a setup that allows users to physically interact with virtual content by means of a Phantom force-feedback device. As discussed in subsection 4.4.5, this device has similarities with a small robot arm (cf. Vallino and C. Brown, 1999) with a thimble at the end. Placing their finger in the device's thimble, the participant can feel the surface of a virtual object, experience its weight and dynamic forces, as well as move the object around within the real environment. In their demonstration, participants can, e.g., experience a virtual globe, spin it around its axis, feel the difference between water and land, or move the virtual cube around in real space with their finger.

Billinghurst (2001)[6] also critiques the then existing possibilities of interacting with virtual content in AR. Similar to (Vallino and C. Brown, 1999), he asserts that "*interaction* with AR environments has been usually limited to either passive viewing or simple browsing of virtual information registered to the real world. Few systems provide tools that let the user interact, request or modify this information effectively and in real time" (Billinghurst, 2001, p. 1, emphasis in original). He consequently introduces the concept of "tangible AR interfaces" as a means to change this. Tangible interfaces take the form of physical objects that have virtual objects linked (registered) to them. Consequently, a user can interact with virtual objects by manipulating the corresponding physical object. The resulting interactions can both mimic real-world interactions, as well as take novel imaginative forms. Both happens, e.g., in the SharedSpace Siggraph 99 project. Here, physical game cards with AR markers (see subsection 3.1.2 for information on markers) on them provide a physical counterpart to the virtual content that is associated with them. The cards can be picked up and moved around to view the attached virtual objects from different perspectives. Furthermore, participants can place corresponding virtual cards together, causing interactions between the virtual objects. These ideas are particularly interesting to us because the implemented interactions only partially

[5] In our opinion, only being able to passively view virtual content implies that these environments provided *no* possibilities for interaction with this content. However, the idea of viewing virtual content from different perspectives is often seen as a form of interaction with the content. This could be because the virtual content indeed adapts to and reacts to the participants movement on a technological level. Here, we approach the topic from a perceptual perspective. Hence, an object that appears static and that does not seem to react to one's movement is not considered interactive because it does not appear to be interactive.

[6] and later Billinghurst, Kato, and Poupyrev (2008) as well as Billinghurst, Grasset, et al. (2009)

mimic real-world interactions and also introduce what the authors refer to as "table magic". For instance, if a physical card is tilted, the virtual object is supposed to slide across the card surface, which mimics physical real-world interactions. However, if a card is shaken, a virtual object can appear on the card or the object can change into another object—naturally, this is something which could not happen in a purely physical world. Kato et al. (2000) put it like this:

> Some of these commands simulate physical phenomena in the real world and other simulate table magic. In all these cases we establish a cause-and-effect relationship between physical manipulation of the tangible interface object and the behavior of the virtual images. (p. 118)

In the following section, we will address such interactions that take imaginative forms rather than mimic our physical reality.

5.1.2 *Imaginative Interactions*

Virtual objects do not have to adhere to physical laws. Hence, they can behave differently and potentially, also interact with and react to the real world in new and imaginative ways. Unfortunately alternative forms of interaction have gained very little attention in the context of AR so far. In the following, we will review research and practical projects that show that believable imaginative interactions might be possible. Unlike the projects in the previous section, the reviewed examples generally focus on interaction between a participant and virtual content. This is the case because existing research has paid little attention to imaginative interactions between virtual content and the real environment in general.

One of the few examples that build on new forms of interaction are those AR projects that facilitate some form of x-ray vision and that allow participants to have a look inside or see through physical objects. An example is the system by Bajura, Fuchs, et al. (1992), which visualizes ultrasound echography data within the womb of a pregnant woman and thus, let's a doctor see through parts of her physical body. Next to the medical domain, this concept is also common in outdoor mobile augmented reality applications. For instance, the mobile AR tools by Bane and Hollerer (2004) make it possible to view the area behind physical walls. As noted by Kalkofen et al. (2009), such projects that make it possible to see hidden or occluded objects seemingly go beyond the physical laws of light propagation. In our opinion, such projects can be seen as a counterpart to projects that simulate realistic optical effects and in particular, realistic occlusions between virtual and real objects.

Another approach where optical effects in AR defy the laws of our physical world is the use of magic mirror setups. The underlying idea is that the real environment includes a mirror that presents a mirrored and augmented version of real world to the participant. Magic mir-

rors defy physical laws in the sense that a mirror reflects something that is not actually in front of it. Typically, the magic mirror takes the form of some kind of digital screen, such as a computer monitor. For instance, M. Kim and Cheeyong (2015) have proposed such a system in the fashion context. Here the idea is that users can see themselves in the mirror with different outfits, make-up, and hair styles. Another example comes from the artist Sobecka, who has created a magic mirror that allows viewers to see themselves in a new way: An animal head appears on top of their own head and mimics their movement and expressions (see figure 5.3). In addition to mimicking the viewer, the animal occasionally creates its own expressions. The artist reports that viewers feel compelled to follow along and enact these animal movements.

In our opinion, the interaction between magic mirrors and the real world has both imaginative and imitative qualities. On the one hand, magic mirrors act like real mirrors and present what is in front of them. In this sense, they can be considered to imitate real-world interactions. On the other hand, the name "magic mirror" suggests that there is something mysterious or supernatural going on. If the magic mirror is experienced as something magical rather than natural (e.g., because it reflects things that are not really there), they can be considered "imaginative".

Figure 5.3: In Sobecka's mirror, the viewer sees an animal overlaid on their own reflection. Image from http://www.gravitytrap.com/artwork/perfect-creatures. Printed under fair use.

Although few projects focus on imaginative interactions between the virtual content and real objects, there are some projects that allow a participant to interact with virtual objects in new ways using physical objects. An example is the above-mentioned Siggraph 99 project (Kato et al., 2000). In this context, we have already seen that shaking a physical card can cause a virtual character to appear or to change into another character. Similar interaction possibilities have been explored with the so-called *MagicCup* interface (Billinghurst, Kato, and Myojin, 2009). This tangible interface allows participants to cover virtual objects with the cup. The MagicCup then "holds" the virtual object and can be used to interact with it. Interaction using the cup often mimics physical interactions (e.g., one can move a virtual object around by

moving the cup). However, the MagicCup also allows for a form of interaction that defies physical laws: by shaking the cup, the object inside is deleted. A similar concept has been explored with the so-called "Magic paddle" (Kawashima et al., 2001). Here, participants use a small real paddle to interact with virtual furniture. Like with the Siggraph 99 project (Kato et al., 2000) and the *MagicCup* (Billinghurst, Kato, and Myojin, 2009), some interactions mimic real-world interactions whereas others could never exist in a purely physical world. For instance, virtual models can be removed from the space by hitting them with the paddle.

Aside from using simple physical objects, participants can also interact with virtual content using some sort of digital or virtual interface. This form of interaction is, e.g., part of the mobile game *GeoBoids*, which will be explained in more detail below. The game allows players to catch virtual bird-like creatures by making a swiping gesture on their phone's touch-screen.

In addition to using physical objects, digital interfaces and virtual controls, some AR projects furthermore allow the users or participants to affect the virtual by means of hand gestures. Such gestures can mimic the movements one would make to affect an actual physical object, but they can also allow for new and additional forms of interactions that are not possible in the real world. Such ideas are, for instance, realized in the work by Hürst and Van Wezel (2013). They allow users to interact with virtual objects that appear in the real world when the scene is viewed through a mobile's screen. In addition to viewing the objects, users can, for instance, scale small virtual objects up and down by approaching the object and then increasing and decreasing the distance between two fingers. On the one hand, such interactions would be impossible with most physical objects, and in this sense, can be considered to suspend physical laws. On the other hand, we would likely make a similar gesture to transform a real rubber band. Also, we are quite used to making similar gestures to scale digital documents on the screens of touch-screen devices. In this sense, the interaction can be considered to mimic interactions we know well from the digital domain rather than from the physical world.

Another yet different form of 'imaginative interaction' is part of the iOS application *Konstruct* (see Alliban, n.d.). Here, real sounds create virtual objects in the space. The resulting three-dimensional sculptures can be viewed in the real environment through the screen of a mobile device. In the case of the Konstruct app, the underlying idea is that a user produces these sounds themselves, for instance, by speaking whistling into the microphone. However, the same mechanism of sound seemingly creating matter can of course also be triggered by any other sound in the environment.

It often is difficult to say whether interactive behavior mimics the real world, or takes a new imaginative form. This is especially difficult

when it comes to behavioral interactions, that not necessarily contradict any laws of nature. A form of interaction that, in our opinion, falls somewhere in the area in between imaginative and imitative is part of the mobile AR game *GeoBoid* by Lindeman, G. Lee, et al. (2012). In their game, players are surrounded by flocks of virtual geometric creatures called GeoBoids. These creatures are represented both visually as well as by means of spatialized audio using the player's phone. As mentioned before, players can catch the birds by making a swiping gesture on the phone's touch-screen. However, in addition, the game also allows for sound-based interaction with the birds. Players can scare the flock by whistling at a certain pitch and for a certain duration. Whereas the idea of scaring animals by means of sound is certainly something we know from the real world, the idea that a certain pitch has to be held for a certain amount of time is still quite different from how we scare real animals.

Finally, the popular game *Pokémon Go* gives us another reason to believe that imaginative interactions can be believable. In this game, players can catch Pokémon by throwing so-called Poké Balls at them. When a ball hits the creature, the creature "magically" appears to be captured inside of it (some argue their matter is transformed into pure energy). Although these interactions take place between two virtual objects, we can easily imagine similar scenarios where one of the two objects (e.g., the ball) would be a physical object.

As these examples show, interactions between the virtual and the real can differ quite a lot from the interactions we encounter in a purely physical world. This can be explained by the fact that virtual objects do not have to follow physical laws. As a consequence, they can, for instance, appear out of nothing, change their size, color, shape, teleport, or disappear entirely. If such actions are linked to actions of a real participant or physical object, this facilitates new forms of interaction that have no equivalent in a physical world.

5.1.3 Preliminary Insights

Our theoretical exploration provides us with preliminary answers to our key questions: Both imitative and imaginative interactions between the virtual and the real in AR are possible. However, judging from largely technology-focused descriptions, it is difficult to tell whether or to what degree they are also experienced as believable.

With respect to imitating real-world influences and interactions, much work focuses on creating realistic optical effects, such as shadows and reflections. Furthermore, some work addresses realistic collisions and other dynamic influences between virtual and real objects. For instance, several projects show that basic real-world interactions, such as a virtual ball that bounces on a physical object, can be simulated (e.g., Chae and Ko, 2008; S. Kim et al., 2011; Valentini and Pez-

zuti, 2010). Interactions between participants and virtual content also often mimic interactions that we know from the real world. This is, for instance, the case when we push a virtual puck with a real mallet when playing Airhockey (Ohshima et al., 1998), when we fire a virtual bullet by pulling the trigger on a physical (toy) gun (Piekarski and Thomas, 2002) or when we carry a virtual spider on our hand (Corbett-Davies, Dünser, Green, et al., 2013).

A main concern when it comes to imitating real-world influences is to make real objects react to virtual objects. Paradoxically, making virtual objects react to the real environment in a realistic manner can cause real objects to seemingly behave unrealistically (see S. Kim et al., 2011). If, e.g., a virtual ball hits a real bowling pin and changes its course, the real bowling pin might seem to behave weirdly if it is not affected by the collision at all. A key question is thus how virtual objects can affect the real world.

We have encountered a few projects that address this challenge and where virtual elements actually affect real objects. For instance, some projects extend real objects with electronics to make them 'movable' by virtual characters (Aoki et al., 2005; Kang and Woo, 2011). Furthermore, we have seen interfaces that allow virtual domino stones to knock over real stones (Leitner et al., 2008). These forms of interaction are interesting because here, virtual objects cause a real, physical change in the world.

In addition to projects where virtual objects *actually* affect the real world, we have also encountered various scenarios where virtual objects only *seemingly* affect the real world. This is, for instance, the case when a virtual object appears to cast a shadow onto the real world, while in reality, the real world remains unaffected.

When it comes to realizing new forms of influences and interactions, most existing projects concern interactions between a participant and virtual content. For instance, we have encountered projects where users can make sounds to create virtual visual 3D sculptures (Alliban, n.d.), where participants can make virtual objects disappear by shaking a corresponding physical object (Billinghurst, Kato, and Myojin, 2009) and where users can resize virtual objects by making gestures with their fingers (Hürst and Van Wezel, 2013). In addition, we have seen projects that allow participants to interact with virtual elements using digital touch-screens (Lindeman, G. Lee, et al., 2012), or virtual controls (Schmalstieg et al., 2002). These projects build on interactions we know from the digital domain, rather than from the physical world. In their entirety, these projects illustrate that various forms of interaction that differ from how we interact with physical objects are possible. However, the question of whether these interactions also are believable has received little explicit attention so far.

We have encountered examples that focus on interactions between the real world and virtual content as well as examples that specifically

focus on the interaction between a participant and virtual content. As many projects show, the two are closely intertwined. For instance, both the car racing game as well as the forklifting game by Namee et al. (2010) involve a participant that interacts with a real object that, in turn, interacts with virtual objects in the environment. Likewise, the artwork *Beyond Pages* involves interaction between a virtual light switch and a real lamp, but also a participant who interacts with the virtual switch. Furthermore, interactions between virtual and real objects can facilitate interaction between a participant and virtual content. We have seen this, e.g., in the project by Kang and Woo (2011), where one can play with a virtual character that pushes or pulls a little toy cart by interaction with the cart. We hence believe it makes sense to consider the participant, the virtual content and the physical aspects of the environment as an integrated whole.[7]

[7] A question that arises in this context is whether the interaction between a participant and something real constitutes AR if there is no real environment. In line with our definition in chapter 3, we see the real environment as an important element of an AR experience.

Our theoretical exploration also reveals some gaps in existing research. Little work seems to incorporate imaginative influences between virtual and real objects. Another area that has received almost no attention so far are non-visual or multimodal aspects of interaction between the virtual and the real. We believe it is important to consider all senses because we also foresee some non-visual responses when virtual and real objects interact. For instance, we might expect to hear sounds if virtual raindrops hit the window.

As we have seen, some researchers consider interaction between the virtual and the real not simply a *possibility* but rather a *necessity*. We, too, believe that for virtual and real objects to appear as if they existed in the same space, they should be able to affect each other. For instance, we would expect a real window to break when it is hit by a virtual ball and expect a virtual creature to get wet when it rains. However, in contrast with some existing views, we are not convinced that such interactions always have to mimic real-world interactions—rather we are inspired to explore other possibilities as well.

Although our review has provided us with preliminary answers, many questions remain. In particular, little work has addressed the issue of how virtual objects can affect the real world when imitating physical interactions. Furthermore, few imaginative interactions between virtual content and the real world have been realized. We will address both topics as part of our practical exploration.

5.2 *Practical Exploration*

In order to explore if and how the virtual and real can interact, we conduct a small series of experiments. We divide this exploration into two main categories: (1) Imitative interactions, which focuses on simulating real-world interactions and (2) Imaginative Interactions, which focuses on influences that have no equivalent in a solely physical world but ideally, are believable nonetheless. In both of the two categories,

we present three explorations.

As mentioned, the discussion of the practical exploration is solely based on our own experiences with the different scenarios and does not include any empirical research with participants.

Setup

Unless specified otherwise, our setup used for the exploration is based on relatively cheap conventional office equipment rather than dedicated AR technology. On the hardware side, our setup consists of a monitor, a computer, loudspeakers, and a webcam that provides a live-view of the environment. On the software side, the project uses self-written Max/MSP/Jitter software (see `https://cycling74.com/products/max/`), which makes use of Max's built-in physics engine. The software is used to integrate virtual objects into the view of the real environment. As a result, participants can see virtual objects in the real on-screen environment on the monitor as well as potentially hear virtual objects via the speakers. Unlike typical AR setups, the setup is fixed, and only shows the augmented environment form one static point of view, namely the fixed position of the webcam. Whereas many existing AR projects are interested in exploring the technological possibilities, we want to explore the *conceptual* possibilities. For this goal, this simple setup is sufficient.

5.2.1 *Imitative Interactions*

The first experiments explore whether and to what degree virtual and real objects can (appear to) physically interact like real objects.

Exploration 1: Bouncing Ball

Our first simulation recreates an arguably simple real-world interaction between two objects: a ball that bounces on a surface. As mentioned, similar experiments have been conducted by Valentini and Pezzuti (2010), S. Kim et al. (2011) and Chae and Ko (2008). We hope to validate that this type of interaction indeed is possible as well as believable.

In order for the ball to react to its real surroundings, we have created a virtual reconstruction of the environment in our self-written software and aligned it with the real scenery. (This means that the real desk has a virtual desk as a counterpart. The virtual desk is invisible, but it is positioned at the exact location of the real desk.) Furthermore, we have assigned virtual physical properties such as mass and restitution to the virtual elements and applied gravitational forces (using the Max/MSP/Jitter physics engine).

When we start the simulation and view the environment through the screen, a virtual ball appears to bounce on the desk in front of us (see figure 5.4). On first sight, the experiment appears to be a success:

it is possible to simulate certain existing real-world interactions in AR. However, on second thoughts, it becomes clear that the virtual ball does not have any effect upon the real table. If a real ball hits a table, the collision also goes hand in hand with a distinct sound. However, the simulation remains silent. The example only shows that the real can influence the virtual and leaves us wondering whether and how the virtual can influence the real as well.

In order to overcome this limitation, we adapt the experiment and extend the physical desk with virtual sounds. For this, we use recorded sound samples of a real ball that hits the desk. These samples are consequently triggered when the virtual ball and the virtual representation of the table collide. Every time the ball hits the table, one of several recorded sounds is randomly chosen and played back. We furthermore use the magnitude of the collision to calculate the volume of the playback.

Unfortunately, we have to admit that the result is not convincing (to us) yet. Maybe, because the result is almost but not entirely realistic, we are irritated by the fact that something 'is a bit off'.[8] In our experience, the condition with no sound was more believable than a sound that is not exactly what one would expect.

Of course, the fact that we are not able to create a realistic sound response does not show that this is impossible. If anything, we still expect this can be achieved by more carefully considering what sounds would match the visual impression.

Irrespective of the unconvincing result, the example shows that sometimes, we need to extend real objects (the table) by means of non-visual virtual content (the sound samples) in order to simulate reality. This is not surprising: reality is not just something we see. Consequently, we cannot simulate it realistically while only considering visual aspects.

With respect to real-world applications, it would be desirable to not require pre-recorded sounds. A solution that could be explored in the future is using physical modeling of sound (Cook, 2002) to extend physical objects with virtual sounds.

[8] Similar experiences have been described with respect to almost but not entirely human-like robots and in the context of the so-called uncanny valley effect (e.g., Seyama and Nagayama, 2007). However, as this effect is associated with human-likeness, it would be rather surprising to find something similar when it comes to simple physical objects.

Figure 5.4: A virtual ball is bouncing on a real table. Four snapshots from the live-view. Image © Hanna Schraffenberger and Edwin van der Heide.

Exploration 2: Falling Dominoes

As a next step, we deliberately chose a scenario that seems bound to fail. In this experiment, a virtual ball collides with a row of real dominoes (see figure 5.5). Like in the previous example, the desk has a virtual counterpart, which allows the virtual ball to roll over its surface. Likewise, there is a virtual representation of the first domino stone.

Initially, the screen shows the virtual ball heading towards the real dominoes. Just like in the previous example, the virtual ball reacts to the real: the moment it hits the first stone, the ball changes its direction and rolls back. However, unlike in the real world, the stones do not fall.

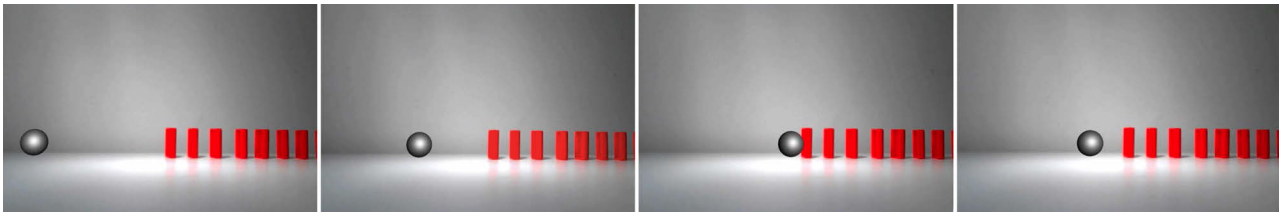

Figure 5.5: A virtual ball approaches a row of dominoes (frame 1 and 2), hits the first stone (frame 3) and consequently changes its course (frame 4). Unlike in the real world, the stones do not fall. Image © Hanna Schraffenberger and Edwin van der Heide.

Based on our own impression, this behavior contradicts our expectations and is not believable at all. Neither the behavior of the virtual nor the behavior of the real elements appears to be credible. We seem to expect a realistic response and want to see the stones falling. Clearly, problems can arise due to the fact that the virtual cannot directly affect the real. As mentioned, this problem has been identified earlier by S. Kim et al. (2011) who call this the 'physical artifact'.

We hope to overcome this limitation by making the virtual object affect the real world. As we have seen in our review of existing work, there are several possible ways in which the virtual can affect the real. For instance, we have seen the use of physical portals that can enable interactions between virtual and real dominoes (Leitner et al., 2008). However, we have another idea: we introduce a physical counterpart that extends the virtual object.

Just like a virtual desk has been acting as a virtual counterpart for the real desk in previous examples, this time, a real ball acts as the physical counterpart for the virtual ball. This is realized by analyzing the camera-image and replacing the real ball with a virtual one. As a result, while looking at the scenery directly, one sees a physical ball. The screen, however, shows a virtual ball instead.

In our own experience, this approach is interesting on three levels. First of all, the setup allows the participant (and in this case, author) to tangibly, intuitively and naturally interact with the object, and roll it towards the dominoes.

Secondly, the result simply looks fascinating: when the virtual ball hits the first domino stone, the stones start falling (see figure 5.6). Although we expect the stones to fall on some level (see exploration 2), we also expect them not to fall on another level. After all, a virtual ball can typically not apply real forces to physical objects. Based on our own experience, we believe part of the fascination is based on the fact that we can see that the ball is not real and yet, affects the real world. We assume that things would look less intriguing, had we created a virtual ball that looks so real that it can be mistaken for a real ball.[9]

[9] From a technological perspective, this example is "nothing new". Like many AR applications, a virtual object (here, a virtual ball) is superimposed onto a camera feed and spatially aligned with a real object (here, another ball). However, on an experience level the result is quite fascinating—likely because we are not used to seeing virtual objects physically affect the real world.

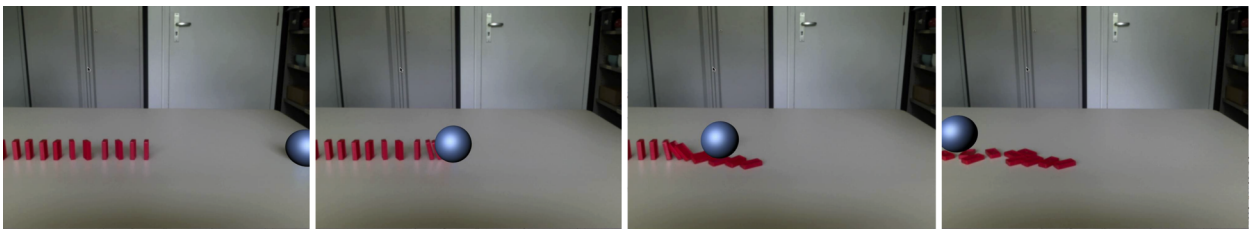

Figure 5.6: A virtual ball approaches a row of dominoes (frame 1) and hits the first stone (frame 2). This causes all domino stones to fall (frame 3 and 4). Image © Hanna Schraffenberger and Edwin van der Heide.

Third, the quality of the execution in terms of visual and spatial realism is rather low—yet we experience a strong sense that the virtual object is present in the real environment. The virtual ball looks a bit too big and, at times, appears at somewhat weird positions in the image (in front rather than behind of the real stones). In our experience, we get a strong sense that the ball is present in the real environment even though it sometimes does not look like it if we analyze the scene visually. In our experience, the interaction between the virtual and the real is such a strong perceptual cue that the virtual object is part of the environment and the contradicting visual information does not seem to matter much.

What makes this approach particularly worthwhile is that the virtual ball can potentially also display behavior that purely physical ball cannot exhibit. For instance, it can change its texture according to its speed, or change its color when it hits a real object. What is more, we can also replace the real ball with any other virtual object and thereby give it qualities that the real object does not have. This approach is found in the earlier mentioned product *Sphero* (2011). Here, a robot ball that is turned into a virtual beaver when it is viewed with a dedicated smartphone app (cf. J. Carroll and Polo, 2013).

Of course, the idea of extending a virtual sphere with a real sphere and the idea of extending a real ball with a virtual sphere are interchangeable. To make this point explicit, we have also applied this concept to a Newton's cradle (see figure 5.7. In our setup, one of the cradle's real spheres is augmented with a virtual counterpart. As a result, it looks like a virtual ball in the resulting digital view of the environment.

Judging from our own experience, watching this Newton's cradle

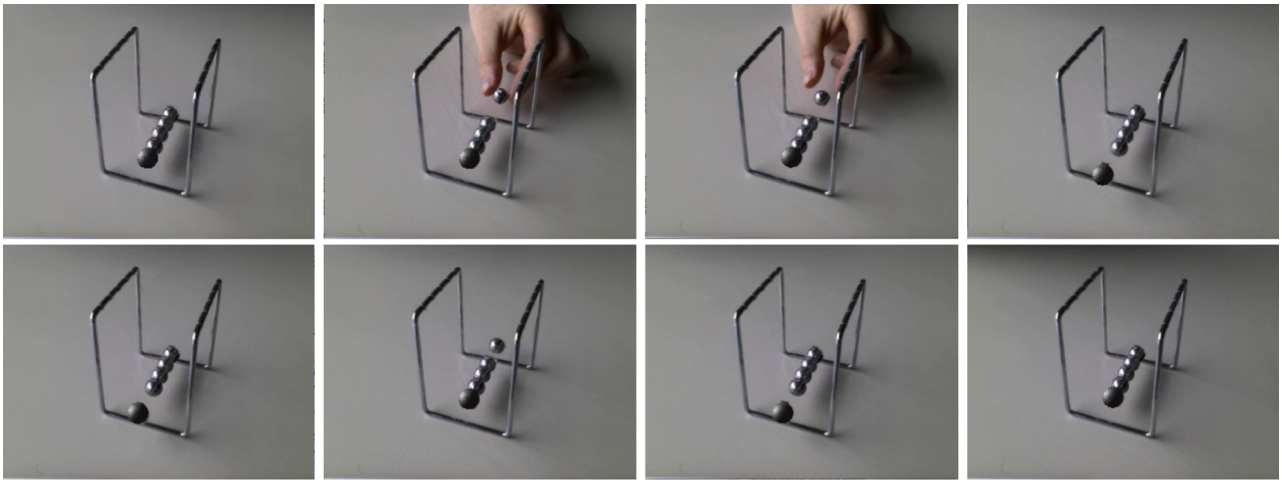

Figure 5.7: One of the physical spheres in a Newton's cradle is overlaid with a virtual sphere. The setup results in a realistic interaction between the spheres and allows for participant interaction with all spheres. Image © Hanna Schraffenberger and Edwin van der Heide.

swing is quite irresistible, maybe because seeing a virtual ball physically interact with the real balls contradict our expectations ("a virtual ball cannot push a real ball") in a pleasant way. Like in the case of the domino stones, we believe the interaction is even more fascinating because we are aware that one of the spheres is virtual. Also here, we expect the resulting image to be much less interesting in cases where the virtual ball is mistaken for a real ball.

Exploration 3: A Door That Opens for Virtual Objects

With this last exploration, we move from physical-like forces to behavioral interactions (section 4.10). The underlying idea is that real objects in the environment can sense and react to the behavior of virtual objects. We explore this idea in the context of a real automatic door that opens when it is approached by a virtual object (see figure 5.8).

This exploration uses a somewhat different setup than the previous examples. We use a steerable toy that approaches the door. Due to sensors above the door, the door opens when the toy moves in front of it. In order to turn the physical toy into a virtual object, an AR marker is placed on top of the toy. A combination of the Unity game engine and the Vuforia AR SDK is used to detect the marker in a live webcam feed of the environment. When the marker is detected, it is replaced by a virtual sphere. On the laptop screen, this makes it look as if a virtual ball was rolling towards the real door. The result works as expected: we see a virtual sphere that approaches a real door, and a real door that opens just like it would open for a real object or person.

Of course, our implementation of this scenario is not ideal. To realize this concept, we do not need a physical object or toy to open the door. Ideally, a virtual object could be made to appear in the environment, and if it approaches the door, a signal that opens the door

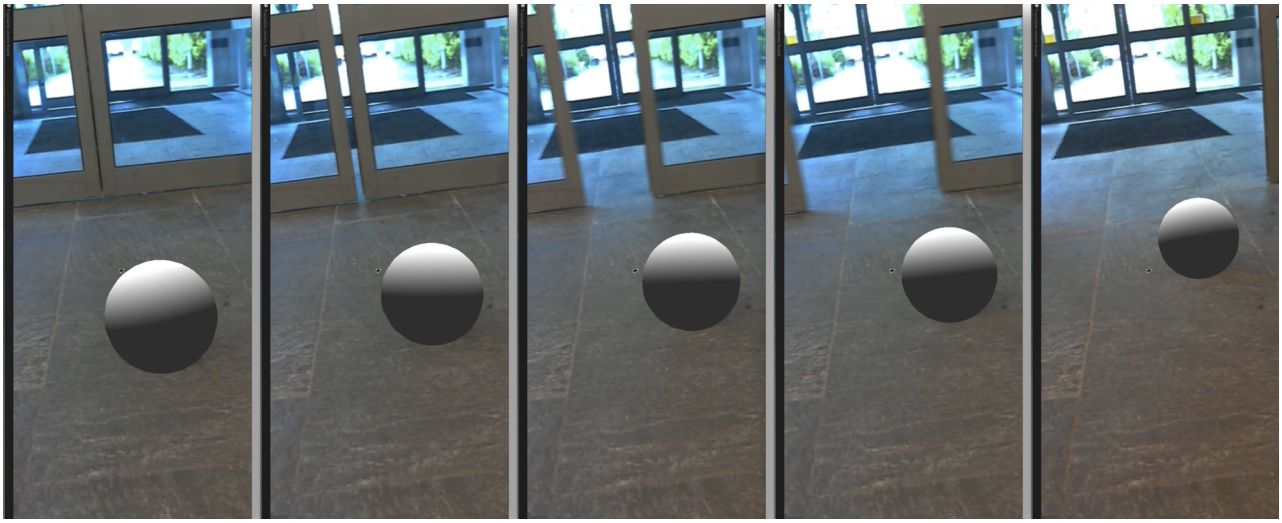

Figure 5.8: A real door opens for a virtual sphere. Image © Hanna Schraffenberger and Edwin van der Heide.

could be sent to the door. After all, the motion sensor above the door also sends such a signal to the door when a real person approaches it. Ideally, the virtual object would also react to the door, and e.g., stop and wait until it is open until it moves on.

In our opinion, this example is interesting because it makes use of a real object that senses the world and reacts to what it senses. In our experience, the fact that the real object seemingly senses the virtual object contributes to the feeling that the virtual object is actually part of the environment. Simply put: if the door sees the object, it has to be there!

We can imagine similar interactions between virtual objects and other real objects that are controlled by signals. As suggested by Barakonyi et al. (2004), virtual objects could, e.g., interact with printers, digital instruments and interactive robots as all of these objects can be queried for status information and controlled with commands.

5.2.2 *Imaginative Interactions*

The preceding imitative explorations have shown that in AR, at least some real-world interactions can be simulated between virtual and real objects. In the following, we explore influences between the virtual and the real that do not mimic reality, but instead, try to bring new imaginative and physically impossible forms of actions and reactions to the real world.

As discussed, the virtual does not have to obey physical laws (cf. Sutherland, 1965). Hence, it can behave in novel and *physically impossible* but—according to our hypothesis—*nevertheless believable* ways. Unfortunately, our review has revealed little examples to support this point. However, just as we accept imaginary objects with their own behaviors in books, computer games and movies, we expect that we

can accept different sorts of objects and different forms of behavior in AR.

Whereas interactions based on imitating the real world present the challenge of making real objects behave realistically, imagination-based relationships between the virtual and the real not necessarily entail such a challenge. If we do not mimic specific realistic interactions, it might not be a problem if the real world has an effect on the virtual but is not affected in return. Accordingly, we believe an interesting direction to explore is the possibility of realizing simple one-directional influences rather than complex two-way interactions.

One challenge when it comes to exploring imaginative influences is the vast realm of possibilities. We can, for instance, easily imagine a virtual ghost that—unlike any real object—floats through real walls.[10] Likewise, we can envision virtual objects that change their size, color or shape when they collide with a real object or teleport to another position when they hit something real.[11] Furthermore, we can take inspirations from books and movies and, for instance, envision a virtual character that, unlike Pinokio, grows a longer nose whenever it hears a lie. Ultimately, the range of possible influences is only limited by our own imagination.

[10] One could argue that this is a lack of influences rather than a novel form of interaction. However, as real objects cannot move through real walls, we believe this is an exciting possibility to explore—no-matter whether one terms it interaction or not.

[11] For instance, the virtual object could 'take over' the color of the real object.

Of course, we cannot explore all possibilities at once. In the following, we explore this idea of physically impossible but nonetheless believable interactions in the context of *attractions* between virtual and real objects. The combination of virtual and real objects allows us to create new forms of 'magnetism' or attractive forces that cannot exist in a purely physical world. The choice to focus on attractive forces is simply based on the curiosity of the author, who always considered magnetism a fascinating phenomenon in the real world and who is curious to explore variations of this phenomenon in the context of a mixed virtual-real environment. In the future, the idea of inventing new laws can be explored in many other contexts. Whereas the virtual objects presented in the sections below still have quite some similarities with real-world objects (e.g., with physical marbles), chapter 6 takes the idea of imaginative AR and new forms of virtual objects even further.

As mentioned, the implemented attractions solely establish *influences* between real and virtual elements rather than some form of *interaction*. However, a participant can react to the behavior of virtual objects as well as affect their behavior through interactions with the real environment. In this sense, the virtual, the participant and the real world interact with one another. This is why we title the section *Imaginative Interactions* and place this idea under the umbrella of interactions between the virtual and the real.

Exploration 1: Attractive Colors

The first scenario explores a new and imagined way in which virtual objects can react to the real environment. In this experiment, virtual objects are attracted by real objects that have a similar color. (However, real objects are not attracted by the virtual objects in return.)

The concrete setup to explore this idea includes a cloud of small virtual spheres—half of them with bluish colors and the other half with yellowish colors. Furthermore, it includes one blue and one yellow rubber ball. As soon as the rubber balls enter the scene, they attract spheres of similar colors (see figure 5.9). According to our own impression, the relationship between the virtual and the real is easily understood, intuitive and believable although the virtual spheres do not imitate the behavior of real objects and do not obey the same physical laws as real objects. As expected, the fact that the virtual has no influence on the real is not a problem.[12]

[12] The virtual objects appear rather light in relation to the physical objects. This means that this lack of influence might be experienced as realistic.

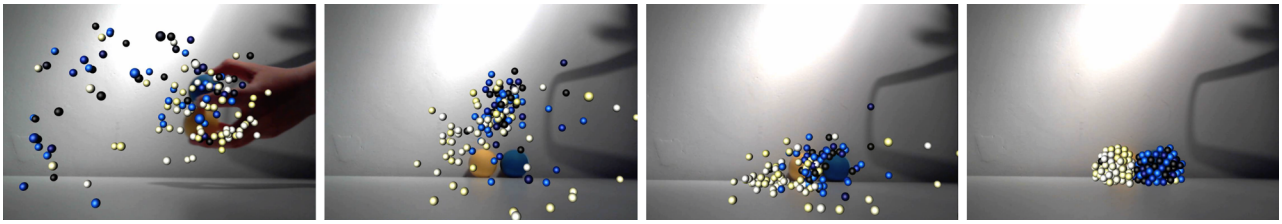

Figure 5.9: The colored rubber balls attract virtual spheres of similar colors. Image © Hanna Schraffenberger and Edwin van der Heide.

Even though the example is based on physical forces (attraction), it leaves behind the realm of realistic physical interaction as we know it. It becomes clear that the virtual does not have to behave like a real object in order to be believable.

Furthermore, the example shows that influences between the virtual and the real bring great possibilities for interaction between a participant and the virtual content: If physical objects influence virtual ones, the participant can interact with the virtual elements simply by interacting with physical objects.

Although imaginative interactions play an important role, the project also includes some imitative interactions. For instance, the virtual spheres roll over the table like real objects and block each other's way like real spheres do.

A nice aspect of this example is that the real balls can interact with each other, and, e.g., collide. These interactions between the real spheres, in turn, affect the virtual spheres that also interact. If the participant also starts to intervene, this setup can lead to complex chains of cause-and-effect that involve influences (1) between the two real spheres, (2) between real spheres and virtual spheres, (3) between virtual spheres, (4) between the participant and the real spheres, and indirectly, (5) between the participant and virtual spheres.

It should be noted that although this example aims at exploring

potential color-based relationships, the virtual and real objects also shared the same shape. This might have an influence on our experience and, e.g., cause the scenario to appear more convincing. In the future, it might be interesting to look at the attraction between similar colors and similar shapes independently.

Exploration 2: Attractive Light

As a next step, we explore a setting where small randomly colored virtual spheres are attracted by real light. A small desk lamp serves as a light source that can be turned on and off. This allows a participant to interact with the virtual spheres by interacting with the physical lamp (see figure 5.10).

In the default state, the virtual spheres are placed on the desk. However, when a participant turns on the light, the virtual spheres move upwards, as if they were pulled up by the light. If the light remains on, the spheres cluster around (in front of) the light source.

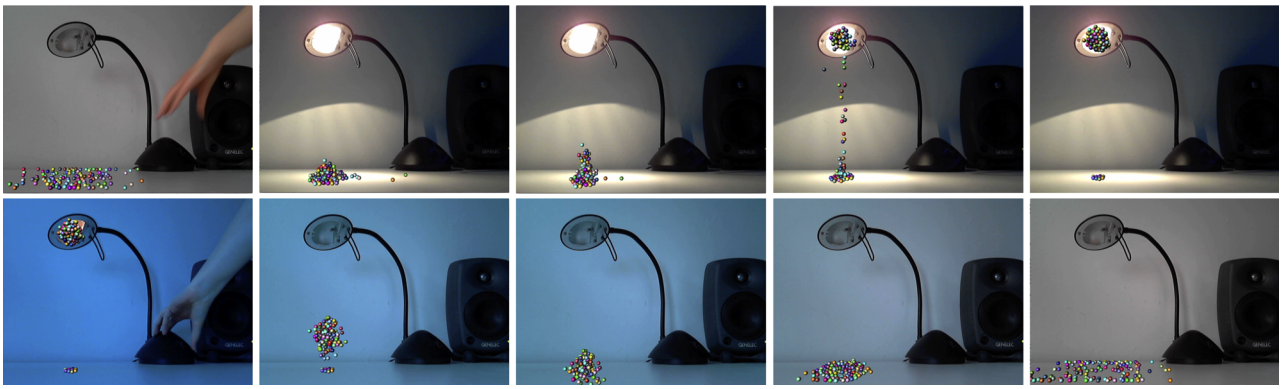

Figure 5.10: Virtual colored spheres are attracted by light. The top row shows what happens when the lamp is turned on (from left to right) and the bottom row shows what happens when the light is turned off. Image © Hanna Schraffenberger and Edwin van der Heide.

Admittedly, this project has some weak spots in its execution. First of all, when the virtual objects are pulled up, they can gain so much momentum that they 'shoot over the target', and momentarily move above the lamp, when one would actually expect a collision with the lamp. The fact that we expect a collision shows that we would actually like to see some imitative real-world interactions, even if the virtual objects display imaginative behavior. Secondly, the automatic color adjustment of the webcam feed is activated by the light changes, which causes the image of the environment to undergo weird color shading changes (e.g., the view becomes blueish when the light is turned off). Yet, in our opinion, the interaction between the participant and the lamp and the consequent response of virtual objects to the light, create a very strong sense of virtual objects being present in the space.

In our opinion, this scenario can be considered a gray area between both realistic and non-realistic interactions. On the one hand, simple physical objects like marbles are not attracted by and fly towards

the light. On the other hand, the experience is supported by imitative interactions: the virtual spheres appear to collide with and roll over the desk like real spheres. Furthermore, we know similar (yet not identical) behavior from animals, such as moths who are attracted by light. It thus should be noted that the virtual spheres look like physical objects—not animals. Their appearance implies no creature like qualities. For instance, they do not have any apparent senses that would allow them to register the light. In addition, they also have no apparent way to move or act in the environment based on their own intentions. Instead, they look like tiny physical spheres. They roll over the real table as if they had a mass and they collide and interact with one another like physical spheres would. Yet, the resulting movement of the spheres occasionally reminds us of little creatures. In particular, when they are clustered in front of the light, they remind us of animals in a huddle that all try to achieve the same thing: getting close to the light. In the future, it would be interesting to explore when the implementation of imaginative laws leads to movements and actions that are interpreted as an active, goal-oriented *behavior* rather than physical cause-and-effect relationship.[13].

In the future, this example could be extended with sounds (especially when the spheres and the desk collide). A question that remains open is if such sounds would have to mimic the sound of small spheres dropping on a desk in order to be believable.

We can imagine this light-based form of interaction as a very intuitive and easy-to-learn way for participants to interact with virtual content. For instance, a participant could use a flashlight to attract virtual objects, move them around, and turn off the flashlight in order to let go of them.

[13] Similar topics have been explored in the context of cybernetics for a long time. For instance, in the early 1950s, Grey Walter has presented so-called *Machina speculatrix*. These simple mobile robotic vehicles were equipped with two miniature radio tubes, a touch- and a light-sensor and two motors, and appeared to exhibit the "exploratory, speculative behavior that is so characteristic of most animals" (Walter, 1950, p. 43).

Exploration 3: Magic Hands

This small project focuses on interaction between a participant and virtual content and addresses the issue of how participants can move virtual objects in space. Like previous examples, the project builds on the idea of new forms of attractions: Here, the hands of the participant (author) are tracked with a Microsoft Kinect camera and programmed to attract virtual objects. Unlike in previous examples, the computer monitor serves as 'magic mirror'. To see themselves interact with the virtual objects, the participant has to look at the monitor. The monitor shows a mirrored view of the environment that also contains the virtual balls.

To explore the ideas of having 'magnetic hands' that attract virtual objects, a large number of approximately tennis-ball sized virtual spheres are placed in the participant's/author's room. If the participant moves the hands close to the balls, they move towards and stick to the hands. By quickly moving the hands away, it is possible to let go of the balls. The screenshots of the resulting view in figure 5.11 show

this effect. In addition to sticking to the hands of the author, the balls also seemingly defy physical laws in the sense that they move through the space in slow motion.

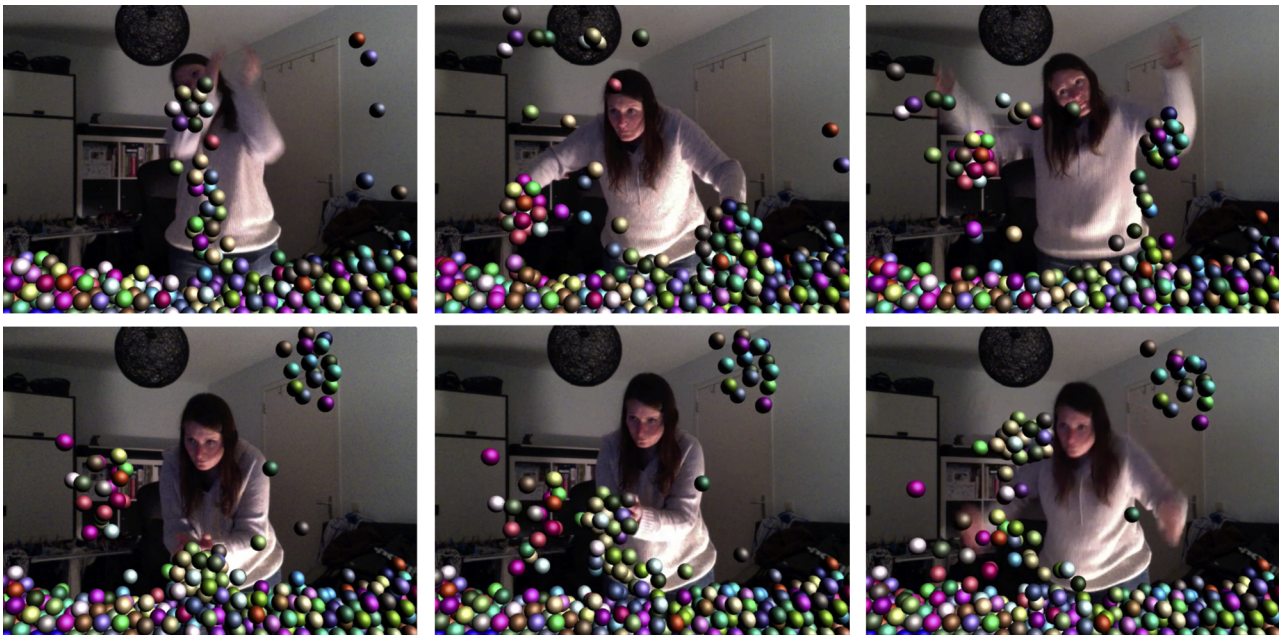

Figure 5.11: Virtual colored spheres are attracted by the hands of the author/participant. Quickly moving the hands away from the virtual balls allows one to let go off them. Image © Hanna Schraffenberger and Edwin van der Heide.

According to our experience, playing around with this form of interaction is extremely intuitive and the setup is very enjoyable. One of the observations that stuck with the author, is that the author intuitively adapted her movement, and started to largely move in slow-motion (however, using short and sudden movements to let go of the spheres). Seeing herself on the screen, she felt a strong urge to move in a believable way, and not, e.g., create motions that contradict the characteristic movement of the virtual spheres. This, e.g., also included holding the hands in a way that they looked properly covered by the spheres. It would be interesting to see if the same urge to behave believably and the resulting slow-motion movement is also found in other people.

We believe that 'magic hands' can serve as an easy means to facilitate interaction between virtual and real objects. For instance, one hand could be made to attract virtual objects whereas the other hand could be made to repel them. We expect this to allow for a lot of basic interactions, such as picking up and moving virtual objects in space.

5.3 *Discussion, Conclusions and Future Directions*

We have asked whether interaction between virtual content and the real world can imitate interactions that we know from the physical realm as well as take on new imaginative forms that have no equivalent

in a purely physical world. Our theoretical and practical exploration has shown that both is possible but also challenging.

When it comes to realizing realistic interactions, we are facing a surprising dilemma: by making virtual objects behave like real objects, actual objects in the environment behave unrealistically. As S. Kim et al. (2011, cf.) show, this can, for instance, happen when a virtual ball collides with a real paper cup without causing this cup to move. A key question that arises in this context is how virtual objects would be able to affect the real world. We have identified four key ways to address this challenge and allow virtual objects to affect the real world.

First of all, the virtual can affect the real if the real world is extended by means of electronics such as actuators. The virtual can then move or transform the real by controlling these actuators. Likewise, one can modify those real objects that already are equipped with electronics and that already react to the environment. For example, we have suggested that automatic doors can be modified to also open when something virtual approaches them.

Second, we can extend physical objects with *virtual* qualities. This can make it seem *as if* real objects reacted to virtual objects. We have proposed a simple example where a real desk is extended by means of virtual sounds and thereby, can react (resonate) when it is hit by a virtual object. This idea could be taken to the extreme: For instance, if a virtual ball would hit a real window, we could make it look as if the window were broken, and if a virtual ball would hit a real cup, we could make it look as if the cup were moving. Of course, it is questionable if we should call a real window that appears to be broken, or a real cup that appears to move while it actually stands still, *real*. We propose to call such objects that unite both virtual and real qualities *augmented objects*. (The concept of augmented objects will be discussed in more detail below.)

Third, just like real objects can be extended with virtual qualities, virtual objects can be extended with physical properties. This, too, results in an *augmented object*. We have shown this in the case of a virtual sphere that we have extended with a physical ball.

Fourth, we can present virtual content in the real world directly and thereby utilize real-world interactions. For instance, we can play back the sounds of a virtual creature on loudspeakers. In this case, the sounds actually resonate in the environment. Likewise, the light of a virtual sun might be projected onto a wall directly, in which case the light from the projector actually lights up (and maybe even warms up) the space. In such cases, the boundaries between the virtual and the real blur: unlike the virtual creature and the virtual sun, the light and sounds actually exist in the environment, and thus, will interact with it.

While we have shown that that virtual objects can affect the real world, it also has become clear that this is not always necessary. We

also can establish believable influences where the real affects the virtual but where the virtual does not affect the real in return. For instance, we have created an imaginative scenario where real objects attract virtual objects that have a similar color. Judging from our experience, this is convincing even if the real world is not affected by the virtual spheres.

Our review and exploration have shown that interaction between the virtual and the real can take place both in the virtual domain as well as in the physical domain. For instance, the voice of a virtual character can be played back in a church on a loudspeaker. In such a case, it is *actually* affected by the characteristics of the surrounding space and the interaction thus takes place in the physical domain. However, we might also play back the sound on headphones and apply virtual reverb to the voice. This would only make it seem *as if* the voice were affected by the church and the interaction would thus take place in the virtual realm. Both forms can be desirable. For instance, when shooting virtual bullets, the bullets should arguably only *seemingly* have an effect on the world. On the other hand, when playing bowling with a virtual ball and real pins, it might be quite entertaining if the real pins would actually fall down upon being hit. When designing AR environments, we thus have to carefully consider these possibilities.

As discussed in chapter 3, a common goal in AR is to make it seem as if virtual objects were a part of the real environment. Judging from our own experience, interactions between the virtual and the real can play a crucial role in making it seem as if a virtual object were actually present in the real world. For instance, the fact that an automatic door opens for a virtual object can reaffirm our impression that the object is part of the space. Likewise, virtual objects that stick to our hands can give us the feeling that they are present in space 'with us'. In the future, it would be interesting to further explore the effect of such interactions on the experienced presence of virtual objects in real space. We expect that the lack of interactions between the virtual and the real can also harm the illusion of virtual objects being part of the real environment. If, for instance, virtual rain does not cause the real world to get wet, we might not feel like the rain is a part of the environment.

As our review has shown, it is sometimes assumed that virtual objects should obey to physical laws and interact with the world in the same manner as real objects. Judging from our current exploration, this is not always necessary. In our opinion, what matters is not whether something can exist in the real world but whether it is believable.

5.3.1 *Augmented Objects*

As mentioned above, we have proposed to extend virtual objects with physical properties and to extend physical objects with virtual proper-

ties and have coined the result *augmented objects*.[14] We believe that augmented objects can play an important role in facilitating interactions between the virtual and the real. Augmented objects share qualities with both virtual and real objects and consequently, can interact with both virtual elements and the real world. Our exploration has shown that augmented objects can be used to apply forces to the real world just like real objects. Furthermore, we expect that augmented objects can change their visual appearance, such as color, texture, size and to some degree shape just like virtual objects. Similarly, they can likely create sounds that differ from the sound of physical objects. However, at the same time, augmented objects also lack many of the possibilities of virtual objects: An augmented object can, e.g., not hover through space, teleport or move through walls. It might also cause a natural sound that cannot easily be 'removed' from the environment. Which of these qualities are considered advantages and which are considered disadvantages likely differs from project to project and depends on the actual context and goal of a project. In the future, it would be interesting to explore the possible manifestations of augmented objects more systematically.

The idea of working with *augmented objects* in AR is not new. In our theoretical exploration, we have encountered examples that extend physical objects with virtual qualities. For instance, Bau and Poupyrev (2012) have extended physical objects with virtual textures. Likewise, Sphero—the robot ball that can take the form of a virtual beaver—is based on the idea of extending a physical object with a virtual character. In addition, we have seen projects that extend virtual content with a physical dimension. So-called 'tangible interfaces' (Billinghurst, 2001; Billinghurst, Kato, and Poupyrev, 2008) link a physical and tangible element to virtual objects. However, to the best of our knowledge, the AR research community has only addressed such concepts in the context of interaction between a *participant* and virtual content. In contrast, we have also focused on interaction between the real general environment and virtual content. This has shown that augmented objects can interact with physical objects as well as allow participants to interact with them.

The concept of augmented objects—objects which combine virtual and physical qualities—also raises more fundamental questions about the nature of virtual content: One could argue that every perceivable virtual object is an augmented object. This is because every perceivable virtual object has some physical counterpart that allows us to sense the object. For instance, a virtual sun might be shown to a participant via pixels on an AR head-mounted display and virtual rain might be represented with the sound of raindrops. The sun and the rain thus have a physical representation in the form of light or sound. If we follow this line of thought, every virtual object is actually an augmented object and falls somewhere onto a spectrum between entirely real and

[14] An augmented object is similar to the hybrid objects defined in section 4.5. However, when it comes to hybrid objects, the virtual and the real complete one another and both serve as an integral part of the object. When it comes to augmented objects, the virtual adds something additional to a real object but not necessarily completes it. We hence see augmented objects as a more general umbrella that describes objects with both a virtual and a real component.

entirely virtual.

5.3.2 *Participant-Content Interaction*

In existing research, interaction between a *participant* and virtual content and interaction between virtual content and the real *environment* are often addressed independently. In this chapter, we have proposed to see the participant as a part of the real environment. In our opinion, this approach has been successful: our explorations have resulted in scenarios that facilitate both. We have seen that if virtual objects interact with the real world, participants can interact with the virtual content by interacting with the real environment. For instance, if virtual objects are attracted by light, a participant can interact with them by using a flashlight, closing the curtains or turning on the light. Likewise, if virtual objects are attracted by real objects that have the same color, a participant can interact with virtual objects by moving around real objects. This shows that by establishing influences between virtual objects and the real environment, we can also facilitate the often-desired interaction between a participant and virtual content. We hope to further explore the complex dynamics between a participant, their real surroundings and virtual elements in this space in the future.

5.3.3 *Future Directions*

When it comes to imaginative influences and interactions, a key issue is believability. Presumably, not everything that we can imagine is also credible. So far, our exploration has only taken the individual experience of the author into account. Of course, this experience might be biased and does not necessarily tell us anything about how other people would experience the different scenarios. In the future, it would be desirable to explore the different scenarios with other participants and address whether they are experienced as believable. Another topic that could be addressed in this context is the 'physical artifact' (S. Kim et al., 2011). More generally, it would be interesting to know how we experience real objects that appear to display an 'unrealistic' behavior. How do we feel when a real mirror does not reflect the virtual creatures that sits in front of it? How do we experience a real cup, when it does not move after being hit by a virtual object (cf. S. Kim et al., 2011)?

A more general goal for the future will be to understand what factors contribute to whether the interaction between virtual and real objects is perceived as credible. We assume that what is experienced as believable is closely linked to our expectations as well as to the appearance and behavior of a virtual object. Presumably, if a virtual object looks so realistic that we mistake it for a real object, we also expect it to also affect the real environment just like a real object. However, virtual objects do not have to look like real objects and in such cases, we

might not have such expectations. As mentioned in the introduction of this chapter, we expect that the believability of an object's behavior is linked to the situation and context in which it is presented. For instance, we expect different behaviors to be believable in a gaming context than in an educational context. A possible next step would be to place our proposed interactions in different contexts (e.g., as part of a game) and evaluate whether they are believable in different situations. More generally, future research can address whether and how the context in which a behavior takes place affects the believability of this behavior.

This chapter has focused on the behavior of virtual and real objects *with respect to one another*. In the future, it would be interesting to also address believable *behaviors* of virtual objects in the real world that are not related to any real object. For instance, it would be interesting to know whether we find it believable if virtual objects teleport, change their size, color, shape without any apparent real cause.[15]

So far, our exploration of imaginative influences has only addressed new forms of attraction between the real world and virtual content. However, many more forms of imaginative laws and behaviors can be explored. As mentioned, virtual objects could change their shape or color whenever they collide with a real object. Likewise, they could teleport to another position. It might also be interesting to play with the time-dimension, and e.g., rewind the movement of a virtual object after it collides with a real object. Future research can address whether such interactions are credible and meaningful.

Furthermore, our exploration has focused on simple physical objects so far. In the future, it would be interesting to explore the more specific case of virtual characters that have their own senses and pursue their own goals.

In the course of this exploration, we have focused on visual AR. In the future, it would be especially interesting to explore the possibilities of non-visual, multimodal and crossmodal interactions and to include other physical properties of the environment. For example, real wind might move virtual leaves and the temperature of the environment might affect the behavior of virtual creatures.

This chapter has focused on interactions between the virtual and the real. Although many virtual objects used in our explorations have reacted to the real world in new and imaginative ways, they have still imitated real objects and the physical world in many ways. In particular, the virtual spheres used in our experiments still *look* like a lot like spheres that actually can exist physically. This makes us wonder whether and to what degree virtual objects can differ from real objects. We will explore this question in the next chapter and attempt to create an object that does not look, feel or behave like any real object.

To conclude, we propose there are two worthwhile approaches to creating influences and interactions in AR: On the one hand, the imita-

[15] With respect to realistic and imaginative qualities of virtual objects, it would also be interesting to implement virtual versions of so-called "supernormal stimuli" (see, e.g., Staddon, 1975). A supernormal stimulus is an imitation of a real-world stimulus that differs substantially from the imitated stimulus and that is not encountered naturally in the real world. Supernormal stimuli have, e.g., been studied in the context of animal biology, where it has been demonstrated that animals can respond stronger to such (e.g., exaggerated) imitations than to the natural imitated stimuli. Augmented reality opens up many possibilities to design virtual supernormal stimuli and place them in the real environment. Studying participant responses to such stimuli might be particularly interesting in the context of biology, psychology and perception research.

tion of reality. On the other hand, the creation of imaginative realities. Whereas the first has gained a lot of attention, the latter has gained little attention so far. We hope that our research serves as a fundamental step towards new kinds of realities in which the virtual and the real interact in unique and creative ways.

6 The Invisible Cube: Introducing Sonically Tangible Objects

As the previous chapter has shown, augmented reality often mimics reality. Researchers strive for photorealism and aim for scenarios where virtual objects cause the same occlusions and shadows as physical objects (see, e.g., Gibson and Chalmers, 2003; Madsen, Jensen, and Andersen, 2006). Similarly, scientists include physics simulations to make virtual objects adhere to physical laws and move like real objects (e.g., Chae and Ko, 2008; S. Kim, Kim, and Lee, 2011). In line with this, many AR projects allow users to interact with virtual content in the same way as they would interact with real physical objects (e.g., Corbett-Davies, Dünser, Green, et al., 2013; Ohshima et al., 1998). Simply put, much AR research and development focuses on making virtual objects as real as possible.

There is nothing wrong with this aim. Many contexts, such as training environments or exposure therapy sessions require a realistic setting to be effective. However, the imitation of reality is not the only path towards creating meaningful AR experiences. A unique power of virtual objects is that they do not have to look, feel or behave like real objects. In our opinion, AR research can harness this power and create imaginative forms of realities that offer new and unique experiences. Accordingly, our research follows another direction: Instead of imitating reality, we want to create new experiences that have no equivalent in a purely physical world. We are interested in how augmented reality scenarios can differ from strictly physical, 'unaugmented' environments.

In the previous chapter, we have explored this idea with respect to interaction between the virtual and the real. However, the virtual objects used in this preceding exploration still have mimicked physical objects—both with respect to appearance and behavior. For instance, we have used virtual spheres that look like physical spheres and that seemingly are affected by gravity (as well as other imaginative forces). In this chapter, we want to focus in on the idea that virtual objects do not have to mimic physical objects and explore alternative manifestations of the virtual. To do so, we create a virtual object that does not look, feel or behave like any real object. Although we are not in-

terested in imitating reality, our project shares important goals with existing AR research: it tries to convey the presence of a virtual object in an otherwise real environment.

Like in the previous chapter, we build on the idea that virtual objects can differ from real objects. However, in this chapter, we take this concept one step further. If virtual objects have different qualities than real objects, they might also be perceived differently from how real objects are perceived. For instance, while most real objects can be seen, we might not sense a virtual object simply by looking at it.

Building on the ideas that virtual objects (1) can differ from real objects and (2) can be perceived differently from how we perceive real objects, we have developed a new kind of virtual object—the so-called sonically tangible cube. Unlike real objects, this cube is invisible and does not provide tactile feedback. However, touching the virtual cube triggers binaural sounds that appear to originate from the exact spot where it is touched. Our initial experiments show that through this sonic feedback, virtual objects can gain an almost-tactile quality and appear as if they were actually present in real space. It is this idea of making virtual objects both tangible and present through spatial sonic feedback that makes "sonically tangible objects" unique.

Several questions have fueled the development of the invisible cube and our research into sonically tangible objects. For instance, we were wondering if it is possible to leave out the tactile component in tangible perception. If there is no tactile stimulation, would the virtual object still be perceived as part of real space—and if so, would it be experienced as an object with a tactile or physical component? We were intrigued by how one experiences an object that provides no tactile sensations. Most importantly, however, we were eager to learn more about how virtual objects in an AR environment can differ from real objects.

While we provide preliminary answers to these questions, the main contribution of this chapter lies in the proposed concept of sonically tangible objects. This concept, to the best of our knowledge, is new. So far, inferences regarding the perceptual qualities of the invisible cube are based on informal testing and on our subjective experience with the cube.

The central idea—that a virtual object can be tangible but not tactile—calls for a distinction between the terms tangible and tactile. In this chapter, things are called tangible, if they can be perceived by touching (being in contact with) them.[1] Only objects that also stimulate the tactile receptors (as found in the skin and tissue) are referred to as tactile. This understanding creates room for objects that are tangible but not tactile.

[1] We use the word "touch" to refer to gestures where the body is (brought) in contact with an object.

Unlike in previous chapters, we will discuss issues related to implementation and technology in more detail. This is because this knowledge is necessary to reproduce and experience sonically tangible ob-

jects. Another reason for this is that different concepts, such as making sounds appear from specific positions in the real environment, are strongly interlinked with the chosen implementation, such as the use of binaural recordings (this recording technique will be explained in more detail later on).

One primary goal of this exploration is to create a new type of object and with it, a new type of experience. In line with this, we focus on the practical exploration of sonically tangible objects rather than on theory. However, by pursuing this idea, we also hope to learn more about the possible manifestations of AR. Furthermore, we hope to better understand the concept and context of sonically tangible objects. This is why we will have a look at sonically tangible objects from different perspectives and place it in the context of related research.

The chapter is structured as follows: In the following section 6.1, we focus on the practical side of sonically tangible objects. We share choices made and insights gained during the development of the invisible cube, describe the used setup and implementation and discuss our experience with it. In section 6.2, we place the cube in the context of pertinent research and compare our project with related work. Because our project is multi-disciplinary, we consider research from various fields, such as augmented reality, tangible interaction and perception. The chapter ends with a reflection on the project and possible directions for future research (section 6.3).

6.1 *The Sonically Tangible Cube*

The sonically tangible cube is a virtual object. It is unlike any real object in the sense that it is non-tactile, invisible and lacks many physical properties, such as weight and temperature. It does, however, have sonic and spatial properties such as a shape, loudness and sound texture. Although the cube has no tactile component, its presence can be perceived through touch. When fingers enter the cube, sound appears to originate from the spot where the virtual object is touched. The resulting sonic feedback not only corresponds to the fingers' positions but also fits the movement of the fingers. Fast finger movements result in more agitated soundscapes while slower movements cause less dense, more distinct feedback. As the cube is virtual, fingers can move through it and explore its inner texture.

6.1.1 *Implementation and Setup*

The virtual cube is 20 cm x 20 cm x 20 cm of size and it floats 10 cm above the otherwise empty work desk of the author. The technical setup consists of a Leap Motion Controller (see `www.leapmotion.com`), which detects the position of the participant's fingertips in real space. The Leap Motion is placed on the desk and senses hand movement

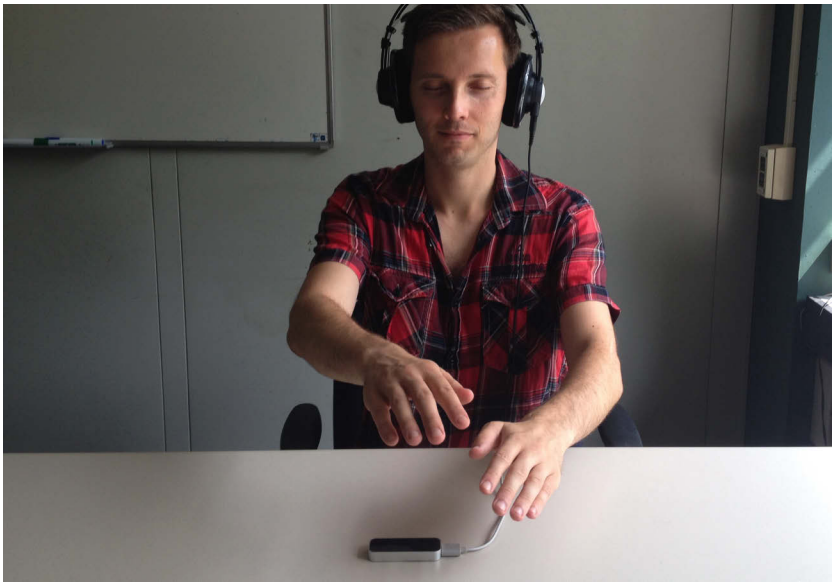

Figure 6.1: A colleague explores the virtual, invisible and non-tactile cube. A Leap Motion Controller is used to track the position of his fingertips.

above the device (see figure 6.1). A custom *Max* (2014) patch, which runs on an Apple Mac mini, interprets the data provided by the Leap Motion. Interfacing with the Leap Motion device is realized with a Max external object 'aka.leapmotion' by Akamatsu (2014). In our current setup, the frame rate of the Leap Motion device is around 57 fps when the office is naturally lighted and slightly above 200 fps when the amount of interfering infrared light is reduced by darkening the room. The Max patch evaluates whether and where the participant is touching the cube on the basis of the fingers' coordinates. If the fingers are located within the 20 cm x 20 cm x 20 cm area that has been defined as the cube, their movement triggers pre-recorded binaural sounds. This interpretation of the finger position works for every finger independently and allows the participant to explore the cube with up to ten fingers at a time.

A constraint of the current setup is that the sound only matches the fingers' position if the participant is sitting at the right spot and directly facing the cube.[2] Also, due to the frame rate of the Leap Motion device, very fast hand-movement can cause a mismatch between the hand-position and the spatial information of the triggered sound. Moreover, finger movement is sensed best, if the hands are held horizontally. Fingers that are not in a direct line of sight from the Leap Motion device cannot be sensed.

[2] This is the case because the pre-recorded sounds are positioned in relation to the listener. If a sound originated from right in front of a listener while being recorded, it also appears right in front of the listener when it is played back. If one moves while listening to a recording, the recorded sounds move along. Hence, for the cube to remain at the intended position, the participant can not move. The underlying binaural recording technique will be explained in more detail in section 6.1.2.

6.1.2 *Development and Choices*

The sonically tangible cube was developed iteratively during the course of several months. In the course of the project, the author acted as a researcher, developer and participant. Additionally, colleagues

were occasionally asked to provide feedback and describe their experience with the cube.

From the beginning, we have explored the idea of making virtual objects tangible and present through sonic feedback. The topic of (in)visibility was left aside for future research and hence, many evaluations have been conducted with closed eyes. This choice was made because in a very early stage, it became clear that it was easier to concentrate on the sonic aspects with closed eyes. Furthermore, not seeing an object in the space seemed to potentially interfere with the audible information. We thus decided to first only focus on the tangible and audible experience. Aside from this, four determining observations and decisions were made concurrently in the early stages of the development process. These important choices concerned (1) the shape of the object, (2) the recording technique, (3) the recorded material and (4) the sound design (including the mapping between movement and sound)—together, these choices determine the properties of the invisible object.

The Shape: A Cube

One of the most fundamental early decisions regards the shape of the object. We have started out with several simple geometric shapes. For instance, we have used a laser pointer and triggered a sound every time the laser was interrupted to see if this would evoke the experience of touching a virtual string or line. Furthermore, we have tested triggering sounds whenever the fingers crossed a predefined plane in the real environment. We expected that this might create the feeling of 'crossing a virtual border'. However, our initial experimentation indicated that it is very difficult to experience a plane or a line. Running one's hands freely through a three-dimensional object and exploring both its borders and inner texture offered the most intriguing, tactile-like experience and promised to convey an object's presence in space best. Given that we deemed it best to begin the exploration with a *simple* 3D object with a very clear geometry, we decided to focus on a cube-shaped virtual object.

The Recording Technique: Binaural Audio

Another crucial decision concerns the sonic aspect of the project. In the beginning, simple synthesized clicks were played back in mono (feeding the identical signal to both the left and the right channel) through closed Beyerdynamics DT 770 Pro headphones whenever a virtual object was touched. This was done in order to learn about the effects of linking movement in a certain area in the environment to a very basic sonic response. However, our initial trials showed that the resulting experience was closer to being informed or 'being told' that one's hand had entered a predefined space and there was no direct

sensory experience of an object in space. While there was a clear link between the movement and the sound, it only felt as of one triggered the sound through movement and not as if the sound originated from the virtual object.

This experience did not come as a complete surprise. After all, interacting with real objects and materials—crumbling paper, scratching on a surface, typing on a keyboard or moving the mouse—causes sounds that originate from the objects themselves and from the position where the objects are touched. Based on this, we assumed that in order to convey the presence of an object at a certain position in space, it could help to make it seem as if the object's sounds originate from its position.

Here the idea of using binaural audio to achieve this effect came into play. We had previously encountered this technique when reading about and listening to the work of Cardiff (n.d.), who uses binaural recordings for her audio walks (see also section 3.1.3). Binaural audio is based on the notion that hearing makes use of two signals: the sound pressure at each eardrum (Møller, 1992). If these two signals are recorded in the ears of a listener (or a dummy head), the complete auditive experience—including the three-dimensional spatial information of the sounds—can be reproduced by playing the signals back at the ears. By making binaural sound recordings in which the sounds originate from the spatial position of the virtual cube, and playing these recordings back via headphones when the fingers enter this area, the resulting sounds will sound as if they originate from the location of the fingers/object.

Based on our theoretical considerations, the use of binaural audio seemed promising. However, as the author had little practical experience with this recording technique, many questions remained open: Would binaural recordings indeed allow us to make sounds appear from any position within the room? Would these sounds be distinguishable from 'real' sounds. Would the experience really be that different from listening to mono-sounds or stereo recordings? In order to get a better idea about the actual potential and limitations of binaural recordings, we decided to first investigate its qualities with some simple experiments. For example, we recorded the sound of someone knocking on the closed office door and the sound of the ringing phone while working in the office. We intentionally chose sounds that originate from objects that are (already) physically present in the space. Likewise, we made sure to use sounds that are not accompanied by visible changes (it is not possible to see whether the phone is ringing or whether someone is standing behind the door).

From these initial experiments, it became clear that binaural audio indeed can convey the desired experience of 'something happening' in the real space and 'something being present in the space'. Listening back to the recordings while working, the sounds seemed to

originate from those exact spots where they originally had happened and sounded equally real as the original sound. Consequently, the author often was in a state of doubt: had someone knocked on the door, was someone passing in front of the office door, did someone open the door to the toilet across the hallway, was someone actually calling? The virtual ringing of the phone was practically indistinguishable from a real call. Similarly, it was almost impossible to tell whether someone actually was waiting behind the door, or whether a knocking sound was played back from the recordings. In fact, a simple virtual knock proved powerful enough to communicate the presence of something or someone in real space.

Although the sounds were extremely realistic, there was also a simple and safe way for the author to distinguish between real and virtual events: Whereas real sounds remained unaffected by head movement, virtual sounds would move along! This happened because binaural recordings are positioned in relation to the listener. If a sound originated from right behind a listener while being recorded, it also appears right behind a listener when it is played back. If one moves while listening to a recording, the recorded sounds move along.

From our initial exploration, we concluded that binaural recordings could create a strong sense of something happening in the real environment. However, we also realized that in order to make a sound originate from a certain spot in space, we would have to determine the position and orientation of the listener. This is why we decided that for this initial realization of sonically tangible objects, participants would sit in the author's office chair looking straight ahead without moving their head.

The choice for binaural audio went hand in hand with a switch to the open AKG K702 headphone. Due to the open nature of the headphones, the recorded sounds mix in with the sounds naturally present in the environment. This additionally supports the experience that the virtual object inhabits our real physical space rather than a virtual or separate space.

The Recordings: Crumbles of Aluminum-Foil in a Plastic Bag

What should the virtual sonic object sound like? The choice of using binaural recordings introduced the question of what to record. We were searching for sounds that (1) are abstract (do not invoke the idea of a specific real object), (2) have a tactile quality (indicate touching) and (3) support the idea of a non-solid object/material that allows the fingers to move through it. Several different sound sources were tested: for example, foils, paper, plastics, packaging materials from everyday objects, rattles and empty bottles. All sounds were produced by interacting with the materials with the hands and fingers. This choice was based on the assumption that sounds that actually are cre-

ated by hand/finger movement are more likely to fit the exploratory hand gestures of the participant and are more likely to create a tactile-like experience (similarly to how the sound of squeaking nails on a chalkboard can be an almost-tactile, physical experience, even if someone else is scratching the board). For the current implementation of the sonically tangible cube, we have settled on the sound of aluminum foil, produced by squashing a tiny plastic bag filled with small crumbles of the foil (see figure 6.2).

Figure 6.2: We recorded the sounds of squashing a tiny plastic bag filled with small crumbles of the aluminum foil (the image shows a recreation).

It was clear that this 'crumbling sound' has to be recorded at the intended position of the virtual cube. However, it soon became apparent that recording only one sample at the location of the cube was not convincing because the sound did not seem to move when the fingers were moving. We thus determined that when, e.g., moving one finger from the left side of the object to the right side of the object, the sound should move along with the finger. More generally, we decided sounds should seemingly originate from the exact position where the cube is touched (at the fingertips of the participant). To achieve this, we divided the cube into 64 sub-cubes of 5 cm x 5 cm x 5 cm (see figure 6.3) and recorded five-second samples of aluminum foil sounds at all 64 positions within the cube.

The 64 recordings were made with a ZOOM H4 audio interface and two DPA 4060 microphones. The microphones were placed slightly above the ear-entrance of the author and the sound was recorded with a basic Max patch. For the recordings, the author successively produced the desired sound by squashing the little plastic bag and rubbing the aluminum crumbles against each other at each of the 64 sub-

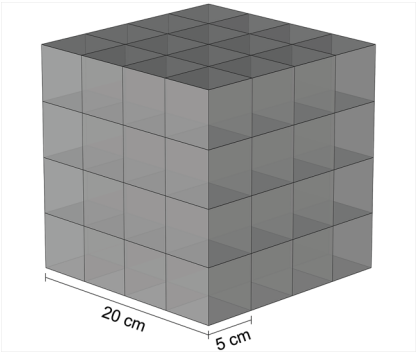

Figure 6.3: The sonically tangible cube was divided into 64 sub-cubes. A binaural recording was made at all 64 positions. Image of the cube contributed by Wim van Eck.

areas. Aside from this, the author was sitting motionlessly in front of the desk, facing the cube just like participants do during the experience (see figure 6.1).

Sound Design and Mapping: Movement Causes Sound

When a participant interacts with the cube, the positions of his/her fingers determine which of the 64 recorded audio samples are played back. If a finger is placed in a sub-cube, the corresponding recording is activated. However, first tests showed that simply playing back the recordings resulted in a sound that only matched the fingers' positions, but not the different variations in hand and finger movement (slow, fast, no movement, etc.). Because this felt not convincing yet, we experimented with more complex settings that map the movement of the fingers to parameters in the sound design.

Our current implementation knows two sound design settings. The first setting makes use of granular synthesis. Granular synthesis makes use of very short snippets of audio; so-called grains. These grains typically are between 1 and 50 ms long and can be layered (Roads, 1988). Playback parameters such as speed or volume can be varied for each grain individually. In our sound design, all grains are taken from the binaural recordings and between 10 ms and 20 ms long. A random offset is used to vary the position in the binaural recording from where each grain is taken. When a grain is played back, it is randomly varied slightly in pitch/playback speed.[3] As a result, every grain sounds differently. (This appeared to be crucial for the believability of the experience.) The key concept in this first setting is that the change of a finger's position triggers the playback of an audio grain taken from the corresponding 5-second recording. This means that fast finger movements trigger a lot of grains successively, whereas no movement does not trigger any grains.

The second setting follows a similar underlying idea. However, instead of working with short grains, the entire 5-second recording is layered, varied and looped if there is movement in a sub-cube. In this setting, a faster movement activates more layers. For instance, moving slowly, one would only listen to one version of the recording. Moving

[3] By varying the speed, the pitch automatically changes as well. Furthermore, the changes in playback speed also in influence the spatial characteristics of the sounds. However, as those variations were minimal this effect was perceptually negligible.

a bit faster, a second (varied) version of the recording would join in, starting from a different position and playing in a slightly different speed.[4]

[4] Here, too, the changes in playback speed influence the spatial characteristics of the sounds. However, these variations were again small and the effect negligible.

The settings have many similarities: Both result in a louder, more complex and dense soundscape if the finger moves fast and in a softer, less dense but more distinct soundscape if the movement is slow. As this happens for each finger individually, the number of fingers used by the participant has a similar effect: The more fingers are involved in the exploration, the denser the resulting sound. For either setting, some movement is necessary to 'excite' the virtual cube and to elicit its sounds. No movement results in silence, even if the hand is placed in the cube. However, as it is impossible to keep one's fingers completely still, occasional slight trembling of the digits will cause corresponding sound output.

There are also differences between the settings. In particular, the two settings differ with respect to the textural nature of the sound. Whereas the granular synthesis results in a more gritty and rough soundscape, the layered loops produce a thinner, airier sound texture. Furthermore, the granular synthesis sounds a bit more abstract and less like a real-world recording. (This fits well with our intention not to mimic the real world.)

Ultimately, the choice of the recording technique, the chosen material we recorded and the chosen sound design and mapping all determine aspects of the sonic qualities of the cube—together, they determine how it sounds to touch the object.

6.1.3 *Experiencing the Cube*

Ideally, the previous section has provided some insight into how the invisible cube sounds. However, how does the cube *feel*—does touching the cube really feel different from simply moving one's hands through thin air? Do we experience the cube as present in space, do we perceive it as tangible? It is important to systematically investigate this by performing experiments with a group of unbiased participants in the future. In the following, we compare experiences with the cube to experiences we know from the everyday world. These comparisons are objective in the sense that some similarities between the cube and existing real-world phenomena are simple facts. However, in addition, the author also provides her own subjective account of how it feels to interact with the cube.

On some level, experiencing the cube has indisputable similarities with moving one's hand through a beam of light. When touching a beam of light, one can clearly see the beam's presence in space but one cannot feel it. Similarly, in the case of the cube, one can hear the cube's presence in space, but one cannot feel it. In both cases, there is no tactile stimulation on our fingertips. At the same time, the cube can

also be considered an opposite to a beam of light. Whereas we can see a light beam but cannot experience it with touch gestures, we cannot see the cube but we can experience it with touch gestures.

Given those similarities and differences, does touching the cube actually feel like touching a beam of light? Judging from the experience of the author, similarities between the two indeed exist. In particular, both seem to provide a quite fascinating 'dissonance between the senses'. In other words, both create an intriguing experience where one sense tells us 'something is there', whereas another sense tells us 'nothing is there'.

In addition to touching light, experiencing the cube also has some undeniable similarities to feeling out a physical object blindly with one's hands. After all, it is only through the physical act of touching that we can perceive the cube in the first place. There is no notion of the object unless one is in contact with it. Also, like in typical haptic perception, the experience of the object takes time and happens through exploratory gestures with one's fingers. Furthermore, touching a real object can cause sounds at the position where the object is touched. The same happens when touching the sonically tangible cube.

So does touching the cube feel like touching a physical object blindly? When exploring the cube, the author indeed often was reminded of blindly interacting with a physical object. However, this might have been caused by the simple fact that the cube was primarily explored with closed eyes. Although exploring the cube with closed eyes was a conscious decision, the author noticed that closing the eyes also came naturally when exploring the cube. This likely happened because the presence of the virtual cube was experienced as stronger and more convincing when the eyes were closed. Possibly, the author tried to avoid the conflict between what she saw and what she heard. (However, as mentioned above, the conflict between senses also was experienced as fascinating.)

According to the author's experience, exploring the cube also feels similar to exploring a physical object blindly because a model of the object's shape emerges in one's mind over time. Of course, when interacting with the cube, the author was always aware that she was dealing with a cube of a certain size. However, she would still repeatedly reconstruct its shape and form an internal representation of the object over time. When exploring the cube, a mental image of a cubic cloud floating in the space before the researcher regularly emerged. In some way, the cube seemed to be 'just air'. However, at the same time, it also seemed clear that 'something was there'.

Although there are obvious similarities between the way we blindly interact with physical objects and the way we perceive the cube, there are also quite some key differences. One can, for example, not hold, move and turn the cube. Instead, it is possible to move right through

the cube and explore its inner texture and structure. Unlike with physical objects, it is impossible to simply follow the contour of a sonically tangible object and to explore its shape that way (cf. Lederman and Klatzky, 1987). Rather, the contour can be perceived by repeatedly crossing (zigzagging around) the border of the object and moving in between the sonic space of the cube and the silent space surrounding it.

In the author's personal view, these differences between the cube and real objects are the most fascinating aspects of the project. According to her experience, perceiving and interacting with the cube indeed feels different than interacting with any real-world object. In particular, zigzagging around the borders of the cube feels fascinating. In addition, drawing three-dimensional shapes with one single finger inside of the cube is intriguing (see figure 6.4). This might be the case because here the relationship between movement and sound is experienced most clearly. However, the fact that the virtual cube is different from any real object also has a downside. The cube (at times) can evoke the frustrating feeling that it is impossible to really 'get hold of it' or to grasp it. In line with this, the author often wondered how she could move the cube around.

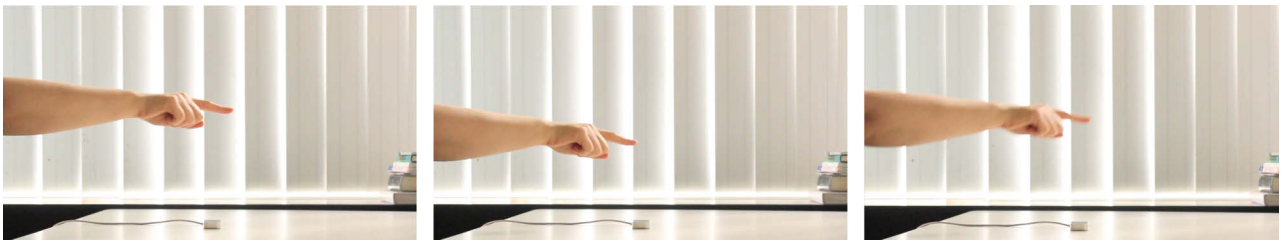

Figure 6.4: The author interacts with the cube.

Interacting with the cube is not only similar to touching light, or interacting with an object blindly—it also has undeniable similarities with playing gesture controlled open-air instruments such as the Theremin. The Theremin is played by moving one's hands in the space between two antennas. The position of the hands determines the sound. When interacting with the cube, movement of the hands in space similarly results in sonic output that corresponds to the position of the hands. The author has never played the Theremin or other gesture-based instruments. This makes it impossible to compare the experiences subjectively. However, in the author's experience, playing an instrument and exploring the cube are comparable; likely due to the cube's "sonic expressiveness". In fact, the cube seems to allow for some (but only very limited and basic) forms of musical expression. These become particularly apparent when holding the hand completely still inside the cube and then shaking the hands to a variable extent (see figure 6.5). By doing so, one can create a variety of

rustling noises and exert a high amount of control over the resulting sound textures and volume. However, when compared to many existing instruments, there is little possibility for sonic variation and no possibility to control the pitch. (In this regard, the cube seems more comparable to the family of shaker instruments.) However, it is easy to imagine a similar Leap Motion setup that focuses on music generation rather than on communicating the presence of a virtual object in space.[5].

[5] Not surprisingly, the potential of the Leap Motion device in the context of new digital musical instruments has explored by several researchers in the past (e.g., Han and Gold, 2014; E. S. Silva et al., 2013)

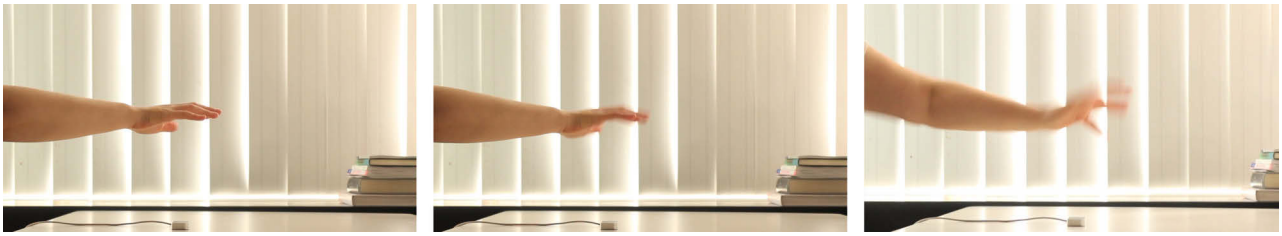

Figure 6.5: The author shakes the hand to a variable extent to produce different types of rustling sounds.

In the course of this study, the cube was mostly explored with closed eyes. This raises the question in how far the cube is experienced as part of (and in relation to) the real environment. When the author interacted with the sonically tangible object, the cube was certainly experienced as present in 'her environment'. In other words, the cube seemed to exist in the same space and environment as the author. Because the author felt present in her office (even with closed eyes), the cube also was experienced as existing in this office. However, the relationship between the cube and other elements in the office space was often experienced as rather weak. For instance, the fact that the cube was floating over the office table felt more like conscious knowledge rather than like a perceptual experience. At times, the link between the cube and the real environment seemed stronger when exploring the cube with open eyes. However, opening the eyes, in turn, seemed to weaken the experienced presence of the cube in space.

When it comes to experiencing the cube as a part of the real environment, an interesting observation was made while fine-tuning the sound design. At this stage, the author often interacted with the cube but also had a computer monitor placed on the desk to make changes to the mapping. The used patch (program) included a visual representation of the cube, as well as visual representation of the author's fingers (see figure 6.6). When this visualization was shown on the screen, the cube was no longer experienced as part of the real environment. Rather, it immediately felt as if one were reaching into the space depicted on the monitor and touching the virtual object shown on the screen. This happened although the sound was recorded at (and therefore indicated) a position in front of the screen.

A question that needs to be addressed is whether the tactile sensation is missing when exploring the cube. In the author's opinion,

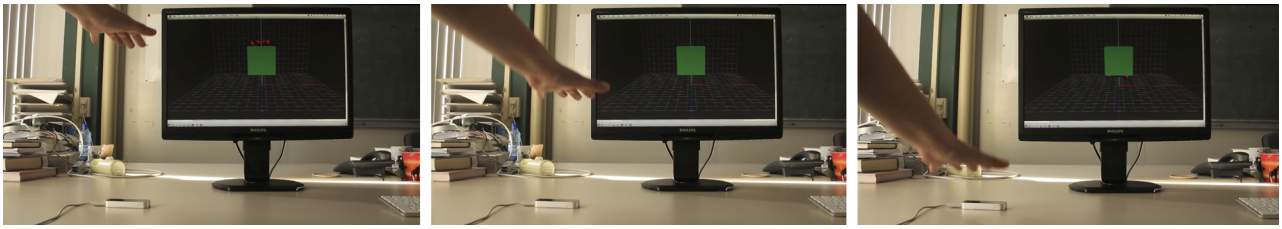

Figure 6.6: During the development of the sound-design, a monitor was placed on the desk. As soon as a visual representation of the cube and the fingertips was presented, the cube no longer felt like a part of the real environment. Instead, it felt as if one were reaching into the space depicted on the monitor and as if one were touching the displayed cube in this space.

this is not the case. Just like we do not miss a tactile sensation when listening to music or looking at an object, the author was not missing a tactile sensation when experiencing the cube. However, as mentioned above, hearing oneself touch the cube without feeling a tactile sensation on the fingertips felt somewhat like a contradiction between senses.

During this project, the author often wondered what bystanders would experience if they witnessed the author's (admittedly awkward looking) interaction with the cube. With respect to this, the author was presented with some surprising initial insights after presenting a video of her interaction with the cube at a colloquium. Naturally, the video (in combination with stereo sound) was not able to convey the experience of interacting with sonically tangible objects directly. Yet, an interesting observation was made: After watching the author interact with the cube and listening to the sonic results on speakers (rather than headphones), some audience members were under the impression that something was present in the space above the Leap Motion and that the author was interacting with this 'thing'. Other audience members, however, indicated that in their opinion, the space was clearly empty. Of course, these reactions are based on a video and were shared in a very informal context. As such, they do not yet answer how the project is experienced by bystanders that are present in the same space as the cube. However, some feedback from the audience suggests that the cube can not just be experienced directly, but also indirectly—by witnessing someone else's interaction with it.

To conclude: It remains difficult to put the experience of the cube in words. However, to the author one thing seems clear: Touching the cube is different from simply moving one's hands through thin air. When we move our hands through air, we feel nothing but empty space. The cube, however, is experienced as something that is present and as something that can be touched and that invites playful exploration. It seems to inhabit the space, albeit in a non-physical way. Although the experience is not tactile in the traditional sense, it definitely has tactile-like aspects. According to the author's experience, running once fingers through the object feels like 'something is here that can be touched'.

6.2 *The Cube in Context*

Our project is multi-disciplinary; it draws from and contributes to various fields of research, such as augmented reality, tangible interaction and perception. In this section, we take a second look at the cube from various different perspectives and discuss the project in the light of related research.

6.2.1 *Augmented Reality*

We have arrived at the concept of sonically tangible objects as part of our trajectory researching augmented reality. The invisible, sonically tangible cube can be seen as an AR project because virtual content is presented in and relates to the real environment. More specifically, the project is concerned with the presence of virtual content in real space, and like many AR projects, aims at making it seem as if additional virtual objects were part of the otherwise real environment. (We have termed this form of AR *presence-based AR* in section 3.4.) In this general sense, the sonically tangible cube relates to all projects, that aim to convey the presence of virtual objects in real space.

In existing research, the presence of virtual objects in the environment is typically conveyed by visually displaying the object in the space. Our cube is similar to traditional visual virtual objects because both do not have a tactile component. A difference between the sonically tangible cube and traditional visual virtual content is that visual content can be experienced without touching it. When experiencing visual virtual content, the absence of tactile stimuli only plays a role if and when one actually touches the object. When experiencing the cube, the absence of tactile stimuli is always apparent when it is perceived.

Our project explores whether the presence of virtual objects can be experienced through a combination of touch gestures and spatial sound. Hence, our project specifically relates to those AR projects that use sound and/or tangible interaction to convey the presence of (invisible) virtual objects in real space. A project where the presence of something virtual is perceived tangibly is Sekiguchi et al.'s (2005) so-called Ubiquitous Haptic Device. The little box, when shaken, conveys a feeling of a virtual object being inside the device. Similarly, a wearable haptic device by Minamizawa, Fukamachi, Kajimoto, Kawakami, and Tachi (2007), called the Gravity Grabber, allows participants to perceive the ruffle of the water in a glass, although he/she actually is holding an empty glass. More examples of objects that convey the presence of virtual objects in real space by means of tangible cues can be found in subsection 4.4.3.

Projects that let a participant experience the presence of "something that is not really there" by means of sound have been discussed in detail in subsection 4.4.2. An example is the SoundPacman game by

Chatzidimitris et al. (2016). This game makes use of spatialized sound in order to give game elements a position in the real environment and to communicate the location of virtual ghosts to the player. Other examples are Cilia Erens' (see Erens, n.d.) and Janet Cardiff's (see Erens, n.d.) sound walks. Both artists use binaural recordings of everyday sounds that blend in with the sounds present in the real environment when the participant navigates the space and listens to the composition on headphones.

Given that our implementation was inspired by Cardiff's approach, it is maybe not surprising that a discussion of Janet Cardiff's work by Féral (2012) also helps our understanding of sonically tangible objects. Féral defines "presence effects" as the feeling that an object (or body) is really there, even when one knows that it is not. This not only plays a role in Cardiff's works but also relates to the experience of the sonically tangible cube. While the ears make it feel as if the cube is present, the lack of tactile (and visual) stimuli informs us that nothing is there.

Our project shares important goals with much existing AR research. Most importantly, it shares the common goal of making it seem like a virtual object existed in real space. However, at the same time, our project also differs from most AR research in one important regard: We do not try to make it seem as if the virtual object was a real, physical object. In fact, the cube aims to be different from any real object. In contrast, many existing AR projects aim at creating a scenario where the user cannot distinguish between what is real and what is virtual. For instance, Azuma (1997) writes "[a]fter the basic problems with AR are solved, the ultimate goal will be to generate virtual objects that are so realistic that they are virtually indistinguishable from the real environment" (p. 380). In line with this, Vallino (1998) suggests that ideally, "[v]irtual and real objects are visually indistinguishable" (p. 20) and R. Silva et al. (2003) point out that "[a]lthough many AR applications only need simple graphics such as wireframe outlines and text labels, the ultimate goal is to render the virtual objects to be indistinguishable from the real ones" (p. 10).

We clearly do not share this 'ultimate' goal. Instead, our goal is to create an object that is so different from everything we know from the real world, that it evokes a new kind of experience. We thus believe, AR research has to distinguish between two goals: (1) making it seem as if a virtual object was a real, physical, material object and (2) making it seem as if a virtual object were really *present* in the space. Whereas those two goals typically go hand in hand in existing AR research, we only share the second goal. We do not mind if the object is experienced as 'virtual'. Ideally, when touching the cube, one would think "something is here, but it is not like any real object I know".

6.2.2 *Tangible and Embodied Interaction*

The cube deals with (in)tangibility, requires active bodily engagement and it explores the possibilities of a tangible experience without tactile stimuli. As such, our research relates to the field of tangible and embodied interaction.

The term 'embodied interaction' has been coined by Dourish (2004), who defines embodied interaction as "the creation, manipulation and sharing of meaning through engaged interaction with artefacts" (p. 126). Research into tangible and embodied interaction is an entire research field in itself, which in turn draws from and contributes to many other fields. However, one key aspect that links the sonically tangible cube to this field is the explicit link between interaction and experience. As Dourish (2004) points out, tangible interaction builds on the idea that we experience the world through directly interacting with it, and that acting in the world happens through exploring the opportunities it offers for action (see, e.g., p. 18).[6] This idea is also at the basis of the sonically tangible cube. One can only experience the cube through interaction with it, and perceiving the cube happens through action in the space and by exploring the possibilities for interaction. There is no way to perceive or experience the cube passively or without a body.

[6] In its original context, this point is part of a larger argument, namely that this is true for both tangible interaction and social computing. Accordingly, one of the main points of Dourish's book is that tangible and social computing share embodiment as a core element.

On the one hand, one could argue that this focus on perception through interactions sets the invisible cube apart from those common virtual objects that we can only perceive by looking at them. On the other hand, one could argue that all forms of perception—also seeing an object—entails interaction with the world. In fact, more recent theories of perception suggest that all perception is active. For instance, Noë (2004) states that "perceiving is a way of acting. Perception is not something that happens to us, or in us. It is something we do" (p. 1). Our point with respect to this is that perception is not always *experienced* this way. Seeing or hearing something often feels like it 'happens to us', whereas exploring the shape of an object blindly typically involves conscious active exploration. With the sonically tangible cube, tangible interaction and tangible perception are explicitly designed to be experienced as one and the same thing.

In our opinion, the sonically tangible cube can not only be placed in the context of AR, but also in the context of the growing body of embodied interaction research that is concerned with "bodily action, human experiences, and physicality, in the context of interaction with and through a world comprised of computationally mediated artifacts and environments" (Antle et al., 2011, p. 9). A specific contribution of our project in this context is a new form of tangibility that is evoked through a combination of spatialized sound and touch gestures and which does not involve any tactile stimulation.

We believe that an important link to explore further is the relation-

ship between presence and tangibility. In our experience, tangibility and presence are closely linked. More specifically, presence is a necessary condition for tangibility. We can only touch an object if it is present. If we touch an object, we and the object are both present in the same space—at least in a mediated way.

6.2.3 *Human-Computer Interaction*

One possible area of application for sonically tangible objects is the field of Human-Computer Interaction (which overlaps with the above-mentioned field of tangible and embodied interaction), and in particular what Chan et al. (2010) call *intangible displays*. Intangible displays are visual virtual interfaces that appear in mid-air, in front of a user's eyes. Aside from simply displaying information they also allow for interaction: Users can touch virtual objects, such as buttons, with their physical hands. However, intangible displays do not provide tactile feedback when they are touched.[7] Chan, Kao, Chen, Lee, Hsu, and Hung (2010) address this lack of tactile feedback by providing visual and audio feedback. In their experiments, they played short sounds whenever participants touched the surface of the intangible display. While similar, their project differs from ours in the sense that sound is used to *inform* the user about the fact that they have successfully touched the object. The sound serves as feedback and not as an integral part of the object.

[7] As mentioned, when referring to the cube, we use the term tangible rather than intangible. We dot so because although there is no tactile feedback, the cube is perceived by touching (being in contact with) it.

Although originally not intended this way, the concept of sonically tangible objects could be used to improve the interaction with intangible displays. It could increase the spatial presence of the display, provide better feedback about the user's hand position and movement through the display and is likely to make the "the awkward feeling of 'touching' a mid-air display" (Chan et al., 2010, p. 2626) less awkward and more tactile-like.

Another related HCI project is the so-called BoomRoom (Müller et al., 2014). In this room, sounds seem to originate from certain spots in real space (this is realized with a circular array of 56 loudspeakers and Wave Field Synthesis). These sounds can be 'touched' in order to grab, move and modify them. Although related, their project differs in the sense that it focuses on the localization and direct manipulation of sound rather than on the presence and tangibility of virtual objects.

6.2.4 *Perception*

Our project relates to perception research, and in particular research into haptics, tactile illusions and cross-modal interactions as well as sensory substitution.

Haptics

The sonically tangible cube is perceived by explorative hand gestures. This links it to the field of haptics. Haptic perception typically involves active exploration (Lederman and Klatzky, 2009). Haptics is commonly understood as a perceptual system that derives and combines information from two main channels: kinesthetic perception and cutaneous sensation (Lederman and Klatzky, 2009). Cutaneous sensation is derived from the receptors that are found across the body surface and that allows us to feel, for example, pressure or temperature. The kinesthetic channel refers to perception of limb position and movement in space, which is derived from the receptors embedded in muscles, tendons and joints.

Kinesthetic perception also plays a key role in the perception of the virtual cube—it provides the participants with the information about where and how fast their fingers are moving in space. This awareness is crucial in order to link what one hears to one's movement in space. What makes the perception of the sonically tangible cube different from common haptics is the lack of cutaneous feedback (including tactile sensations). Rather than 'feeling something at the position where they touch an object' the participants 'hear something at the position where they touch the object'.

An aspect that would be interesting to explore systematically in the future is the use of exploratory gestures that one applies to explore sonically tangible objects. In "Hand movements: A window into haptic object recognition", Lederman and Klatzky (1987) identify six typical movement patterns that are used to explore the properties of real objects and link them to the type of knowledge they reveal about an object. For instance, Lederman and Klatzky (1987) show that we learn about the exact shape of an object by following its contour. As our project shows, other types of gestures apply when it comes to sonically tangible objects. For instance, we experience the shape by zigzagging in and out of the object. Furthermore, it is possible to explore the inner texture of an object, but there is no mass, weight or temperature to explore. In the future, it would be interesting to learn more about hand gestures used to explore the properties of *virtual* objects, and the role sound can play in conveying their properties. Based on our experience with the sonically tangible cube, we believe the sound caused by an exploratory gesture can play an important role, especially when it comes to textural properties of virtual objects.

Tactile Illusions and Cross-Modal Interactions

The sonically tangible cube aims to create a tactile-like experience. There are several studies that indicate that sound can influence actual tactile experiences. The "Parchment-skin illusion" (Jousmäki and Hari, 1998) shows that modifying the sounds that accompany hand-

rubbing can influence the tactile sensation of the skin. It was found that accentuating the high frequencies can lead to the experience of a higher level of skin roughness. Hötting and Röder (2004) have discovered another auditory-tactile illusion. In their experiment, one tactile stimulus was accompanied by several tones. As a result, participants reported that they perceived more than one tactile stimulus. What sets these illusions apart from our cube is that in both cases, the participants were presented with a tactile stimulus.

A study that suggests that a tactile experience can be evoked without presenting any tactile stimuli has been reported in the context of Virtual Reality. Biocca, Kim, and Choi (2001) suggest that visual cues can cause haptic illusions. In their experiment, participants reported that they felt physical resistance when manipulating virtual objects in a virtual environment although the interface contained no haptic displays and the environment provided no direct stimulation to the haptic channel. However, it has to be noted that the participants were wearing gloves that allowed them to manipulate objects by pinching their fingers together. Hence, also here, tactile stimuli (from pinching the fingers together and from the gloves) were present. Thus, it remains unclear whether the visual cues evoked the tactile experience or simply altered present tactile sensations.

Sensory Substitution

One could argue that sonically tangible objects allow us to hear sounds instead of feeling a tactile sensation. In this sense, the cube relates to projects that use sound to substitute touch. One such sensory substitution system is F-Glove (Hafidh, Osman, Alowaidi, El-Saddik, and Liu, 2013). This haptic substitution system aims at helping patients that suffer from the symptoms of Diabetic Peripheral Sensory Neuropathy, such as sensory loss at the fingertips and resulting difficulties with manipulating objects. F-Glove uses audio feedback to inform the patients of the pressure they apply to objects. The volume of the sound is mapped linearly to the applied pressure. Unfortunately, it is not clear whether the system simply informs the patients of the pressure they use via sound or whether they start experiencing pressure directly, via the auditory sense. Naturally, the experience of the cube is quite different from not having a sense of touch, as your hand can simply reach through the virtual sonic object.

6.2.5 Open Air Instruments and Sound Installations

Our project relates to the field of sonic interaction. In particular, it relates instruments and installations that use hand or body gestures in free space to produce sound, such as the above-mentioned Theremin. Like our research, such gesture instruments and installations are based on a mapping between body movement and sound.

The artwork *Very Nervous System* (1986-1990) by David Rokeby (1986-1990) is an early example of an interactive sound installation where body movement in open space generates sound. However, the sound of such artworks and instruments like the Theremin usually does not appear to originate from the location of the movement, which is a key difference from sonically tangible objects. Furthermore, with few exceptions, they do not (try to) express the presence of virtual objects in space.

One exception—an instrument that actually does convey the presence of virtual objects in space—is the invisible drumkit by Demian Kappenstein and Marc Bangert (*The Invisible Drums of Demian Kappenstein and Marc Bangert* 2011). In their invisible setup, each virtual drum is placed at its regular position in space. Hitting the invisible virtual drums triggers pre-recorded samples of a real drum set. The position of the sticks and the speed of the movement determine which sample is triggered. Similarly to the cube, the virtual drum kit becomes perceivable through the interaction. Furthermore, based on our own experience with a video of the performance, the presence of the virtual drum-kit in space also becomes apparent through witnessing (perceiving) this interaction. However, to the best of our knowledge, sounds do not seemingly originate from the location of the drums. This makes the project fundamentally different from the invisible cube.

6.2.6 *Science Fiction*

The sonically tangible cube is unlike any real object and it is experienced differently from how we perceive real objects. We can find similar ideas of objects that differ from physical objects in fiction and science fiction. For example, the film Ghostbusters (Reitman, 2004) features ghosts whose presence can be sensed with the help of custom devices. A key difference between any fictive objects we know from movies, stories and books is the fact that one experiences our cube's presence in space directly, through a new sensory combination of touch and sound. We are not aware of any such object being described or depicted in stories. However, we believe there is much to learn about how virtual objects could look or behave from the domain of fiction.

6.3 *Reflection and Outlook*

With the sonically tangible cube, we have introduced a prototype of a sonically tangible object and a new, sound-based form of augmented reality. The proposed cube is invisible and non-tactile. According to our experience, it is nonetheless perceived as spatially present in our real, physical environment. This suggests that virtual objects do not have to look or feel like real objects in order to be a believable part of

our real, physical space.

The virtual cube is non-tactile and yet tangible. The experience of the cube can be seen as one possible answer to the question of how it could feel to touch an object that provides no tactile feedback. According to our impression, the virtual sonic object offers an almost-tactile experience that has no equivalent in a purely physical world. However, this still has to be confirmed by experiments with unbiased participants.

The current implementation of the cube primarily serves as a proof of concept. While we are happy with its current state, we have many ideas on how to improve the cube and explore the concept of sonically tangible objects further. For instance, it would be interesting to allow the participant to manipulate the cube in space. During presentations of the project, we have repeatedly encountered the question whether one could, e.g., move the cube in space or resize the object.

Concerning the sonic qualities, future experiments can reveal which sounds are most suitable for creating tactile-like experiences and possibly test whether sounds that are created with the hands work best. It would be interesting to find out more about how to sonically represent imaginary material and communicate different densities, textures and shapes with sound.

So far, we have chosen to work with binaural recordings. In the future, it will be valuable to explore computational methods for simulating the sounds' origins in space. If this is successful, it will be much easier to allow participants to move through space freely and experience the cube from different positions. Furthermore, it will be simpler to create polymorphic sonically tangible objects of different shapes and sizes and to place them at various positions and in different spaces.

One aspect that was left aside so far is the topic of (in)visibility. This offers several intriguing directions for future research. For example, we are eager to learn how participants interpret the absence of visual clues. On the one hand, it might lead to a contradiction between senses: "I can hear it, but I see that nothing is there". On the other hand, the lack of visual stimuli could be interpreted as a property of the object: "Something is there, it is invisible". Furthermore, it would be interesting to compare the experience of the cube with open and closed eyes, and, as an additional condition, also add a visual dimension to the cube (e.g., by means of a head-mounted display) to learn more about the influence of (in)visibility on the experience.

One limitation of this research is that so far, our inferences are based on informal tryouts and our own subjective experience with the cube. Our experience might not fully represent how others perceive the cube and we cannot entirely rule out the possibility that it is influenced by the expectations and hopes we have for the project. We hope to extend the presented research and conduct experiments with unbiased participants in the future.

Whereas this study focuses on the experience of the author when interacting with the cube, another interesting direction to pursue on the future is how the cube is experienced by a bystander who witnesses the interaction. Based on our experience when presenting the video of the author's interaction with the cube, we expect that seeing the interaction and hearing the sonic result (even if it is played back on speakers) can create the impression of an invisible object being part of the space. In this context, it would also be interesting to know if bystanders imagine some kind of tactile stimulation when witnessing the interaction.

A limitation of the current setup is that the participant cannot freely move their head while experiencing the cube. This is due to the use of binaural recordings. Another constraint that stems from the fact that recordings were used, is that the cube can only be experienced in the particular office of the researcher and at the particular spot where the sounds have been recorded. Playing the sounds back in another room or at another position would likely sound weird. This is because the qualities of the room (such as the fact that the room is small and windows reflect the sounds from the left) have shaped the recordings. (If the sounds were, e.g., played back in a big room, one would immediately notice the lack of reverb.)

In many ways, sonically tangible objects are the culmination of our preceding research into augmented reality. The project builds on various ideas discussed in earlier chapters. For instance, it takes up our claim that AR is not just something we see, and that virtual content can take non-visual and multimodal forms (cf. chapter 3). It furthermore builds on spatial relationships between the virtual and the real, which have been identified as an important aspect of existing AR projects in chapter 3. In addition, it focuses on the presence of an added, virtual object in real space—a common form of AR that we have discussed in detail in section 4.2. What is more, the project takes up the issue of interaction between participants and virtual content, which has been a topic of interest in the preceding chapter. Similarly, it builds on the idea that virtual objects do not have to follow physical laws and can differ from real objects, that likewise has been a major point of interest in chapter 5.

Although the project has emerged within the context of AR research, it also raises questions that go beyond the field of AR and that fall outside our own area of expertise. For instance, it would be interesting to learn more about what happens on a perceptual level. Are sound and kinesthetic information combined, similarly to how cutaneous information and kinesthetic information are integrated in traditional haptic perception? Can the combination of spatial sound and kinesthetic information lead to cross-modal interactions? What happens if the spatial information of the audio does not match the position of the fingers? Do we perceive the lack of tactile stimuli as

"something missing" or do we fill in this information? We have put much emphasis on describing the concept in a way that allows others to reproduce it and we want to invite researchers to join our investigation of sonically tangible objects.

Subject: Interview

------- Forwarded message -------
From: Hanna Schraffenberger
Date: 2013/4/3 15:53
Subject: Interview
To: Lev Manovich

Dear Lev,

Maybe you remember me from Facebook. I work at the Augmented Reality Lab in The Hague and I am one of the editors of the AR[t] magazine. When I read your article “The Poetics of Augmented Space”, I realized that I would like to interview you about augmented reality for our magazine. A short time ago, I finally also read your book “The Language of New Media”. As a consequence, I’d like to interview you even more. So I hope you’ll agree to an interview?

Best regards,
Hanna

P.S. The questions are in the attachment.

What is Augmented Reality?

To begin with, I would like to ask you what you consider augmented reality (AR) to be. In "The Poetics of Augmented Space" (Manovich, 2006) you describe AR as "the laying of dynamic and context-specific information over the visual field of a user" (p. 222). It would be great if you'd address the topic once more. Firstly, because our readers might not have read your article. And secondly, because I think that this point of view unnecessarily limits AR to the visual sense.

In "The Poetics of Augmented Space", you mention Janet Cardiff's audio walks as great examples of laying information over physical space. These walks are designed for specific walking routes. While navigating the environment, one gets to listen to a mix of edited sounds that blend in with the sounds of the surroundings, as well as spoken narrative elements and instructions such as where to go and what to look at (see Cardiff, n.d., 1991). In contrast to 'typical' visual AR, the user is presented with auditory information that relates to the immediate surrounding space. Personally, I would call this augmented reality. Wouldn't you?

Augmented Space

What is special about AR compared to other forms of Augmented Space? In your article "The Poetics of Augmented Space" you discuss the concept of Augmented Space. Augmented Space refers to all those physical spaces that are overlaid with dynamic information such as shopping malls and entertainment centers that are filled with electronic screens and all those places where one can access information wirelessly on phones, tablets or laptops. Besides AR, you mention several other technological developments in the context of Augmented Space, among which, for example, monitoring, ubiquitous computing, tangible interfaces and smart objects. Is AR just one of many related recent phenomena that play a role in overlaying the physical space with information? What's special about AR compared to other forms of Augmented Space?

What else can be augmented?

Something I really like about your article is that you see augmentation as an idea and a practice rather than a collection of technologies. However, so far, you have only discussed the augmentation of space. I was wondering whether you have considered other manifestations of augmentation as well. I don't think augmentation is limited to a space or an environment. I'd even say that often it's not the space that is augmented, but something else.

For example, you mention software that performs tasks according to the mood, pattern of work, focus of attention or interests of their user. However, I am doubtful whether our experience of a space is affected by this kind of information. Let's imagine that my phone registered that I have been sitting still for a long time and reminds me to take a short break to stretch my legs. This information relates to one individual in the space (me), to the activity the person is performing (sitting still), but I don't think it has anything to do with the surrounding space. Hence, I might consider it an augmentation of the activity (not moving, sitting still) or an augmentation of the user (me), but I don't consider it an augmentation of space.

Edwin van der Heide (my colleague and supervisor) and I have recently given this topic a lot of thought, and we were fascinated by the questions: "What is actually augmented in augmented reality? What else can (we imagine to) be augmented?" We came up with the answer, that in AR, something virtual augments something real. More specifically, the virtual augments that to which it relates. In our view, space is one of the possibilities, but likewise, we have considered things like augmented objects, augmented humans, augmented perception, augmented content and augmented activities. What is augmented depends on what the additional content relates to. I am curious whether you'd agree. Do you think that all forms of augmentation bring along an aug-

mentation of space or influence our experience of the immediate surrounding space?

Information and space—one coherent gestalt?

In "The Poetics of Augmented Space" you raise a question that intrigues me a lot. Do the real space and the dynamically presented information add up to one single coherent phenomenological gestalt or are they processed as separate layers?

I am a bit of a sound-person and it has always fascinated me that sometimes the sounds of a radio seem to mix in with environmental sounds. For example, the ticking of a red streetlight might perfectly mix in with the rhythm of the song that is currently playing. Listening to a radio play, an event could sound so real and so nearby, that I'd turn around, just to find, that nothing is happening there. But of course, most often, the sound of the radio just exists as a separate, independent layer of content. The voice of the newsreader doesn't mix with the voice of my colleague, nor does it relate to my environment. Most of the time, a song is just a song and has nothing to do with the surrounding space— until someone starts dancing or tapping their foot. So judging from my experience of listening to the radio, information and the surrounding space can be perceived as one single mixed thing as well as independently. But besides these two options, there are more possibilities. For example, the newsreader might tell me about a traffic jam and thereby inform me about my immediate physical space. Here, the information and my spatial surroundings aren't perceived as a single gestalt, but nevertheless, there is a relationship between both. I think the same is true for Augmented Space. Often, information and space might be related, even when they don't add up to one phenomenological gestalt. So some questions I'd like you to answer with respect to Augmented Space are: When information and space are perceived independently from each other—would you still call these occurrences Augmented Space? What if information and space are perceived as separate but related layers? And more fundamentally: When and why do information and space add up to one single gestalt?

New Media

One of the main questions I want to ask you is: What makes augmented reality special? I have posed that question with respect to other forms of augmented space. I'd like to ask it again with respect to the history of new media.

Personally, I don't think of AR as a recent phenomenon. Of course, there are more and more so-called AR applications, AR technologies and new media works that work with AR. However, when we consider the concept of AR, we find examples that date back centuries. An example of ancient AR is the Pepper's Ghost trick (Lamers, 2013). It uses a second room, glass and special lighting in order to let objects seem to appear, disappear or morph into each other in an otherwise real, physical environment.

But even if the concept isn't new, current manifestations of AR might still bring something new and special to the table. If we look at contemporary AR and compare that with other forms of new media, what's special about it and what isn't?

AR and the second space

From *The Language of New Media* (Manovich, 2001), I understood that throughout media history, the screen was used to separate two absolutely different spaces. For example, this function of the screen applies equally to renaissance paintings and to modern computer displays. When we imagine a typical AR scenario in which virtual objects are integrated into a real scene (e.g. a virtual bird is sitting on a real tree) there is no second space. It's the same physical space, which appears to contain both virtual and real elements. Is this a fundamental change in visual culture?

AR and the quest for realism

The quest for realism in computer graphics is something that has always bored me. You note that new technological developments illustrate how unrealistic the previous existing images

were. At the same time, they remind us that current images will also be superseded. I was wondering: How does AR fit in the widespread aspiration towards realism? On the one hand, visual AR could be considered a huge step back. The 3D models that are usually integrated into real space don't come close to the level of photorealism we know from cinema. On the other hand, the virtual leaves the realm of virtual space and enters our real physical environments—with respect to that the images might be experienced as more realistic than ever...

Will AR take the quest for realism to a new level? I can imagine, when striving for realism, the virtual things that appear to exist in our physical space should not only look like real things—ideally, they also feel like them, smell like them, taste like them and behave like them. Will photorealism be traded in for a form of realism that encompasses all senses? Do you think new media will develop towards a more multimodal form?

AR and cinema

In *The Language of New Media*, you relate different forms of new media—e.g. Virtual Reality, websites and CD-ROMs—to cinema. How about the relation between AR and cinema?

I'm certainly not a cinema expert, but I guess most of what we see in visual AR has been present in cinema for a long time. For example, AR research is very concerned with registering virtual objects in real space. As far as I understand it, this can be seen as an analogy to compositing in films: an attempt to blend the virtual and the real into a seamless whole 'augmented' reality. Do you agree?

You oppose compositing to montage: while compositing aims to blend different elements into a single gestalt, montage aims to create visual, stylistic, semantic, and emotional dissonance between them. Do we have montage in AR as well? For instance, you give the example of montage within a shot, where an image of a dream appears over a man's sleeping head. The same could easily be done in AR. So I would think, AR can learn from cinema both with respect to compositing and with respect to montage. However, I also wonder: Does cinema use other techniques to create fictional realities that are not (yet) used in AR? Does AR use techniques that might be adapted by cinema in the future?

AR as spatialized databases

One of the main claims in *The Language of New Media* is that at their basis, all new media works are databases. You argue that what artists or designers do when creating a new media work, is constructing an interface to such a database. More specifically, you write about the elements of a database:

"If the elements exist in one dimension (time of a film, list on a page), they will be inevitably ordered. So the only way to create a pure database is to spatialize it, distributing the elements in space." (p. 238)

In AR, virtual and real elements are distributed in real space. Can we understand this as a pure database? What are the consequences of working with spatialized elements? What are the inherent limitations and possibilities when working with this form? (I can imagine it has consequences, e.g. for storytelling? As you point out, we cannot assume that elements will form a narrative when they are accessed in an arbitrary order.)

AR and future research

With *The Language of New Media*, you did not only provide a theory of new media; you also pointed your readers towards aspects of new media that were still relatively unexplored at that time and you suggested directions for practical experimentation. Are there certain aspects of augmented reality you consider especially interesting for future experiments and explorations?

Part III

Conclusion

7 Conclusion

In the last seven chapters, we have addressed the nature and possible manifestations of augmented reality. We have explored AR both theoretically as well as practically and we have applied an unconventionally broad perspective. The investigation has led to various new insights. In this final chapter, we summarize our main results and reflect on our findings. We revisit some of the questions that have surfaced during this trajectory and that we can answer now, after having had a critical look at existing research and after having worked with AR ourselves. Furthermore, we present suggestions for designing AR environments as well as possible directions for future AR research.

7.1 What Is Augmented Reality?

One of the main goals of this thesis was to understand the nature of AR and to answer the question what augmented reality is. So, what is augmented reality? In our opinion, AR is an environment in which a participant experiences a relationship between the virtual and the real. More specifically, AR is concerned with relationships between the virtual and the real physical environment. Since real environments are multimodal by nature, AR environments are also multimodal, even when the virtual content is only mediated by one modality. The relationships between the virtual and the real set AR apart from those environments where the virtual and the real merely coexist and where both are experienced as independent from one another.

In the following, we apply this definition to questions that have surfaced throughout this thesis. We place our view of AR in the context of existing research and emphasize differences. This will illustrate how our understanding of AR differs from common notions in three ways.

7.1.1 From Technologies to Experiences

One of the most prominent understandings of AR in existing research is the idea of AR as a technology. But is AR a technology? According to our definition, the answer is no. Although we believe that technologies enable AR, we do not treat AR as a technology. In our opinion, a reason to change towards a more environment- and experience-focused view is that the ultimate purpose of AR technologies is to allow people

to experience virtual content in relation to an otherwise real environment. If we ultimately aim at creating certain environments and experiences, why define the field in terms of the technologies that enable them rather than in terms of the environments and experiences we are interested in? An environment- and experience- focused definition will hold, even if enabling technologies change or take unforeseen forms.

If one accepts that AR is characterized by the experience of virtual content in relation to the real world, a definition in terms of enabling technologies becomes unfeasible. For one, there is no one single kind of technology that creates such experiences. To mention just a few examples, we have seen projects where a participant listens to prerecorded audio on a simple mobile CD player. Likewise, we have seen setups that allow a participant to see virtual content in real space with a head-mounted display and projection-based setups that present virtual content in the real environment directly. In addition, we have encountered devices that use electric current to change a food's taste or the tactile feeling of a real object. In our opinion, the main thing these various technologies have in common is the experience they evoke.

Furthermore, the same type of technologies can be used for characteristic AR experiences as well as for other purposes. For instance, a CD player can be used to listen to audio walks where virtual sounds mix in with the real environment. However, one can also use it to listen to music and to isolate oneself from the real surroundings. Likewise, we might use a projector to present a movie on a wall, but we can just as well use it to project a slowly expanding crack onto the wall that looks as if actually existed in the real environment. As this shows, the technology alone does not determine whether we are dealing with AR experiences or not.

7.1.2 *From Vision to Multimodal Environments*

Existing understandings of AR are often focused on what a user or participant *sees*. Accordingly, AR is commonly understood in terms of virtual *imagery* that is overlaid onto a user's or participant's *view* of the world. In contrast, our definition of AR suggests that we have to approach AR from a multimodal and all senses-encompassing point of view. We have identified many reasons for this. First of all, we believe that a participant experiences virtual content in relation to the physical world. This physical world is multimodal. As such, the resulting environment entails both virtual as well as multimodal real elements. As briefly mentioned in section 7.1, AR is inherently multimodal because an AR environment includes the multimodal real environment. Aside from this, our definition leaves room for virtual content to take on non-visual and also multimodal forms. In our opinion, there is no good reason to exclude such virtual content from the domain of AR.

Even if one disagrees with our notion—and defines AR in terms of visual overlays—it makes sense to treat AR from a multimodal perspective. This is because also solely visual additions can affect our non-visual impressions of the world. For instance, we have encountered a project where the visual information changes how real objects feel. If we solely focus on what a participant sees, we might affect a person's non-visual experience of the world without being aware of this. We believe that the combination of these arguments makes a compelling case to treat AR in a senses-encompassing way.

7.1.3 From Registration to Relationships

Many AR scenarios are realized by means of an interactive system that aligns virtual and real elements in 3D. If we are to believe general opinions and widespread definitions, this alignment or registration process is necessary for AR. In contrast, our definition does not require registration. Instead, it focuses on relationships between the virtual and the real. Registration can be such a relationship but other possibilities exist as well.

There is no doubt that spatial links between the virtual and the real are at the heart of many AR applications. However, our main argument to define AR in broader terms is that other types of relationships also lead to the augmentation of the physical world. Most notably, the virtual can relate to the real world on a content-level, and e.g., affect our experience of the environment by informing us about our surroundings.

One might disagree with this opinion. However, even if one approaches AR in terms of interactive systems that spatially align virtual and real content in 3D, it still makes sense to look beyond *spatial* registration. This is because such interactive systems typically aim at making it seem as if virtual content existed in the real environment. This goal, however, is not only a matter of spatial alignment. Many other relationships between the virtual and the real can potentially contribute to or harm this underlying goal. For instance, we can imagine that the presence of a virtual creature in the real environment is much more convincing if this creature listens and reacts to the sounds in the environment. At the same time, the illusion of it being present in the space might be harmed if the creature is not affected by real light sources or by real wind, if it is not reflected in real glossy surfaces or if it remains dry when it rains. If we look at current AR research, this idea is acknowledged, but primarily explored with respect to optical effects between the virtual and the real, such as illumination, reflections and shadows. Other types of relationships have still received little attention.

Arguably, a strength of this definition is that it is broader than most common views on AR. We hope this broader perspective will free prac-

titioners and researchers alike from restricting ideas, such as the association of AR with *visual* overlays, and thereby inspire and facilitate new and different forms of both AR and AR research. However, one might also argue that our definition is too broad. For instance, according to our definition, food with synthetic additives or the use of air fresheners in a real space could be considered examples of AR. Likely, few readers will agree with such a broad notion of AR. However, we believe that considering such extremes is important because it shows us how normal and commonplace synthesized information has become in our everyday lives already—possibly, we will be equally casual about the presence of virtual objects in real space in the future. On the other hand, we believe that for many purposes, a more narrow definition will better describe the actual focus of an AR project. In this respect, the many proposed subforms of AR (see section 7.2) can be used to describe AR projects more narrowly. [1]

[1] However, where necessary, our definition could also be refined by using a different definition of the virtual. E.g., defining the virtual in terms of computer-generated simulations would exclude examples such as the use of air fresheners and food additives, but likewise, exclude analog audio recordings.

Although we have reached a firm conclusion, our claims should not be taken as proven facts. The question of what AR is—to some degree—will always remain a matter of opinion. We have supported our opinion with arguments. Yet, many might disagree with our view of AR. This is not a problem. However, we hope to nonetheless convey that there is a family of environments in which participants experience relationships between the virtual and the real and that it makes sense to approach this collection of environments as a cohesive field. These points should hold, independently of whether the reader agrees to see this as part of the AR field or not.

7.2 *What Forms Can AR Take?*

A second question that has fueled our exploration is what forms AR can take. The answer to this question depends on the chosen perspective and point of interest. On a fundamental level, we have identified two forms of AR:

- **Presence-based AR**: Here, a participant experiences the presence of virtual content in the real environment. In other words, virtual content seemingly exists in real space, rather than, e.g., on a screen or in a separate virtual world.

- **Content-based AR:** In this form of AR, the virtual relates to the real environment content-wise. This is, e.g., the case when virtual content informs us about our real surroundings or when it tells a story about the real environment.

In both presence-based AR and content-based AR, virtual content is presented in and relates to a real physical environment.

Another way to distinguish between different forms of AR is based on how this virtual content affects its real surroundings. Based on the

role that the virtual content plays in the real environment, we distinguish between the following sub-forms of AR:

- **Extended reality:** Here, the virtual supplements the real. The environment appears to contain more/additional information.
- **Diminished reality:** In this case, the virtual removes the real elements from the perception of the participant. As a result, there seems to exist *less* content in the surroundings.
- **Altered reality:** In this form of AR, the virtual transforms the apparent qualities of the real world. For instance, the virtual might alter the perceived size or shape, weight or texture of real objects. As a consequence, the participant not necessarily perceives more or less information, but instead, perceives different information.
- **Hybrid reality:** Here, the virtual completes the real. It does not serve as 'something additional' and optional but rather is an integral part of an object or environment. A hybrid object/environment would be considered incomplete without the virtual component.
- **Extended perception:** In this case, the virtual translates already present and real but unperceivable aspects of the environment into virtual but perceivable information. As a result, the participant can perceive more aspects of the environment. For instance, a participant might be able to hear radioactive radiation. This form of AR differs from other manifestations because it is primarily concerned with augmenting a participant's perception rather than with augmenting the environment.

Finally, we can distinguish between two different manifestations of AR with respect to how the augmented environment compares to the real world:

- **Imitative augmented reality:** This form of AR mimics reality and, e.g., aims at presenting virtual objects that look and behave like real objects. The ultimate goal of much research in this context is to create AR environments that are indistinguishable from real environments.
- **Imaginative augmented reality:** This type of AR takes the form of new and imaginative environments that have no equivalent in a purely physical world. Research in this context explores the fact that virtual objects do not have to look, feel or behave like real objects.

It should be noted that the above-described forms of AR are neither exhaustive nor exclusive. The different forms can be combined. Even seemingly opposing forms can be united in one AR environment. E.g., an AR environment can mimic the real world when it comes to gravity but also allow virtual objects to move through real walls.

7.3 New Forms of AR

We have started out this trajectory with two main aims: First of all, advancing AR research through a better understanding of AR. Arguably, the above-summarized theory fulfills this goal and contributes to this end. In addition, we have set out to facilitate, create and explore new forms of AR. We have pursued this goal in two contexts.

7.3.1 Introducing New Laws

First of all, we have explored new forms AR with respect to influences between the virtual and the real. Here, we have shown that AR does not have to adhere to physical laws. Instead, we can introduce new laws. Of course, this does not mean that we can make real objects float through space or allow people to walk through physical walls—real elements still follow the laws of our physical world. However, virtual objects can behave differently, and react to the real world in new and imaginative ways. We have demonstrated this by introducing imaginative attractive forces. For instance, we have created an environment where virtual objects are attracted by real objects of a similar color or by light. We see a lot of potential in realizing imaginative influences between the virtual and the real and hope to explore this research direction further in the future.

7.3.2 Introducing New Objects

A second way in which we have explored new forms of AR is by designing a novel kind of virtual object, namely the so-called sonically tangible cube. As we see it, sonically tangible objects do not look, feel or behave like any real object, and they are also perceived differently from how we perceive real objects. Sonically tangible objects can appear to exist in real space, but unlike real objects, they are invisible and non-tactile. The underlying concept is that 'touching' such a virtual object triggers binaural sounds that originate from the exact spot where the object is touched. Our initial experimentation has suggested that this sound-based approach can convey the presence of virtual objects in real space and result in almost-tactile experiences. We believe that when it comes to creating new forms of AR, a main direction to pursue is working with new types of virtual content that does not try to mimic real objects.

In our opinion, the combination of our practical and theoretical exploration reveals many concrete insights into what AR is and what else it potentially can be.

7.4 Pending Questions

Our review of existing AR literature (chapter 2) has raised questions that we can answer now—by applying our definition and by looking back at the preceding chapters. For instance, we have seen that little consensus exists on what is augmented in AR. In accordance with our proposed definition, we suggest that the virtual augments that to which it relates. More importantly, the virtual and the real relate to, add to and augment one another. During this trajectory, we have, among others, encountered scenarios where the virtual augments a specific physical object, where the virtual augments the general environment, where virtual content augments humans and where it augments media content presented in books or music playing on the radio.

Another question that has surfaced in the beginning and that we can answer now concerns the role of the participant. Do we have to be present in an augmented environment to experience AR? Is AR something, we can watch on television or is it something we have to interact with and engage with more actively? According to our definition, AR results from experiencing relationships between one's real surroundings and virtual content in this environment. This entails that the participant is part of the environment. However, just like we can experience some aspects of a physical environment in a mediated form, we might also be able to experience some aspects of augmented reality in a mediated form—for instance, when watching a video of someone else's AR experience online.

7.5 Limitations and Concerns

This thesis focuses on the conceptual characteristics and possibilities of AR. In contrast, technological issues, such as how to technically implement AR or advance AR systems, fall out of the scope of this thesis. Although we have addressed AR both in depth as well as in breadth, our research has some limitations.

With respect to methodology, one limitation is that our practical observations and propositions are based on our own, subjective experiences. For instance, we have assessed that virtual objects do not have to behave like real objects in order to appear as a believable part of real space. Likewise, we have concluded that sonically tangible objects create an almost tactile-like and new experience. However, these propositions are largely based on our own experience. Naturally, our own experiences might have been biased. We cannot rule out that our initial expectations and intentions have contributed to our resulting experience. Furthermore, people are different and our own experience might not represent how other participants perceive AR. These issues are especially relevant because we have argued that AR is the result of the *experienced* relationships between the virtual and the real. It hence

would be very desirable to study how others experience our proposed AR scenarios.

Another limitation concerns the technological implementations. So far, most of our projects have been realized with rather cheap equipment in a controlled office environment. Furthermore, we have limited the complexity of all projects by determining one fixed point of view from which the augmented environment can be perceived. Although we have shown that several concepts are feasible in this specific context, it remains open whether similar ideas can be implemented in real-world settings that are not as predictable and that poses additional challenges, such as a moving participant.

When it comes to our theoretical approach, a concern is that we have made inferences about AR experiences from studying textual or visual descriptions of AR research projects. Unfortunately, such descriptions often focus on other aspects, such as the technological workings of an AR system. Hence, our assumptions about the resulting experiences might not always be correct.

Like every printed publication about AR, our thesis faces the challenge of describing a fast-moving field. There is no way to prevent this: by the time this thesis reaches the reader, AR technology will have advanced and additional relevant publications will have appeared. However, it is also great to see that since originally submitting this thesis and finalizing it, more experience-focused and modalities-encompassing views have emerged. For instance, the recent book "Augmented Human" by Papagiannis (2017) shares our multimodal approach to AR and—like this thesis—looks beyond the mere technological aspects of AR.

In this thesis, we have challenged many prevailing views and opinions about AR, such as the idea that AR overlays virtual imagery onto a user's view. It should be noted that many of the reviewed claims have been presented in the context of a specific AR project and with no aspiration of describing AR in a more general sense. While we have challenged such views on a general level, we do not mean to critique them on an individual level. E.g., the claim that AR technology overlays virtual images onto a user's view makes sense in the context of a project that works with such a technology. It is only natural that many authors only describe what is relevant to their project, rather than the general field of AR. With respect to this, we believe this thesis fills a gap: we are not aware of any AR publication that presents such a comprehensive overview of the general field.

It stands out that many of our conclusions mirror our point of departure. For instance, we have approached AR with the idea that it engages all human senses and subsequently, have arrived at the exact same conclusion. Of course, this raises the concern of circular reasoning—have we arrived at our conclusions because we have been assuming them all along? In our opinion, this is not the case. Rather,

we have explored what AR entails if we apply a broader view. In our opinion, this exploration has revealed a complex but coherent image of the AR landscape that reaffirms the value of our chosen perspective. Hence, we conclude that our point of view does make sense. This, however, does not mean that it is the only valid view. We believe the contrary is the case: our perspective on AR can complement rather than replace existing notions.

7.6 *Creating AR*

AR is not only a research field but also of interest to artists, designers and developers. When it comes to creating AR experiences, we have arrived at some insights that can guide and inform design processes. We will quickly summarize these points:

- Creating AR experiences concerns more than designing virtual content for the real world. Namely, it involves the design of the relationships between the virtual and the real.
- The physical component/environment does not have to be taken for what it is. It can be (re-)designed as well.
- AR environments are not something we see but something we experience with all our senses. Virtual content can take non-visual and multimodal forms and react to non-visual properties of the real world.
- AR environments are not something we consume rather passively, like watching a movie. Instead, they are environments we interact with. AR environments should be designed to facilitate action in and interaction with the environment.
- AR does not have to mimic reality. We can create new forms of environments, introduce new laws and create virtual objects that do not imitate real objects.

To summarize, designers can give shape to the virtual, the real, as well as to the relationship between the two. We hope that a better theoretical understanding of AR will inform AR practice and development and lead to new and exciting AR works.[2]

[2] In this thesis, we have identified various examples of interactive applications that defy prevailing definitions of AR but yet, augment our experience of our physical surroundings. This shows that narrow definitions not necessarily prevent practitioners to think outside of the box and to come up with different forms of (arguably) augmented reality. Yet, we expect that a better and broader understanding of AR will highlight those possibilities and hopefully, inspire even more and new forms of AR.

7.7 *The Future of AR and AR Research*

Our investigation of AR has raised many issues that could be addressed in the future. First and foremost, it would be desirable to conduct empirical studies with unbiased participants. Such experiments could not only be used to validate our findings but also to obtain new insights into AR experiences. In our opinion, it would be particularly

interesting to address the perceptual goals of AR with empirical studies. For instance, AR often aims at making it seem as if virtual objects existed in the real environment. We believe that future research could more consistently measure whether this goal is met. Furthermore, it should systematically explore which factors contribute to the experience of virtual objects existing in real space. For instance, does it harm our experience if virtual objects are not reflected in real-world objects? Does it benefit our experience if virtual creatures react to sounds in the environment? A first step towards this goal will be to develop and adopt methods that can measure the presence of virtual objects in the real environment. While VR research has established and widely adopted questionnaires to measure a participant's presence in a virtual environment (see, e.g., Witmer and Singer, 1998), AR research—to the best of the author's knowledge—does not (yet) have similarly established and adopted methods to measure the perceived presence of a virtual object in real space.[3] Although the question whether virtual objects are experienced as present in real space differs substantially from the question whether participants feel present in a virtual environment, existing VR research on presence and telepresence (e.g., Sheridan (1992), Witmer and Singer (1998), Steuer (1992) and Schubert et al. (2001)), can serve as a point of departure for AR research into the presence of virtual content in real space. This is because many factors relevant for presence in VR might also be relevant for making objects appear as if they were present in the real world. For instance, interactivity and vividness (as proposed by Steuer (1992) in the context of VR) might also play a role in how present virtual content appears in real space.

Another issue that would benefit from an empirical study is the concept of believability. Virtual objects do not have to adhere to physical laws, and AR can take new and imaginative forms. However, not everything that can be realized technologically is also credible. It would be interesting to gain better insights into what forms of AR are accepted as believable, and what factors affect whether an environment is perceived as credible.

On a more general level, we believe future AR projects will benefit from establishing more influences and interactions between virtual content and the real world. First of all, such influences can potentially support the common goal of making it seem as if virtual objects existed in the real environment. Presumably, if a virtual object reacts to a physical object, this can heighten the impression that both objects exist in the same space. What is more, influences between the virtual and the real can facilitate the often-desired interaction between a participant and virtual content: If the virtual reacts to the real world, a participant can interact with virtual objects by interacting with the real world. For instance, in one of our projects, a participant can move real colored objects and thereby, play with virtual colored objects.

[3] A questionnaire for measuring a virtual object's presence in the real world has been proposed by Regenbrecht and Schubert (2002). However, as of 12th February 2018, its adoption in AR research is quite low. To give an impression: According to Google Scholar, Witmer and Singer's paper that proposes a questionnaire to measures a user's presence in a virtual environment currently counts 3362 citations. In contrast, the AR presence questionnaire by Regenbrecht and Schubert (2002), which focuses on a virtual object's presence in the real world, currently has 19 citations.

Furthermore, future projects can benefit from incorporating both multimodal virtual content as well as taking multimodal qualities of the real world into account. This thesis has sketched out ideas that are just waiting to be realized, such as virtual leaves that fly in real wind and virtual creatures that can be lured closer by making sound. In line with this, we believe future projects can take up the idea that virtual elements can sense the world as well as act in and react to the world.

Whereas much AR research and development mimics our physical reality, we believe much potential lies in imaginative forms of augmented reality. If we imitate a real environment, we know in advance how the result will turn out if we succeed. If we try to create something that does not yet exist, the outcome is uncertain and might surprise us. In his vision about the ultimate display (a room in which a computer controls the existence of matter), the "father of computer graphics" Sutherland concludes that an ultimate display "could literally be the Wonderland into which Alice walked" (p. 2). Augmented reality is no ultimate display. Yet, is has the power to transform our everyday reality into a wonderland. We have shown that AR can use new laws, introduce new types of objects into this world and consequently, facilitate new kinds of experiences. In L. Carroll's wonderland, Alice experiences herself grow enormously after eating a magical cake. Consequently, she is so surprised that she momentarily forgets how to speak proper English and exclaims: "Curious and curiouser!" (L. Carroll, 2015, p. 13). As an AR community, let us go down the rabbit hole and make sure things get curious and curiouser!

References

Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., and Hopper, A. (2001). "Implementing a sentient computing system". *Computer*, 34(8), pp. 50–56 (cited on p. 124).

Agusanto, K., Li, L., Chuangui, Z., and Sing, N. W. (2003). "Photorealistic Rendering for Augmented Reality Using Environment Illumination". In: *Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR '03)*. Washington, DC, USA: IEEE Computer Society, pp. 208–216 (cited on p. 134).

Akamatsu, M. (2014). *aka.leapmotion (Version 0.21)*. Max external. Creative Commons Attribution 3.0 Unported License. URL: http://akamatsu.org/aka/max/objects/ (cited on p. 176).

Akman, O. (2012). "Robust augmented reality". PhD thesis. Delft University of Technology (cited on p. 91).

Alliban, J. (n.d.). Online; accessed November 1, 2018. URL: https://jamesalliban.wordpress.com/2011/03/30/konstruct-ar-iphone-app/ (cited on pp. 152, 154).

Antle, A. N., Marshall, P., and van den Hoven, E. (2011). "Embodied Interaction: Theory and Practice in Human-Computer Interaction". In: *Workshop Proceedings CHI 2011*. ACM, pp. 8–12 (cited on p. 189).

Aoki, T., Matsushita, T., Iio, Y., Mitake, H., Toyama, T., Hasegawa, S., Ayukawa, R., Ichikawa, H., Sato, M., Kuriyama, T., et al. (2005). "Kobito: virtual brownies." In: *Siggraph Emerging Technologies*, p. 11 (cited on pp. 146, 154).

ARToolKit SDK (1999). Computer Software. v5.3.2. URL: https://artoolkit.org/download-artoolkit-sdk (cited on p. 43).

Augmented Reality (2005). Oxford English Dictionary (cited on pp. 23, 31, 37, 57).

Augmented reality (n.d.). *Augmented reality—Wikipedia, The Free Encyclopedia*. Online; accessed October 18, 2017. URL: https://en.wikipedia.org/wiki/Augmented_reality (cited on pp. 24, 29).

Auvray, M., Hanneton, S., and O'Regan, J. K. (2007). "Learning to perceive with a visuo—auditory substitution system: localisation and object recognition with 'The Voice'". *Perception*, 36(3), pp. 416–430 (cited on p. 114).

Azuma, R. (1997). "A Survey of Augmented Reality". *Presence*, 6(4), pp. 355–385 (cited on pp. 9, 19, 22, 25, 26, 28, 29, 34, 39–42, 45, 46, 51, 59, 79, 104, 115, 141, 188).

Azuma, R., Baillot, Y., Behringer, R., Feiner, S., Julier, S., and MacIntyre, B. (2001). "Recent advances in augmented reality". *Computer Graphics and Applications, IEEE*, 21(6), pp. 34–47 (cited on pp. 20, 22, 25, 27, 41, 42, 51, 67, 68, 99, 148).

Bach-y-Rita, P. and Kercel, S. W. (2003). "Sensory substitution and the human–machine interface". *Trends*

in cognitive sciences, 7(12), pp. 541–546 (cited on p. 114).

Bajura, M., Fuchs, H., and Ohbuchi, R. (1992). "Merging virtual objects with the real world: Seeing ultrasound imagery within the patient". In: *ACM SIGGRAPH Computer Graphics*. Vol. 26. 2. ACM, pp. 203–210 (cited on pp. 91, 150).

Bajura, M. and Neumann, U. (1995). "Dynamic registration correction in video-based augmented reality systems". *IEEE Computer Graphics and Applications*, 15(5), pp. 52–60 (cited on pp. 41, 42).

Bandyopadhyay, D., Raskar, R., and Fuchs, H. (2001). "Dynamic shader lamps: Painting on movable objects". In: *Proceedings of the IEEE and ACM International Symposium on Augmented Reality*. IEEE, pp. 207–216 (cited on p. 81).

Bane, R. and Hollerer, T. (2004). "Interactive tools for virtual x-ray vision in mobile augmented reality". In: *Proceedings of the Third IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2004)*. IEEE, pp. 231–239 (cited on pp. 115, 150).

Barakonyi, I., Psik, T., and Schmalstieg, D. (2004). "Agents that talk and hit back: animated agents in augmented reality". In: *Proceedings of the third IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2004)*. IEEE, pp. 141–150 (cited on pp. 125, 161).

Bau, O. and Poupyrev, I. (2012). "REVEL: tactile feedback technology for augmented reality". *ACM Transactions on Graphics (TOG)*, 31(4), p. 89 (cited on pp. 4, 31, 60, 105, 108, 148, 149, 169).

Bauer, C. and Waldner, F. (2013). "Reactive music: when user behavior affects sounds in real-time". In: *Extended Abstracts on Human Factors in Computing Systems (CHI'13)*. ACM, pp. 739–744 (cited on p. 103).

Bederson, B. B. (1995). "Audio augmented reality: a prototype automated tour guide". In: *Conference companion on Human factors in computing systems*. ACM, pp. 210–211 (cited on pp. 4, 31, 52, 79).

Bellamy, C. and 6, P. (2012). *Principles of methodology: Research design in social science*. Sage (cited on p. 77).

Benko, H., Wilson, A. D., and Zannier, F. (2014). "Dyadic projected spatial augmented reality". In: *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, pp. 645–655 (cited on pp. 45, 92).

Bennett, C. A. and Rey, P. (1972). "What's so hot about red?" *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 14(2), pp. 149–154 (cited on p. 110).

Berry, P. C. (1961). "Effect of colored illumination upon perceived temperature." *Journal of Applied Psychology*, 45(4), pp. 248–250 (cited on p. 110).

Bertelson, P. and De Gelder, B. (2004). "The psychology of multimodal perception". *Crossmodal space and crossmodal attention*, pp. 141–177 (cited on p. 112).

Bianchi, G., Knoerlein, B., Szekely, G., and Harders, M. (2006). "High precision augmented reality haptics". In: *Proceedings EuroHaptics*. Vol. 6. Citeseer, pp. 169–178 (cited on pp. 5, 61, 97).

Billinghurst, M. (2001). "Crossing the chasm". In: *International Conference on Augmented Tele-Existence (ICAT 2001)*. Citeseer, pp. 5–7 (cited on pp. 149, 169).

Billinghurst, M., Clark, A., and Lee, G. (2015). "A survey of augmented reality". *Foundations and Trends in Human-Computer Interaction*, 8(2-3), pp. 73–272 (cited on p. 40).

Billinghurst, M., Grasset, R., Seichter, H., and Dünser, A. (2009). "Towards ambient augmented reality with tangible interfaces". *Human-Computer Interaction. Ambient, Ubiquitous and Intelligent Interaction*, pp. 387–396 (cited on p. 149).

Billinghurst, M., Kato, H., and Myojin, S. (2009). "Advanced interaction techniques for augmented reality applications". *Virtual and Mixed Reality*, pp. 13–22 (cited on pp. 151, 152, 154).

Billinghurst, M., Kato, H., and Poupyrev, I. (2001). "The MagicBook - Moving Seamlessly between Reality and Virtuality". *IEEE Computer Graphics and Applications*, 21(3), pp. 6–8 (cited on p. 4).
– (2008). "Tangible augmented reality". In: *ACM SIGGRAPH ASIA 2008 courses*. ACM, pp. 1–10 (cited on pp. 78, 149, 169).

Bimber, O., Encamacao, L., and Schmalstieg, D. (2000). "Real mirrors reflecting virtual worlds". In: *Proceedings IEEE Virtual Reality 2000*. IEEE, pp. 21–28 (cited on p. 119).

Bimber, O. and Raskar, R. (2005). *Spatial Augmented Reality: Merging Real and Virtual Worlds*. Natick, MA, USA: A. K. Peters, Ltd. (cited on pp. 23, 41, 45, 51, 92).

Biocca, F., Kim, J., and Choi, Y. (2001). "Visual Touch in Virtual Environments: An Exploratory Study of Presence, Multimodal Interfaces, and Cross-Modal Sensory Illusions". *Presence: Teleoperators and Virtual Environments*, 10(3), pp. 247–265 (cited on p. 192).

Blum, J. R., Greencorn, D. G., and Cooperstock, J. R. (2012). "Smartphone sensor reliability for augmented reality applications". In: *Mobile and Ubiquitous Systems: Computing, Networking, and Services*. Springer, pp. 127–138 (cited on p. 44).

Bondo, J., Barnard, D., Burcaw, D., Novikoff, T., Kemper, C., Parrish, C., Peters, K., Siebert, J., Wilson, E., Wilson, B., et al. (2010). *iPhone User Interface Design Projects*. Apress (cited on p. 103).

Breen, D. E., Whitaker, R. T., Rose, E., and Tuceryan, M. (1996). "Interactive occlusion and automatic object placement for augmented reality". In: *Computer Graphics Forum*. Vol. 15. 3. Wiley Online Library, pp. 11–22 (cited on pp. 55, 83, 118, 142, 143).

Broll, W., Lindt, I., Ohlenburg, J., Wittkämper, M., Yuan, C., Novotny, T., Mottram, C., Fatah gen Schieck, A., and Strothman, A. (2004). "Arthur: A collaborative augmented environment for architectural design and urban planning" (cited on p. 118).

Brown, L. G. (1992). "A survey of image registration techniques". *ACM computing surveys (CSUR)*, 24(4), pp. 325–376 (cited on p. 21).

Buchmann, V., Violich, S., Billinghurst, M., and Cockburn, A. (2004). "FingARtips: Gesture Based Direct Manipulation in Augmented Reality". In: *Proceedings of the 2nd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia (GRAPHITE '04)*. Singapore: ACM, pp. 212–221 (cited on p. 40).

Bujas, Z. (1971). "Electrical taste". In: *Taste*. Springer, pp. 180–199 (cited on pp. 60, 96, 108).

Caarls, J. (2009). "Pose estimation for mobile devices and augmented reality". PhD thesis. Delft University of Technology (cited on p. 44).

Cardiff, J. (n.d.). *Introduction to the audio walks*. Online; accessed November 7, 2014. URL: http://www.cardiffmiller.com/artworks/walks/audio_walk.html (cited on pp. 34, 47, 48, 51, 127, 178, 198).
– (1991). *Forest Walk*. Audio walk. Online; accessed November 7, 2014. URL: http://www.cardiffmiller.com/artworks/walks/forest.html (cited on pp. 34, 47, 48, 60, 198).

Carroll, J. and Polo, F. (2013). "Augmented Reality Gaming with Sphero". In: *ACM SIGGRAPH 2013 Mobile*. Anaheim, California: ACM, p. 17 (cited on pp. 147, 159).

Carroll, L. (2015). *Alice's adventures in wonderland*. Macmillan Children's Books (cited on p. 213).

Cassell, J., Bickmore, T., Campbell, L., Vilhjálmsson, H., and Yan, H. (2000). "Conversation as a system framework: Designing embodied conversational agents". *Embodied conversational agents*, pp. 29–63 (cited on p. 142).

Caudell, T. and Mizell, D. (1992). "Augmented reality: an application of heads-up display technology to manual manufacturing processes". In: *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*. Vol. 2. IEEE, pp. 659–669 (cited on pp. 3, 4, 30, 34, 57, 58, 80, 86, 90).

Chae, C. and Ko, K. (2008). "Introduction of Physics Simulation in Augmented Reality". In: *Proceedings of the 2008 International Symposium on Ubiquitous Virtual Reality (ISUVR '08)*. Washington, DC, USA: IEEE Computer Society, pp. 37–40 (cited on pp. 118, 145, 153, 156, 173).

Chan, L.-W., Kao, H.-S., Chen, M. Y., Lee, M.-S., Hsu, J., and Hung, Y.-P. (2010). "Touching the Void: Direct-touch Interaction for Intangible Displays". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Atlanta, Georgia, USA: ACM, pp. 2625–2634. URL: http://doi.acm.org/10.1145/1753326.1753725 (cited on p. 190).

Chatzidimitris, T., Gavalas, D., and Michael, D. (2016). "SoundPacman: Audio augmented reality in location-based games". In: *2016 18th Mediterranean Electrotechnical Conference (MELECON)*. IEEE, pp. 1–6 (cited on pp. 45, 60, 65, 80, 84, 92, 93, 102, 123, 188).

Chien, C.-H., Chen, C.-H., and Jeng, T.-S. (2010). "An interactive augmented reality system for learning anatomy structure". In: *Proceedings of the International MultiConference of Engineers and Computer Scientists, IMECS 2010*, pp. 17–19 (cited on p. 3).

Cook, P. R. (2002). *Real sound synthesis for interactive applications*. CRC Press (cited on p. 157).

Corbett-Davies, S., Dünser, A., and Clark, A. (2012). "An interactive augmented reality system for exposure treatment". In: *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality (ISMAR-AMH 2012)*. IEEE, pp. 95–96 (cited on pp. 91, 101, 102, 118, 148).

Corbett-Davies, S., Dünser, A., Green, R., and Clark, A. (2013). "An advanced interaction framework for augmented reality based exposure treatment". In: *Proceedings IEEE Virtual Reality 2013*. IEEE, pp. 19–22 (cited on pp. 148, 154, 173).

Craig, A. B. (2013). *Understanding Augmented Reality: Concepts and Applications*. Newnes (cited on pp. 22, 24–26, 29, 41, 43, 45, 59, 67, 68, 148).

Cui, Y. and Ge, S. S. (2003). "Autonomous vehicle positioning with GPS in urban canyon environments". *IEEE transactions on robotics and automation*, 19(1), pp. 15–25 (cited on p. 44).

Dialectical research (n.d.). *Dialectical research—Wikipedia, The Free Encyclopedia*. Online; accessed November 1, 2018. URL: https://en.wikipedia.org/wiki/Dialectical_research (cited on p. 10).

Dourish, P. (2004). *Where the action is: the foundations of embodied interaction*. MIT press (cited on p. 189).

Doyle, S., Dodge, M., and Smith, A. (1998). "The potential of web-based mapping and virtual reality technologies for modelling urban environments". *Computers, Environment and Urban Systems*, 22(2), pp. 137–155 (cited on pp. 20, 24).

Drascic, D. and Milgram, P. (1996). "Perceptual issues in augmented reality". In: *Electronic Imaging: Science & Technology*. International Society for Optics and Photonics, pp. 123–134 (cited on p. 41).

Drettakis, G., Robert, L., and Bougnoux, S. (1997). "Interactive common illumination for computer

augmented reality". In: *Rendering Techniques '97*. Springer, pp. 45–56 (cited on p. 54).

Eckel, G. (2001). "Immersive audio-augmented environments: the LISTEN project". In: *Proceedings of the 5th International Conference on Information Visualization (IV2001)*. IEEE, pp. 571–573 (cited on p. 93).

Engadget (2016). URL: https://www.engadget.com/2016/03/14/samsung-entrim-4d-vr-motion-headphones/ (cited on p. 110).

Erens, C. (n.d.). *The Audible Space*. Online; accessed November 7, 2014. URL: http://www.cilia-erens.nl/cilia-erens-2/?lang=en (cited on pp. 34, 188).

Ernst, M. O. and Bülthoff, H. H. (2004). "Merging the senses into a robust percept". *Trends in cognitive sciences*, 8(4), pp. 162–169 (cited on p. 61).

Fanger, P., Breum, N., and Jerking, E. (1977). "Can colour and noise influence man's thermal comfort?" *Ergonomics*, 20(1), pp. 11–18 (cited on p. 110).

Feiner, S., MacIntyre, B., Höllerer, T., and Webster, A. (1997). "A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment". *Personal Technologies*, 1(4), pp. 208–217 (cited on p. 85).

Feiner, S., Macintyre, B., and Seligmann, D. (1993). "Knowledge-based augmented reality". *Communications of the ACM*, 36(7), pp. 53–62 (cited on pp. 4, 44, 58, 86, 90, 116).

Féral, J. (2012). *How to define presence effects: the work of Janet Cardiff*. Ed. by G. Giannachi, N. Kaye, and M. Shanks. Routledge, pp. 29–49 (cited on p. 188).

Fischer, J., Bartz, D., and Straber, W. (2005). "Stylized augmented reality for improved immersion". In: *Proceedings IEEE Virtual Reality 2005*. IEEE, pp. 195–202 (cited on pp. 103, 127).

Fujihata, M. (1995). *Beyond Pages*. Interactive installation (cited on p. 146).

Gelenbe, E., Hussain, K., and Kaptan, V. (2005). "Simulating autonomous agents in augmented reality". *Journal of Systems and Software*, 74(3), pp. 255–268 (cited on p. 125).

Gibson, S. and Chalmers, A. (2003). "Photorealistic augmented reality". In: *Eurographics 2003 Tutorial*. Granada, Spain (cited on pp. 134, 173).

Gómez-Maureira, M. A., Teunisse, C., Schraffenberger, H., and Verbeek, F. (2014). "Illuminating Shadows: Introducing Shadow Interaction in Spatial Augmented Reality". In: *Creating the Difference: Proceedings of the Chi Sparks 2014 Conference*, pp. 11–18 (cited on pp. 81, 82).

Gradman, E. (2010). *Cloud Mirror*. Interactive art installation. Online; accessed November 7, 2014. URL: http://www.gradman.com/cloudmirror (cited on pp. 86, 87).

Graham, M., Zook, M., and Boulton, A. (2013). "Augmented reality in urban places: contested content and the duplicity of code". *Transactions of the Institute of British Geographers*, 38(3), pp. 464–479 (cited on p. 24).

Hafidh, B., Osman, H. A., Alowaidi, M., El-Saddik, A., and Liu, X. P. (2013). "F-Glove: A glove with force-audio sensory substitution system for diabetic patients". In: *Proceedings IEEE International Symposium on Haptic Audio Visual Environments and Games (HAVE 2013)*. Washington, DC, USA: IEEE Computer Society, pp. 34–38 (cited on p. 192).

Hampshire, A., Seichter, H., Grasset, R., and Billinghurst, M. (2006). "Augmented Reality Authoring: Generic Context from Programmer to Designer". In: *Proceedings of the 18th Australia Conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments (OZCHI '06)*. Sydney, Australia:

ACM, pp. 409–412 (cited on p. 73).

Han, J. and Gold, N. (2014). "Lessons Learned in Exploring the Leap Motion™ Sensor for Gesture-based Instrument Design". In: *New Interfaces for Musical Expression 2014*. Goldsmiths University of London, pp. 371–374 (cited on p. 185).

Hashimoto, Y., Nagaya, N., Kojima, M., Miyajima, S., Ohtaki, J., Yamamoto, A., Mitani, T., and Inami, M. (2006). "Straw-like user interface: virtual experience of the sensation of drinking using a straw". In: *Proceedings of the 2006 ACM SIGCHI international conference on Advances in computer entertainment technology*. ACM, p. 50 (cited on p. 101).

Hassenzahl, M. and Tractinsky, N. (2006). "User experience-a research agenda". *Behaviour & information technology*, 25(2), pp. 91–97 (cited on p. 11).

Heller, F. and Borchers, J. (2011). "Corona: Audio augmented reality in historic sites". In: *Proceedings 1st Workshop on Mobile AR, MobileHCI, Sweden* (cited on pp. 92, 93).

Herling, J. and Broll, W. (2010). "Advanced self-contained object removal for realizing real-time diminished reality in unconstrained environments". In: *Proceedings of the 9th IEEE International Symposium on Mixed and Augmented Reality (ISMAR 2010)*. IEEE, pp. 207–212 (cited on pp. 27, 81, 99, 100).

– (2011). "Markerless Tracking for Augmented Reality". In: *Handbook of Augmented Reality*. Ed. by B. Furht. New York, NY: Springer New York, pp. 255–272. URL: http://dx.doi.org/10.1007/978-1-4614-0064-6_11 (cited on p. 26).

Hirano, Y., Kimura, A., Shibata, F., and Tamura, H. (2011). "Psychophysical influence of mixed-reality visual stimulation on sense of hardness". In: *Proceedings IEEE Virtual Reality 2011*. IEEE, pp. 51–54 (cited on pp. 62, 105, 112).

history and archive - Stedelijk Museum Amsterdam (n.d.). URL: http://www.stedelijk.nl/en/artours/history-and-archive (visited on 01/22/2017) (cited on pp. 53, 87).

Ho, H.-N., Iwai, D., Yoshikawa, Y., Watanabe, J., and Nishida, S. (2014). "Combining colour and temperature: A blue object is more likely to be judged as warm than a red object". *Scientific reports*, 4, p. 5527 (cited on pp. 62, 105, 112).

Höllerer, T. and Feiner, S. (2004). "Mobile augmented reality". *Telegeoinformatics: Location-Based Computing and Services* (cited on p. 79).

Hopper, A. (1999). "The Royal Society Clifford Paterson Lecture". *Sentient Computing. AT&T Lab Cambridge Technical Report 1999.12, 1999*, pp. 1–10 (cited on p. 124).

Hoshi, T., Takahashi, M., Iwamoto, T., and Shinoda, H. (2010). "Noncontact tactile display based on radiation pressure of airborne ultrasound". *IEEE Transactions on Haptics*, 3(3), pp. 155–165 (cited on p. 94).

Hoshi, T., Takahashi, M., Nakatsuma, K., and Shinoda, H. (2009). "Touchable holography". In: *ACM SIGGRAPH 2009 Emerging Technologies*. ACM, p. 23 (cited on pp. 94, 95).

Hötting, K. and Röder, B. (2004). "Hearing cheats touch, but less in congenitally blind than in sighted individuals". *Psychological Science*, 15(1), pp. 60–64 (cited on p. 192).

Hugues, O., Fuchs, P., and Nannipieri, O. (2011). "New Augmented Reality Taxonomy: Technologies and Features of Augmented Environment". In: *Handbook of Augmented Reality*. Ed. by B. Furht. Springer New York, pp. 47–63 (cited on p. 59).

Hürst, W. and Van Wezel, C. (2013). "Gesture-based interaction via finger tracking for mobile augmented reality". *Multimedia Tools and Applications*, 62(1), pp. 233–258 (cited on pp. 152, 154).

Huzaifah Bhutto (2012). *Google Glasses Project*. YouTube video. Online; accessed January 8, 2016. URL: https://www.youtube.com/watch?v=JSnB06um5r4 (cited on p. 46).

Hwang, A. (2006). *Albert Hwang*. YouTube channel. Online; accessed November 7, 2014. URL: https://www.youtube.com/user/phedhex (cited on p. 133).
– (2012). *How to Dance Liquid: Rails*. Video. Online; accessed November 7, 2014. URL: http://youtu.be/Wwxwc_2vDKQ?t=7m50s (cited on p. 134).

Iesaki, A., Somada, A., Kimura, A., Shibata, F., and Tamura, H. (2008). "Psychophysical influence on tactual impression by mixed-reality visual stimulation". In: *Proceedings IEEE Virtual Reality 2008*. IEEE, pp. 265–266 (cited on pp. 62, 106, 112).

IKEA Place (2017). Mobile application software. Version 1.13. URL: https://itunes.apple.com/us/app/ikea-place/id1279244498 (cited on p. 40).

Inception - The App (2016). Mobile application software. Version 2.0. URL: https://www.inception-app.com (cited on pp. 33, 103).

Irawati, S., Green, S., Billinghurst, M., Duenser, A., and Ko, H. (2006). "" Move the couch where?": developing an augmented reality multimodal interface". In: *IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2006)*. IEEE, pp. 183–186 (cited on p. 78).

iReport, C. (2010). URL: http://ireport.cnn.com/docs/DOC-491232 (visited on 03/15/2017) (cited on p. 115).

ISMAR2015 (n.d.). URL: http://ismar2015.vgtc.org/ismar/2015/info/call-participation/call-fullshort-papers/ (cited on p. 27).

Iwamoto, T., Tatezono, M., Hoshi, T., and Shinoda, H. (2008). "Airborne ultrasound tactile display". In: *ACM SIGGRAPH 2008 new tech demos*. ACM, p. 1 (cited on p. 94).

Jacobs, L. F., Arter, J., Cook, A., and Sulloway, F. J. (2015). "Olfactory orientation and navigation in humans". *PloS one*, 10(6), e0129387 (cited on p. 95).

Jonietz, E. (2010). *Augmented Identity*. URL: https://www.technologyreview.com/s/417655/augmented-identity/ (cited on p. 86).

Jousmäki, V. and Hari, R. (1998). "Parchment-skin illusion: sound-biased touch". *Current Biology*, 8(6), R190–R191 (cited on pp. 62, 191).

Kaczmarek, K. A. (1995). "Sensory augmentation and substitution". *CRC handbook of biomedical engineering*, pp. 2100–2109 (cited on p. 114).

Kalkofen, D., Tatzgern, M., and Schmalstieg, D. (2009). "Explosion diagrams in augmented reality". In: *Proceedings IEEE Virtual Reality 2009*. IEEE, pp. 71–78 (cited on p. 150).

Kán, P. and Kaufmann, H. (2012). "High-quality reflections, refractions, and caustics in augmented reality and their contribution to visual coherence". In: *International Symposium on Mixed and Augmented Reality (ISMAR 2012)*. IEEE, pp. 99–108 (cited on pp. 117, 144).

Kanbara, M. and Yokoya, N. (2004). "Real-time estimation of light source environment for photorealistic augmented reality". In: *Proceedings of the 17th International Conference on Pattern Recognition (ICPR 2004)*.

Vol. 2, pp. 911–914 (cited on pp. 134, 144).

Kang, C. and Woo, W. (2011). "ARMate: An Interactive AR Character Responding to Real Objects". In: *Proceedings of the 6th International Conference on E-learning and Games, Edutainment Technologies (Edutainment'11)*. Taipei, Taiwan: Springer-Verlag, pp. 12–19 (cited on pp. 120, 124, 146, 154, 155).

Kaspar, K., König, S., Schwandt, J., and König, P. (2014). "The experience of new sensorimotor contingencies by sensory augmentation". *Consciousness and cognition*, 28, pp. 47–63 (cited on pp. 114, 115).

Kato, H., Billinghurst, M., Poupyrev, I., Imamoto, K., and Tachibana, K. (2000). "Virtual object manipulation on a table-top AR environment". In: *Proceedings of the IEEE and ACM International Symposium on Augmented Reality (ISAR 2000)*. Ieee, pp. 111–119 (cited on pp. 150–152).

Kawashima, T., Imamoto, K., Kato, H., Tachibana, K., and Billinghurst, M. (2001). "Magic paddle: A tangible augmented reality interface for object manipulation". In: *Proceedings on ISMR2001*, pp. 194–195 (cited on p. 152).

Kim, M. and Cheeyong, K. (2015). "Augmented reality fashion apparel simulation using a magic mirror". *International journal of smart home*, 9(2), pp. 169–178 (cited on p. 151).

Kim, S., Kim, Y., and Lee, S.-H. (2011). "On Visual Artifacts of Physics Simulation in Augmented Reality Environment". In: *Proceedings of the 2011 International Symposium on Ubiquitous Virtual Reality (ISUVR '11)*. Washington, DC, USA: IEEE Computer Society, pp. 25–28 (cited on pp. 26, 55, 119, 134, 138, 143, 145, 153, 154, 156, 158, 167, 170, 173).

Klopfer, E. and Squire, K. (2008). "Environmental Detectives—the development of an augmented reality platform for environmental simulations". *Educational Technology Research and Development*, 56(2), pp. 203–228. URL: http://dx.doi.org/10.1007/s11423-007-9037-6 (cited on pp. 24, 27, 30, 68).

Knoerlein, B., Székely, G., and Harders, M. (2007). "Visuo-haptic collaborative augmented reality ping-pong". In: *Proceedings of the international conference on Advances in computer entertainment technology*. ACM, pp. 91–94 (cited on p. 97).

Köster, E. P. (2002). "The specific characteristics of the sense of smell". *Olfaction, taste and cognition*, pp. 27–43 (cited on p. 95).

Kounavis, C. D., Kasimati, A. E., and Zamani, E. D. (2012). "Enhancing the tourism experience through mobile augmented reality: Challenges and prospects". *International Journal of Engineering Business Management*, 4 (cited on p. 26).

Kunst und Medientechnologie Karlsruhe, Z. für (n.d.). URL: http://on1.zkm.de/zkm/werke/BeyondPages (cited on p. 146).

Lamers, M. H. (2013). "Pre-digital Augmented Reality". *AR[t], Augmented Reality, Art and Technology*, 3. Ed. by H. Schraffenberger, M. Kniveton, Y. Kolstee, and J. Verlinden, pp. 24–25. URL: http://arlab.kabk.nl/ar-magazines (cited on pp. 26, 134, 199).

Lawless, H. T., Stevens, D. A., Chapman, K. W., and Kurtz, A. (2005). "Metallic Taste from Electrical and Chemical Stimulation". *Chemical Senses*, 30(3), p. 185. URL: +%20http://dx.doi.org/10.1093/chemse/bji014 (cited on p. 96).

Layar (2009). Mobile application software. URL: http://www.layar.com (cited on pp. 3, 85).

Lederman, S. J. and Klatzky, R. L. (1987). "Hand movements: A window into haptic object recognition". *Cognitive Psychology*, 19(3), pp. 342–368 (cited on pp. 184, 191).

– (2009). "Haptic perception: A tutorial". *Attention, Perception, & Psychophysics*, 71(7), pp. 1439–1459 (cited on p. 191).

Lee, K. (2012). "Augmented Reality in Education and Training". *TechTrends*, 56(2), pp. 13–21 (cited on p. 3).

Leitch, R. and Tokhi, M. (1987). "Active noise control systems". *IEE Proceedings A-Physical Science, Measurement and Instrumentation, Management and Education-Reviews*, 134(6), pp. 525–546 (cited on p. 100).

Leitner, J., Haller, M., Yun, K., Woo, W., Sugimoto, M., and Inami, M. (2008). "IncreTable, a mixed reality tabletop game experience". In: *Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology*. ACM, pp. 9–16 (cited on pp. 146, 154, 158).

Liberati, N. and Nagataki, S. (2015). "The AR glasses' "non-neutrality": their knock-on effects on the subject and on the giveness of the object". *Ethics and Information Technology*, 17(2), pp. 125–137 (cited on pp. 53, 54).

Lindeman, R. W., Lee, G., Beattie, L., Gamper, H., Pathinarupothi, R., and Akhilesh, A. (2012). "GeoBoids: A mobile AR application for exergaming". In: *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality (ISMAR-AMH 2012)*. IEEE, pp. 93–94 (cited on pp. 61, 96, 123, 153, 154).

Lindeman, R. W. and Noma, H. (2007). "A classification scheme for multi-sensory augmented reality". In: *Proceedings of the 2007 ACM symposium on Virtual reality software and technology*. ACM, pp. 175–178 (cited on pp. 61, 67, 118, 144).

Lok, B. C. (2004). "Toward the merging of real and virtual spaces". *Communications of the ACM*, 47(8), pp. 48–53 (cited on p. 99).

Looser, J., Grasset, R., Seichter, H., and Billinghurst, M. (2006). "OSGART-A pragmatic approach to MR". in: *Proceedings of the 5th IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2006)*, pp. 22–25 (cited on p. 73).

López, M. B., Hannuksela, J., Silvén, O., and Vehviläinen, M. (2014). "Interactive multi-frame reconstruction for mobile devices". *Multimedia Tools and Applications*, 69(1), pp. 31–51 (cited on pp. 48, 49).

MacIntyre, B. (2002). "Authoring 3D Mixed Reality Experiences: Managing the Relationship Between the Physical and Virtual Worlds". *Proceedings of the ACM SIGGRAPH and Eurographics Campfire: Production Process of 3D Computer Graphics Applications-Structures, Roles and Tools*, pp. 1–5 (cited on p. 73).

Mackay, W. E. (1996). "Augmenting reality: A new paradigm for interacting with computers". *La recherche*, 284. Online; accessed November 7, 2014, pp. 1–9. URL: http://www-ihm.lri.fr/~mackay/pdffiles/LaRecherche.English.pdf (cited on p. 29).

Madsen, C. B., Jensen, T., and Andersen, M. S. (2006). "Real-time image-based lighting for outdoor augmented reality under dynamically changing illumination conditions". In: *International Conference on Graphics Theory and Applications, GRAPP 2006*, pp. 364–371 (cited on pp. 83, 117, 119, 144, 173).

Maebayashi, A. (1999). *Sonic Interface*. Sound work (cited on pp. 33, 104).

Mann, S. and Fung, J. (2002). "EyeTap devices for augmented, deliberately diminished, or otherwise altered visual perception of rigid planar patches of real-world scenes". *Presence: Teleoperators and Virtual Environments*, 11(2), pp. 158–175 (cited on pp. 81, 82, 100).

Manovich, L. (2001). *The Language of New Media*. Cambridge, MA, USA: The MIT Press (cited on pp. 16, 199).

Manovich, L. (2006). "The poetics of augmented space". *Visual Communication*, 5(2), pp. 219–240 (cited on pp. 16, 21, 27, 29, 67, 68, 73, 79, 198).

Maranan, D. S., Schiphorst, T., Bartram, L., and Hwang, A. (2013). "Expressing Technological Metaphors in Dance Using Structural Illusion from Embodied Motion". In: *Proceedings of the 9th ACM Conference on Creativity & Cognition (C&C '13)*. Sydney, Australia: ACM, pp. 165–174 (cited on pp. 133–135).

Max (2014). Visual programming language. Verison 6.1.9. URL: https://cycling74.com/products/max/ (cited on p. 176).

McGurk, H. and MacDonald, J. (1976). "Hearing lips and seeing voices". *Nature*, 264(5588), pp. 746–748 (cited on p. 62).

Media Technology MSc Programme - Leiden University (n.d.). URL: http://mediatechnology.leiden.edu (cited on pp. 9, 113).

MediaArtTube (2008). URL: https://www.youtube.com/watch?v=6Ek2DW7aV68 (cited on p. 146).

Meijer, P. B. L. (n.d.). *The vOICe - New Frontiers in Sensory Substitution*. URL: https://www.seeingwithsound.com (cited on p. 114).

Microsoft (n.d.). *Microsoft HoloLens*. Online; accessed March 27, 2015. URL: http://www.microsoft.com/microsoft-hololens/ (cited on pp. 3, 40).

Microsoft Research (2014). *Dyadic Projected Spatial Augmented Reality*. YouTube video. Online; accessed January 8, 2016. URL: https://www.youtube.com/watch?v=Df7fZAYVAIE (cited on p. 45).

Milgram, P. and Kishino, F. (1994). "A taxonomy of mixed reality visual displays". *IEICE TRANSACTIONS on Information and Systems*, E77-D(12), pp. 1321–1329 (cited on pp. 19, 21, 24–26, 28–30, 68, 102).

Milgram, P., Takemura, H., Utsumi, A., and Kishino, F. (1994). "Augmented reality: a class of displays on the reality-virtuality continuum". *Proceedings of SPIE: Telemanipulator and Telepresence Technologies*, 2351, pp. 282–292 (cited on pp. 19, 28).

Minamizawa, K., Fukamachi, S., Kajimoto, H., Kawakami, N., and Tachi, S. (2007). "Gravity Grabber: Wearable Haptic Display to Present Virtual Mass Sensation". In: *ACM SIGGRAPH 2007 Emerging Technologies (SIGGRAPH '07)*. San Diego, California: ACM (cited on pp. 84, 94, 106, 187).

Minamizawa, K., Kamuro, S., Fukamachi, S., Kawakami, N., and Tachi, S. (2008). "GhostGlove: Haptic existence of the virtual world". In: *ACM SIGGRAPH 2008 new tech demos*. ACM, p. 18 (cited on p. 94).

Møller, H. (1992). "Fundamentals of binaural technology". *Applied Acoustics*, 36(3–4), pp. 171–218 (cited on pp. 48, 178).

Müller, J., Geier, M., Dicke, C., and Spors, S. (2014). "The boomRoom: mid-air direct interaction with virtual sound sources". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, pp. 247–256 (cited on p. 190).

Murakami, T., Person, T., Fernando, C. L., and Minamizawa, K. (2017). "Altered touch: miniature haptic display with force, thermal and tactile feedback for augmented haptics". In: *ACM SIGGRAPH 2017 Posters*. ACM, p. 53 (cited on p. 94).

Museum of London: Streetmuseum (2014). Mobile application software. Version 2.03. URL: https://itunes.apple.com/de/app/museum-london-streetmuseum/id369684330?l=en&mt=8 (cited on pp. 4, 127).

Nagel, S. K., Carl, C., Kringe, T., Märtin, R., and König, P. (2005). "Beyond sensory substitution?learning the sixth sense". *Journal of Neural Engineering*, 2(4), R13–R26 (cited on pp. 114, 115).

Nakamura, H. and Miyashita, H. (2011). "Augmented gustation using electricity". In: *Proceedings of the 2nd Augmented Human International Conference*. ACM, p. 34 (cited on pp. 60, 108, 115).

Namee, B. M., Beaney, D., and Dong, Q. (2010). "Motion in augmented reality games: an engine for creating plausible physical interactions in augmented reality games". *International Journal of Computer Games Technology*, 2010, p. 4 (cited on pp. 145, 155).

Narumi, T., Kajinami, T., Tanikawa, T., and Hirose, M. (2010a). "Meta cookie". In: *ACM SIGGRAPH 2010 Posters*. ACM, p. 143 (cited on pp. 5, 112).
– (2010b). "Meta cookie". In: *ACM SIGGRAPH 2010 Posters*. ACM, p. 143 (cited on pp. 63, 64).

Narumi, T., Nishizaka, S., Kajinami, T., Tanikawa, T., and Hirose, M. (2011a). "Augmented reality flavors: gustatory display based on edible marker and cross-modal interaction". In: *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, pp. 93–102 (cited on pp. 63, 64).
– (2011b). "Meta cookie+: an illusion-based gustatory display". In: *Virtual and Mixed Reality-New Trends*. Springer, pp. 260–269 (cited on pp. 5, 63, 64, 108, 109, 112).

Narumi, T., Sato, M., Tanikawa, T., and Hirose, M. (2010). "Evaluating cross-sensory perception of superimposing virtual color onto real drink: toward realization of pseudo-gustatory displays". In: *Proceedings of the 1st Augmented Human International Conference*. ACM, p. 18 (cited on pp. 108, 112, 130).

Newsroom, S. (2016). URL: https://news.samsung.com/global/samsung-to-unveil-hum-on-waffle-and-entrim-4d-experimental-c-lab-projects-at-sxsw-2016 (visited on 03/03/2017) (cited on p. 110).

Niantic, Inc. (2017). URL: https://support.pokemongo.nianticlabs.com/hc/en-us/articles/115015868188-Catching-Pok%C3%A9mon-in-AR-mode-iOS-only- (cited on pp. 55, 124).

Noë, A. (2004). *Action in perception*. MIT press (cited on p. 189).

Normand, J.-M., Servières, M., and Moreau, G. (2012). "A New Typology of Augmented Reality Applications". In: *Proceedings of the 3rd Augmented Human International Conference (AH '12)*. Megève, France: ACM, 18:1–18:8 (cited on pp. 29, 115).

Ohshima, T., Satoh, K., Yamamoto, H., and Tamura, H. (1998). "AR2Hockey: a case study of collaborative augmented reality". In: *Proceedings IEEE 1998 Virtual Reality Annual International Symposium*. IEEE, pp. 268–275 (cited on pp. 145, 148, 154, 173).

Oliver, J. (2008-2010). *The Artvertiser*. Improved reality software platform. Online; accessed November 7, 2014. URL: http://theartvertiser.com/ (cited on p. 100).

Olsson, T. and Salo, M. (2011). "Online user survey on current mobile augmented reality applications". In: *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality (ISMAR 2011)*. IEEE, pp. 75–84 (cited on p. 3).

Omosako, H., Kimura, A., Shibata, F., and Tamura, H. (2012). "Shape-COG Illusion: Psychophysical influence on center-of-gravity perception by mixed-reality visual stimulation". In: *IEEE Virtual Reality Short Papers and Posters (VRW 2012)*. IEEE, pp. 65–66 (cited on pp. 62, 106, 112).

Papagiannis, H. (2017). *Augmented Human*. O'Reilly Media (cited on p. 210).

Pepper's ghost (n.d.). *Pepper's ghost—Wikipedia, The Free Encyclopedia*. Online; accessed August 7, 2016.

URL: https://en.wikipedia.org/wiki/Pepper%27s_ghost (cited on p. 26).

Persa, S.-F. (2006). "Sensor fusion in head pose tracking for augmented reality". PhD thesis. Delft University of Technology (cited on p. 44).

Pessoa, S., Moura, G., Lima, J., Teichrieb, V., and Kelner, J. (2010). "Photorealistic rendering for augmented reality: A global illumination and brdf solution". In: *Proceedings IEEE Virtual Reality 2010*. IEEE, pp. 3–10 (cited on pp. 117, 144).

Piekarski, W. and Thomas, B. (2002). "ARQuake: the outdoor augmented reality gaming system". *Communications of the ACM*, 45(1), pp. 36–38 (cited on pp. 4, 21, 24, 26, 28, 83, 91, 123, 148, 154).

– (2004). "Augmented reality working planes: A foundation for action and construction at a distance". In: *Proceedings of the 3rd IEEE/ACM International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, pp. 162–171 (cited on p. 22).

Pokémon GO (2016). Mobile application software. URL: http://www.pokemongo.com (cited on pp. 9, 96).

Porta, G. della (1658). *Natural Magick*. Online; accessed November 7, 2014. URL: http://www.mindserpent.com/American_History/books/Porta/jportat5.html (cited on p. 26).

Poupyrev, I. and Maruyama, S. (2003). "Tactile interfaces for small touch screens". In: *Proceedings of the 16th annual ACM symposium on User interface software and technology*. ACM, pp. 217–220 (cited on p. 107).

Punpongsanon, P., Iwai, D., and Sato, K. (2015). "Softar: Visually manipulating haptic softness perception in spatial augmented reality". *IEEE transactions on visualization and computer graphics*, 21(11), pp. 1279–1288 (cited on p. 105).

Rani, K. and Sharma, R. (2013). "Study of different image fusion algorithm". *International journal of Emerging Technology and advanced Engineering*, 3(5), pp. 288–291 (cited on pp. 21, 41).

Raskar, R., Welch, G., and Chen, W.-C. (1999). "Table-top spatially-augmented realty: bringing physical models to life with projected imagery". In: *Proceedings of the 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99)*. IEEE, pp. 64–71 (cited on pp. 46, 98, 103).

Raskar, R., Welch, G., and Fuchs, H. (1998). "Spatially augmented reality". In: *First IEEE Workshop on Augmented Reality (IWAR'98)*. Citeseer, pp. 11–20 (cited on pp. 45, 92, 99, 102, 104).

Reality Jockey Ltd. (2013). *Sonic Experiences - RjDj*. Online; accessed November 7, 2014. URL: http://rjdj.me/ (cited on p. 33).

Regenbrecht, H. T. and Schubert, T. (2002). "Measuring presence in augmented reality environments: design and a first test of a questionnaire". *Porto, Portugal* (cited on p. 212).

Regenbrecht, H. T. and Wagner, M. T. (2002). "Interaction in a collaborative augmented reality environment". In: *Extended Abstracts on Human Factors in Computing Systems (CHI'02)*. ACM, pp. 504–505 (cited on p. 40).

Reimann, C. and Paelke, V. (2006). "Computer vision based interaction techniques for mobile augmented reality". In: *Proceedings 5th Paderborn Workshop Augmented and Virtual Reality in der Produktentstehung*. Citeseer, pp. 355–362 (cited on p. 48).

Reiners, D., Stricker, D., Klinker, G., and Müller, S. (1998). "Augmented reality for construction tasks: Doorlock assembly". *First IEEE International Workshop on Augmented Reality (IWAR '98)*, 98(1), pp. 31–46 (cited on pp. 20, 24, 28, 37).

Reitman, I. (2004). *Ghostbusters*. Columbia pictures (cited on p. 193).

RjDj (n.d.). *RjDj—Wikipedia, The Free Encyclopedia*. Online; accessed September 21, 2016. URL: https://en.wikipedia.org/w/index.php?title=RjDj&oldid=740512080 (cited on pp. 103, 126).

RjDjme (2008). *RjDj The mind twisting hearing sensation.* YouTube video. Online; accessed February 17, 2017. URL: https://www.youtube.com/watch?v=WPrIPcyemdM (cited on p. 103).

Roads, C. (1988). "Introduction to granular synthesis". *Computer Music Journal*, 12(2), pp. 11–13 (cited on p. 181).

Roberts, G. W., Evans, A., Dodson, A., Denby, B., Cooper, S., Hollands, R., et al. (2002). "The use of augmented reality, GPS and INS for subsurface data visualization". In: *FIG XXII International Congress*. Washington, DC, USA, pp. 1–12 (cited on pp. 20, 24, 28).

Rokeby, D. (1986-1990). *Very Nervous System*. Interactive sound installation. Online; accessed November 7, 2014. URL: http://www.davidrokeby.com/vns.html (cited on p. 193).

Rosenblum, L. (2000). "Virtual and augmented reality 2020". *Computer Graphics and Applications, IEEE*, 20(1), pp. 38–39 (cited on pp. 21, 24, 26).

Ross, C. (2005). "New Media Arts Hybridity: The Vases (Dis)communicants Between Art, Affective Science and AR Technology". *Convergence: The International Journal of Research into New Media Technologies*, 11(4), pp. 32–42 (cited on pp. 29, 115).

Rozier, J. M. (2000). "Hear&there: An augmented reality system of linked audio". PhD thesis. Massachusetts Institute of Technology (cited on p. 52).

Ruhl, A. (2013). "Vestibular Stimulation as an AR Layer?" *AR[t], Augmented Reality, Art and Technology*, 3. Ed. by H. Schraffenberger, M. Kniveton, Y. Kolstee, and J. Verlinden, pp. 26–29. URL: http://arlab.kabk.nl/ar-magazines (cited on p. 109).

Ruhl, A. and Lamers, M. H. (2011). "Experiments with Galvanic Vestibular Stimulation in Daily Activities". In: *Proceedings of the Chi Sparks 2011 Conference* (cited on p. 109).

Russell, S., Norvig, P., and Intelligence, A. (1995). "A modern approach". *Artificial Intelligence. Prentice-Hall, Egnlewood Cliffs*, 25 (cited on p. 124).

Sano, Y., Hirano, Y., Kimura, A., Shibata, F., and Tamura, H. (2013). "Dent-softness illusion in mixed reality space: Further experiments and considerations". In: *Proceedings IEEE Virtual Reality 2013*. IEEE, pp. 153–154 (cited on pp. 62, 105, 112).

Schäfer, T., Fachner, J., and Smukalla, M. (2013). "Changes in the representation of space and time while listening to music". *Frontiers in psychology*, 4, p. 508 (cited on p. 110).

Scherrer, C., Pilet, J., Fua, P., and Lepetit, V. (2008). "The Haunted Book". In: *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR '08)*. Washington, DC, USA: IEEE Computer Society, pp. 163–164 (cited on p. 91).

Schmalstieg, D., Fuhrmann, A., Hesina, G., Szalavári, Z., Encarnaçao, L. M., Gervautz, M., and Purgathofer, W. (2002). "The studierstube augmented reality project". *Presence: Teleoperators and Virtual Environments*, 11(1), pp. 33–54 (cited on p. 154).

Schraffenberger, H. (2012). "Chasing virtual spooks, losing real weight". *AR[t], Augmented Reality, Art and Technology*, 2. Ed. by Y. Kolstee, H. Schraffenberger, E. Vahrmeijer, and J. Verlinden, pp. 48–51. URL: http://arlab.kabk.nl/ar-magazines (cited on p. 17).

– (2013). "Subject: Interview". *AR[t], Augmented Reality, Art and Technology*, 3. Ed. by H. Schraffenberger,

M. Kniveton, Y. Kolstee, and J. Verlinden, pp. 18–23. URL: http://arlab.kabk.nl/ar-magazines (cited on p. 17).
Schraffenberger, H. (2014). "Hitting imaginary walls, pulling virtual strings". *AR[t], Augmented Reality, Art and Technology*, 5. Ed. by H. Schraffenberger, M. Kniveton, L. Kolstee, Y. Kolstee, J. Verlinden, and R. Wesdorp, pp. 66–71. URL: http://arlab.kabk.nl/ar-magazines (cited on p. 17).

Schraffenberger, H. and van der Heide, E. (2013a). "From Coexistence to Interaction: Influences Between the Virtual and the Real in Augmented Reality". In: *Proceedings of the 19th International Symposium on Electronic Art (ISEA2013)*. Ed. by K. Cleland, L. Fisher, and R. Harley. Sydney, pp. 1–3 (cited on pp. 16, 18).
– (2013b). "Towards Novel Relationships between the Virtual and the Real in Augmented Reality". In: *Arts and Technology*. Ed. by G. De Michelis, F. Tisato, A. Bene, and D. Bernini. LNICST 116. Springer, pp. 73–80 (cited on p. 16).
– (2014). "The Real in Augmented Reality". In: *Proceedings of the Second Conference on Computation, Communication, Aesthetics and X (xCoAx 2014)*. Ed. by M. Carvalhais and M. Verdicchio, pp. 64–74 (cited on p. 16).
– (2015). "Sonically Tangible Objects". In: *Proceedings of the Third Conference on Computation, Communication, Aesthetics and X (xCoAx 2015)*. Ed. by A. Clifford, M. Carvalhais, and M. Verdicchio, pp. 233–248 (cited on p. 17).
– (2016). "Multimodal Augmented Reality: The Norm Rather Than the Exception". In: *Proceedings of the 2016 Workshop on Multimodal Virtual and Augmented Reality (MVAR '16)*. ACM, pp. 1–6 (cited on p. 17).
– (2018). "Reconsidering Registration: New Perspectives on Augmented Reality". In: *Interactivity, Game Creation, Design, Learning, and Innovation. ArtsIT 2017, DLI 2017.* Ed. by A. L. Brooks, E. Brooks, and N. Vidakis. LNICST 229. Springer, pp. 172–183 (cited on p. 17).
– (2014b). "Everything Augmented: On the Real in Augmented Reality". *Journal of Science and Technology of the Arts*, 6(1), pp. 17–29 (cited on p. 17).

Schubert, T., Friedmann, F., and Regenbrecht, H. (2001). "The experience of presence: Factor analytic insights". *Presence: Teleoperators & Virtual Environments*, 10(3), pp. 266–281 (cited on p. 212).

Sekiguchi, Y., Hirota, K., and Hirose, M. (2005). "The Design and Implementation of Ubiquitous Haptic Device". In: *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC '05)*. Washington, DC, USA: IEEE Computer Society, pp. 527–528 (cited on pp. 94, 106, 187).

Seyama, J. and Nagayama, R. S. (2007). "The uncanny valley: Effect of realism on the impression of artificial human faces". *Presence: Teleoperators and virtual environments*, 16(4), pp. 337–351 (cited on p. 157).

Shazam (2008). Mobile application software. URL: http://www.shazam.com/ (cited on p. 87).

Sheffield, E. (2004). "Beyond abstraction: Philosophy as a practical qualitative research method". *The qualitative report*, 9(4), pp. 760–769 (cited on p. 10).

Sheridan, T. B. (1992). "Musings on telepresence and virtual presence". *Presence: Teleoperators & Virtual Environments*, 1(1), pp. 120–126 (cited on p. 212).

Siemens SX1 (n.d.). *Siemens SX1—Wikipedia, The Free Encyclopedia*. Online; accessed August 25, 2016. URL: https://en.wikipedia.org/w/index.php?title=Siemens_SX1&oldid=736142402 (cited on p. 48).

Silva, E. S., de Abreu, J., de Almeida, J., Teichrieb, V., and Ramalho, G. L. (2013). "A preliminary evaluation of the leap motion sensor as controller of new digital musical instruments". *Recife, Brasil* (cited on p. 185).

Silva, R., Oliveira, J. C., and Giraldi, G. A. (2003). "Introduction to augmented reality". *National Laboratory for Scientific Computation, Av. Getulio Vargas* (cited on pp. 29, 188).

Six degrees of freedom (n.d.). *Six degrees of freedom—Wikipedia, The Free Encyclopedia*. Online; accessed December 6, 2016. URL: https://en.wikipedia.org/w/index.php?title=Six_degrees_of_freedom&oldid=753300785 (cited on p. 43).

Sobecka, K. (n.d.). *All the Universe is Full of the Lives of Perfect Creatures | Karolina Sobecka*. URL: http://www.gravitytrap.com/artwork/perfect-creatures (cited on p. 151).

Spence, C. and Youssef, J. (2015). "Olfactory dining: designing for the dominant sense". *Flavour*, 4(1), p. 32 (cited on pp. 24, 30, 68).

Sphero (2011). Robot toy. URL: http://www.gosphero.com/ (cited on pp. 4, 40, 56, 60, 91, 147, 159).

Staddon, J. E. R. (1975). "A note on the evolutionary significance of" supernormal" stimuli". *The American Naturalist*, 109(969), pp. 541–545 (cited on p. 171).

Steuer, J. (1992). "Defining Virtual Reality: Dimensions Determining Telepresence". *Journal of Communication*, 42(4), pp. 73–93 (cited on pp. 68, 212).

Street Museum NL (2013). Mobile application software. Version 2.0. URL: https://itunes.apple.com/de/app/museum-london-streetmuseum/id369684330?l=en&mt=8 (cited on pp. 85, 127).

Sugano, N., Kato, H., and Tachibana, K. (2003). "The effects of shadow representation of virtual objects in augmented reality". In: *Proceedings of the Second IEEE and ACM International Symposium on Mixed and Augmented Reality*. IEEE, pp. 76–83 (cited on pp. 55, 119).

Sutherland, I. E. (1965). "The Ultimate Display". In: *Proceedings of the IFIP Congress*. Online; accessed November 7, 2014, pp. 506–508. URL: http://www8.informatik.umu.se/~jwworth/TheUltimateDisplay.pdf (cited on pp. 26, 137–139, 161, 213).

– (1968). "A head-mounted three dimensional display". In: *Proceedings of the AFIPS '68 Fall Joint Computer Conference, December 9-11, 1968, San Francisco, California, USA - Part I*. vol. 33. AFIPS Conference Proceedings. AFIPS, pp. 757–764 (cited on p. 3).

Takeuchi, Y. (2010). "Gilded gait: reshaping the urban experience with augmented footsteps". In: *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. ACM, pp. 185–188 (cited on p. 106).

The app formerly known as H _ _ r (2016). Mobile application software. Version 0.3.2. URL: http://hearapp.io (cited on pp. 33, 103, 104).

The Invisible Drums of Demian Kappenstein and Marc Bangert (2011). Blog post. Online; accessed November 7, 2014. URL: http://www.moderndrummer.com/site/2011/10/the-invisible-drums-of-demian-kappenstein-and-marc-bangert (cited on p. 193).

Thomas, B. (2009). "Augmented Reality Visualisation Facilitating The Architectural Process". In: *Mixed Reality In Architecture, Design And Construction*. Ed. by X. Wang and M. Schnabel. Springer Netherlands, pp. 105–118 (cited on pp. 22, 24).

Thomas, B., Close, B., Donoghue, J., Squires, J., de Bondi, P., Morris, M., and Piekarski, W. (2000). "ARQuake: An Outdoor/Indoor Augmented Reality First Person Application". In: *Proceedings of the 4th*

IEEE International Symposium on Wearable Computers (ISWC 2000). Washington, DC, USA: IEEE Computer Society, pp. 139–146 (cited on pp. 3, 4, 91).

Unstable Media, V. for the (n.d.). URL: http://v2.nl/archive/works/sonic-interface (cited on p. 104).

Valbuena, P. (2008). *N 520437 E 041900 [the hague city hall*. Site-specific large scale intervention. Online; accessed November 7, 2014. URL: http://www.pablovalbuena.com/selectedwork/n-520437-e-041900/ (cited on pp. 5, 83, 102, 103).

Valentini, P. P. and Pezzuti, E. (2010). "Interactive multibody simulation in augmented reality". *Journal of Theoretical and Applied Mechanics*, 48(3), pp. 733–750 (cited on pp. 145, 153, 156).

Vallino, J. R. (1998). "Interactive augmented reality". PhD thesis. Rochester, NY, USA: University of Rochester (cited on pp. 26, 40, 188).

Vallino, J. R. and Brown, C. (1999). "Haptics in augmented reality". In: *Proceedings IEEE International Conference on Multimedia Computing and Systems*. Vol. 1. IEEE, pp. 195–200 (cited on pp. 41, 97, 149).

van der Heide, E. (2012). "Radioscape—in the context of augmented reality". *AR[t], Augmented Reality, Art and Technology*, 2. Ed. by Y. Kolstee, H. Schraffenberger, E. Vahrmeijer, and J. Verlinden, pp. 18–23. URL: http://arlab.kabk.nl/ar-magazines (cited on p. 120).
– (2000-). *Radioscape*. Immersive electromagnetic environment. Online; accessed November 7, 2014. URL: http://www.evdh.net/radioscape/ (cited on pp. 33, 120, 147).

van Duijn, M. J. (2016). "The lazy mindreader: a humanities perspective on mindreading and multiple-order intentionality". PhD thesis (cited on p. 8).

Van Hoof, J., Mazej, M., and Hensen, J. L. (2010). "Thermal comfort: research and practice". *Frontiers in Bioscience*, 15(2), pp. 765–788 (cited on p. 110).

Van Krevelen, D. and Poelman, R. (2010). "A survey of augmented reality technologies, applications and limitations". *International Journal of Virtual Reality*, 9(2), p. 1 (cited on p. 29).

van Krevelen, D. and Poelman, R. (2010). "A survey of augmented reality technologies, applications and limitations". *International Journal of Virtual Reality*, 9(2), pp. 1–20 (cited on pp. 3, 25).

van Velthoven, L. (2011). *Room Racers*. Interactive installation. Online; accessed November 7, 2014. URL: www.lievenvv.com/Projects/RoomRacers.htm (cited on p. 118).

Veenhof, S. (2010). *MoMA NYC augmented reality exhibition*. https://www.youtube.com/watch?v=b9T2LVM7ynM (cited on p. 92).
– (2016). *Augmented Reality art exhibition MoMA NYC (guerrilla intervention)*. URL: http://www.sndrv.nl/moma/ (cited on pp. 60, 92).

Verlinden, J. C., de Smit, A., Peeters, A. W. J., and van Gelderen, M. H. (2003). "Development of a flexible augmented prototyping system". *Journal of WSCG*, 11(3), pp. 496–503 (cited on pp. 98, 103).

Walter, W. G. (1950). "An imitation of life". *Scientific American*, 182(5), pp. 42–45 (cited on p. 165).

Ward, J. and Meijer, P. (2010). "Visual experiences in the blind induced by an auditory sensory substitution device". *Consciousness and cognition*, 19(1), pp. 492–500 (cited on p. 114).

Watson, T. (2005). *Audio Space*. Interactive installation. Online; accessed November 7, 2014. URL: http://www.theowatson.com/site_docs/work.php?id=15 (cited on pp. 33, 102).

Wei, J., Zhao, S., Nakatsu, R., and Duh, H. B. (2012). "When AR meets food: A structural overview of the research space on multi-facets of food". In: *Proceedings of the IEEE International Symposium on Mixed and*

Augmented Reality (ISMAR-AMH 2012). IEEE, pp. 97–98 (cited on p. 109).

Winzen, J., Albers, F., and Marggraf-Micheel, C. (2014). "The influence of coloured light in the aircraft cabin on passenger thermal comfort". *Lighting Research & Technology*, 46(4), pp. 465–475 (cited on p. 110).

Witmer, B. G. and Singer, M. J. (1998). "Measuring presence in virtual environments: A presence questionnaire". *Presence*, 7(3), pp. 225–240 (cited on p. 212).

Yamada, T., Yokoyama, S., Tanikawa, T., Hirota, K., and Hirose, M. (2006). "Wearable olfactory display: Using odor in outdoor environment". In: *Proceedings IEEE Virtual Reality 2006*. IEEE, pp. 199–206 (cited on p. 61).

Yanagida, Y., Kawato, S., Noma, H., Tomono, A., and Tesutani, N. (2004). "Projection based olfactory display with nose tracking". In: *Proceedings IEEE Virtual Reality 2004*. IEEE, pp. 43–50 (cited on p. 96).

You, S. and Neumann, U. (2001). "Fusion of vision and gyro tracking for robust augmented reality registration". In: *Proceedings IEEE Virtual Reality 2001*. IEEE, pp. 71–78 (cited on p. 41).

Yuen, S., Yaoyuneyong, G., and Johnson, E. (2011). "Augmented reality: An overview and five directions for AR in education". *Journal of Educational Technology Development and Exchange*, 4(1), pp. 119–140 (cited on p. 79).

Zhou, F., Duh, H. B.-L., and Billinghurst, M. (2008). "Trends in Augmented Reality Tracking, Interaction and Display: A Review of Ten Years of ISMAR". in: *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR '08)*. Washington, DC, USA: IEEE Computer Society, pp. 193–202 (cited on pp. 7, 20, 22, 24, 37, 59, 74).

Zokai, S., Esteve, J., Genc, Y., and Navab, N. (2003). "Multiview paraperspective projection model for diminished reality". In: *Proceedings of the Second IEEE and ACM International Symposium on Mixed and Augmented Reality*. IEEE, pp. 217–226 (cited on p. 100).

Summary

In recent years, virtual content has become part of our everyday environment in a profoundly new way. Virtual objects no longer solely appear on the screen of computers, tablets, mobile phones or advertisement screens. Rather they have started to appear right here, in our everyday environment: With the right mobile application, we can view virtual creatures flying through our surroundings or see site-specific information, such as where to find nearby restaurants, floating right in front of us. Wearing a head-mounted display, we can invite virtual characters into our house or turn our living room into the venue of a partially real and partially virtual adventure.

The phenomenon of virtual content being part of and appearing in the real world has a name: augmented reality (AR). By now, a wide variety of so-called AR applications exists. In many respects, these AR applications could not be more diverse. They make use of a broad variety of different technologies, such as headsets, projectors, headphones and tactile displays. In line with this, they present different kinds of sensory stimuli, like visuals, sounds and scents and provide various types of virtual content, among which 3D models, textual information, photographs and sound recordings. They alter our experience of the real world in various ways; for instance, by seemingly removing physical objects from our view or by integrating additional elements into it. They are designed for many different contexts, such as work, entertainment and education. Accordingly, they serve a variety of purposes. For instance, some AR apps are here to inform us about our surroundings while others exist to keep us entertained.

In their totality, existing applications provide a rather blurry picture of AR and therefore raise the questions: *"What is augmented reality?"* and *"What forms can AR take?"*. In this thesis we address these questions and explore the fundamental characteristics and potential manifestations of AR.

In chapter 1, we take an initial glance at examples of AR. We illustrate the diversity of the AR landscape and consequently raise the question "What is augmented reality?".

In chapter 2, we investigate how existing research answers this question. We review existing definitions and descriptions of AR and identify three common ideas about augmented reality: First of all, AR is

generally considered a technology. Second, AR is widely understood in terms of visual virtual overlays that are presented on top of a participant's view of the real world. Third, AR is considered to spatially align virtual content with the real world in three dimensions (this gives virtual objects a position in the real world).

At the same time, our review of existing AR literature also reveals many divergent and *broader* understandings of AR. For instance, we encounter research that also considers non-visual virtual content (such as sound) in the context of AR and researchers that explicitly argue against seeing AR as a technology. In their totality, the various reviewed positions suggest that AR *can* involve technologies that overlay virtual images onto a participant's view and that align these images with the real world. At the same time, we get a strong sense that there is more to AR than such technologies. Thus, the review leaves us wondering, what, if not just a technology, AR is or can be.

In chapter 3, we respond to existing definitions, challenge commonly accepted views and argue for new (or at least different) perspectives on AR. First, we depart from the understanding of AR as a technology. Instead, we claim that AR technology enables augmented reality. We focus on the resulting augmented reality environments and experiences rather than on the technologies that enable them. Second, we approach AR as a multimodal and interactive environment and argue that AR engages all our senses. Rather than focusing on what a user or participant sees, we focus on non-visual, multimodal and interactive aspects of both the real world and virtual content. Third, we see AR as a result of the relationships between the virtual and the real. Whereas AR is generally assumed to involve the spatial alignment of virtual content with the real world in 3D, we suggest that other types of relationships between the virtual and the real are possible, potentially leading to other and new forms of AR. These three ideas are synthesized and culminate in our definition of AR as an interactive and multimodal environment where a participant experiences a *relationship* between virtual content and the real environment.

In chapter 4, we explore and illustrate the different ways in which the virtual and the real can relate to (and thus augment) one another. With this, we address the second key question of this research: "What forms can AR take?". On a fundamental level, we distinguish AR from scenarios where participants do not experience any link between the virtual and the real. We then identify *spatial relationships* between the virtual and the real (here, virtual content seemingly exists in real space) and *content-based relationships* between the virtual and the real (here, the virtual relates to the real on the content-level) as the two core relationships that facilitate AR. Subsequently, we question how virtual content can affect its real surroundings. Based on the role that the virtual content plays in the real space, we distinguish five forms of AR:

Extended reality: scenarios where the virtual supplements the real environment.

2. *Diminished reality:* cases where virtual content seemingly removes real elements from the real environment.

3. *Altered reality:* environments where the virtual information changes the qualities of the real world.

4. *Hybrid reality:* scenarios where the virtual completes a physical environment that would be considered incomplete without the virtual additions.

5. *Extended perception:* cases where unperceivable but real aspects of the real world are translated into virtual information that we can perceive with our senses.

We then focus on scenarios where virtual objects seemingly exist in and extend the real world. We notice that the presence of virtual objects in real space opens up possibilities for influences and interaction between the virtual and the real. On this level, we distinguish among two main forms of relationships between the virtual and the real: *(1) physical relationships* (the virtual and the real seemingly affect each other physically) and *(2) behavioral relationships* (the virtual and the real sense each other and react to one another on a social or behavioral level).

Subsequently, we briefly discuss other possible relationships, such as temporal relationships between the virtual and the real and musical relationships between virtual and real instruments.

Chapter 5 focuses on one particular relationship between the virtual and the real, namely interaction between the two. Based on the fact that virtual objects do not have to adhere to physical laws and cannot directly apply forces to real objects, we ask the following questions: What types of interaction between the virtual and the real are both possible and believable? We explore (1) whether virtual objects can interact with physical objects in a realistic manner as well as (2) whether they can interact in imaginative but believable ways. In order to answer these questions, we follow both a theoretical and a practical approach. We review existing research and AR works, conduct our own initial series of practical experiments as well as reflect upon these experiments. This leads us to the conclusion that virtual and real objects can believably simulate real-world influences as well as influence each other in imaginative ways that have no equivalent in the physical world.

Chapter 6 builds on the idea that virtual objects can differ from real objects. We explore whether virtual objects can also be perceived differently from how we perceive real objects. In order to study and illustrate such possibilities, we develop and present a prototype of what we call sonically tangible objects. More concretely, we present a virtual, invisible and non-tactile cube that is placed in a real, physical

space. This cube can be experienced through exploratory hand gestures and gives sonic feedback. Touching the cube with one's fingers triggers binaural sounds that appear to originate from the exact spot where the object is touched. Our initial experimentation suggests that this sound- and movement-based approach can result in tactile-like experiences and convey the presence of virtual objects in real space. We discuss the concept behind, implementation of and our experience with the sonically tangible cube and place our research in a broader context.

Chapter 7 concludes the thesis. We revisit our main questions (*"What is augmented reality?"* and *"What forms can AR take?"*) and reflect on the answers we have arrived at. Furthermore, we address pending questions that have surfaced during this trajectory (e.g., *"What is augmented in AR?"*) and that we can answer now, after having obtained a thorough understanding of existing research, after having gained hands-on experience with AR and after having developed our own comprehensive theory of AR. In addition, we summarize insights that can guide the design of AR experiences. Moreover, we discuss methodological and technological limitations of our study and present possible directions for future AR research and development. Among other things, we propose to focus less on mimicking our existing reality, and instead, to create new, imaginative and creative forms of AR that have no counterpart in a purely physical world.

Samenvatting

De afgelopen jaren is virtuele content op een geheel nieuwe manier onderdeel geworden van onze alledaagse omgeving. Virtuele objecten verschijnen nu ook hier, in onze fysieke, dagelijkse omgeving in plaats van alleen op schermen van computers, tablets, mobiele telefoons of op digitale reclameborden. Met een mobile app kunnen we bijvoorbeeld virtuele wezens door de wereld zien bewegen of geïnformeerd worden over objecten in de omgeving. Op soortgelijke wijze kunnen we virtuele personages in huis uitnodigen en in de woonkamer deels echte en deels virtuele avonturen beleven door een zogenaamd head-mounted display op te zetten.

Dit fenomeen van virtuele content die in de echte wereld verschijnt en deel wordt van onze omgeving heeft een naam: augmented reality (AR). Inmiddels bestaat er een groot aantal zogenaamde AR-applicaties. Deze AR-applicaties kunnen veel verschillende eigenschappen hebben. Er wordt gebruik gemaakt van een breed scala aan technologieën, zoals headsets, projectoren, hoofdtelefoons en tactiele schermen. Zodoende bieden AR-applicaties verschillende zintuiglijke stimuli aan, zoals visuele informatie, geluiden en geuren, en geven daarbij verschillende soorten virtuele inhoud weer, bijvoorbeeld 3D-modellen, tekstuele informatie, reclame, foto's en geluidsopnamen. Daarnaast veranderen dergelijke applicaties onze ervaring van de echte wereld op verschillende manieren, bijvoorbeeld door fysieke voorwerpen uit onze blik te verwijderen of juist toe te voegen. AR-applicaties zijn bovendien ontworpen voor veel verschillende toepassingen, zoals werk, entertainment en educatie en dienen daarmee verschillende doelen. Sommige AR-applicaties informeren ons over onze omgeving, terwijl anderen juist het doel hebben om ons te vermaken.

Dit scala aan AR-applicaties geeft in zichzelf geen duidelijk beeld van AR. Dit roept de volgende vragen op: *"Wat is augmented reality?"* en *"Welke vormen kan AR aannemen?"*. Het doel van dit proefschrift is om deze vragen te behandelen en de fundamentele kenmerken van AR te verkennen.

In hoofdstuk 1 werpen we een eerste blik op voorbeelden van AR. We illustreren de diversiteit van het AR-landschap en de vraag *"Wat is augmented reality?"* dient zich aan.

In hoofdstuk 2 nemen we deze vraag aan de hand door te onderzoeken hoe bestaand werk deze vraag beantwoordt. We bekijken bestaande definities en beschrijvingen van AR. Daarmee identificeren we drie gangbare ideeën over augmented reality. Ten eerste wordt AR in het algemeen beschouwd als een technologie. Ten tweede wordt AR vaak gezien als een virtuele visuele laag die over ons perceptie van de echte wereld gelegd wordt. Ten derde wordt AR geacht om virtuele content in drie dimensies in de echte wereld correct te positioneren zo dat het erin opgaat.

Naast deze gangbare opvattingen leert het literatuuronderzoek ons ook andere en *bredere* opvattingen van AR. Naast visuele toevoegingen wordt ook niet-visuele content in de AR-context beschouwd. Er zijn onderzoekers die expliciet beweren dat AR geen technologie is. Tezamen kunnen we zeggen dat de verschillende opvattingen suggereren dat AR in ieder geval iets te maken *kan* hebben met technologieën die virtuele inhoud mengen met de perceptie van de gebruiker. Daarnaast krijgen we de indruk dat AR meer is dan alleen technologie. Daarom vragen we ons af wat AR, behalve technologie, is of kan zijn.

In hoofdstuk 3 nemen we de bestaande definities op de hak. Vanuit geaccepteerde opvattingen over AR pleiten we voor nieuwe (of op zijn minst andere) en verfrissende perspectieven op AR. Om te beginnen nemen we afstand van de opvatting dat AR een technologie is. In plaats daarvan beweren we dat AR-technologie slechts iets is wat augmented reality mogelijk maakt. Daarom richten we ons op de AR-omgevingen en ervaringen die daarmee mogelijk worden in plaats van alleen op de technologieën. Ten tweede zien we AR als een multimodale en interactieve omgeving die al onze zintuigen prikkelt. In plaats van te focussen op de visuele waarneming van een gebruiker in AR, richten we ons op niet-visuele, multimodale en interactieve aspecten van zowel de echte wereld als van de virtuele content. Ten derde beschouwen we AR als een resultaat van de relaties tussen het virtuele en het reële. Het wordt algemeen aangenomen dat AR het correct uitlijnen van virtuele content en de echte drie dimensionale wereld omvat. Echter laten wij zien dat ook andere soorten relaties tussen het virtuele en het reële mogelijk zijn en dat deze potentieel leiden tot andere en nieuwe vormen van AR. Met het samenvoegen van deze ideeën ontstaat onze definitie: AR is een interactieve en multimodale omgeving waarin een deelnemer een *relatie* ervaart tussen virtuele content en de werkelijke omgeving.

Het virtuele en het reële kunnen zich op verschillende manieren tot elkaar verhouden. In hoofdstuk 4 verkennen en illustreren we dit. Hierdoor komen we aan bij de tweede hoofdvraag van dit onderzoek: *"Welke vormen kan AR aannemen?"*. Op een fundamenteel niveau houdt AR in dat deelnemers een relatie tussen het virtuele en het echte ervaren. We identificeren twee kernrelaties die AR mogelijk maken: *(a) ruimtelijke relaties* tussen het virtuele en het echte (hier lijkt virtuele in-

houd in de echte ruimte te bestaan) en *(b) inhoudelijke relaties* tussen het virtuele en het echte. Vervolgens vragen we ons af hoe virtuele content de werkelijke omgeving kan beïnvloeden. Op basis van de rol die virtuele content in de werkelijke omgeving speelt, onderscheiden we vijf vormen van AR:

1. *Extended reality:* scenario's waarin virtuele content de echte omgeving aanvult.
2. *Diminished reality:* gevallen waarbij virtuele content elementen uit de echte wereld lijkt te verwijderen.
3. *Altered reality:* omgevingen waarin virtuele informatie de eigenschappen van de echte wereld ogenschijnlijk verandert.
4. *Hybrid reality:* scenario's waarin de fysieke omgeving alleen dan volledig is als de virtuele content toegevoegd wordt.
5. *Extended perception:* gevallen waarin niet-waarneembare, maar daadwerkelijk echte aspecten van onze wereld worden vertaald in virtuele informatie die wel met onze zintuigen kan worden waargenomen.

Daaropvolgend gaan we opnieuw in op scenario's waarin virtuele objecten lijken te bestaan in onze wereld. We merken op dat de aanwezigheid van virtuele objecten in de fysieke ruimte mogelijkheden biedt voor invloeden van en interacties tussen de virtuele content en de werkelijkheid. Op het niveau van interactie onderscheiden we twee belangrijke vormen van relaties tussen het virtuele en het echte: *(1) fysieke relaties* (virtuele objecten en de werkelijkheid lijken elkaar fysiek te beïnvloeden) en *(2) gedragsrelaties* (het virtuele object en de werkelijkheid nemen elkaar waar en reageren op elkaar qua gedrag). Bovendien bespreken we kort andere mogelijke relaties, zoals temporele relaties tussen virtuele content en de werkelijkheid en interactie tussen virtuele en echte muziekinstrumenten.

Hoofdstuk 5 richt zich op één bepaalde relatie tussen virtuele content en de werkelijkheid, te weten de interactie tussen virtuele en echte objecten in de echte ruimte. Omdat virtuele objecten zich niet hoeven te houden aan fysieke wetten en ook niet direct fysieke krachten kunnen uitoefenen op echte objecten, stellen we de volgende vraag: Welke soorten interactie tussen virtuele en het echte objecten zijn zowel mogelijk als geloofwaardig? We splitsen de vraag op en verkennen (1) of virtuele objecten en fysieke objecten op realistische wijze op elkaar kunnen inwerken, en (2) of ze op fantasierijke maar toch geloofwaardige manieren kunnen interacteren. Om deze twee deelvragen te beantwoorden combineren we een theoretische met een empirische aanpak. Uitgaand van bestaand onderzoek doen we een eerste reeks experimenten en reflecteren op de uitkomsten. Onze conclusie is tweeledig. Ten eerste: virtuele en echte objecten kunnen acties en reacties

ons bekend uit de echte wereld op een geloofwaardige manier simuleren. Ten tweede: ze kunnen elkaar ook beïnvloeden op fantasierijke manieren die echter geen equivalent hebben in de fysieke wereld.

Hoofdstuk 6 borduurt voort op het idee dat virtuele objecten van echte objecten kunnen verschillen. We onderzoeken of virtuele objecten ook anders kunnen worden waargenomen dan echte objecten. Om dit te bestuderen en de mogelijkheden te verkennen, ontwikkelen en gebruiken we een prototype dat we het "sonically tangible object" noemen. Concreet betekent dit dat we een virtuele, onzichtbare en niet-tactiele kubus in een fysieke ruimte plaatsen. Met verkennende handbewegingen kan deze kubus worden ervaren. Het aanraken van de kubus met de vingers veroorzaakt zogenaamde "binaurale geluiden" die afkomstig lijken van de exacte plek waar de aanraking heeft plaatsgehad. Onze experimenten laten zien dat met deze opzet tactielachtige ervaringen kunnen worden gecreëert en dat de aanwezigheid van virtuele objecten in de echte ruimte kan worden overgebracht. We bespreken het onderliggende concept, de implementatie van en onze ervaring met het "sonically tangible object" en plaatsen ons onderzoek in een bredere context.

In hoofdstuk 7 sluiten we het proefschrift af met een terugkeer naar onze belangrijkste vragen. Te weten, *"Wat is Augmented Reality?"* en *"Welke vormen kan AR aannemen?"*. We reflecteren op de antwoorden die we hebben gevonden en behandelen openstaande vragen die tijdens dit traject naar voren zijn gekomen. Met de grondige studie van bestaand onderzoek, de praktische ervaring die we hebben opgedaan en met onze eigen uitgebreide theorie over AR kunnen we deze nu beantwoorden. We vatten de inzichten samen waarmee het ontwerp van AR-ervaringen kan worden ondersteund. Het onderzoek uit dit proefschrift heeft een exploratief karakter. Daarom is het zinnig de methodologische en technologische beperkingen van dit onderzoek ook even te behandelen en mogelijke richtingen voor toekomstig AR-onderzoek en ontwikkeling in het veld te identificeren. Daarbij stellen we voor om minder aandacht te besteden aan het nadoen van onze bestaande realiteit, en in plaats daarvan nieuwe, fantasierijke en creatieve vormen van AR te creëren die juist geen equivalent hebben in onze fysieke wereld.

Zusammenfassung

In den letzten Jahren sind virtuelle Inhalte auf völlig neue Art Teil unseres Alltags geworden. Virtuelle Objekte erscheinen nicht mehr nur auf dem Bildschirm von Computern, Tablets, Mobiltelefonen oder auf digitalen Werbedisplays. Sie sind vielmehr auch hier, in unserer sogenannten "echten Welt" anwesend: Mit einer geeigneten mobilen App können wir virtuelle Wesen bewundern, die durch unsere Umgebung fliegen, oder ortsspezifische Informationen, zum Beispiel über nahegelegene Restaurants, direkt vor uns schweben sehen. Ausgestattet mit einem Head-Mounted-Display können wir virtuelle Charaktere in unser Haus einladen oder unser Wohnzimmer in den Schauplatz eines teils realen und teils virtuellen Abenteuers verwandeln.

Das Phänomen, dass virtuelle Inhalte in der realen Welt erscheinen und Teil unserer echten Umgebung werden, hat einen Namen: Augmented Reality (AR); auf Deutsch auch "Erweiterte Realität" genannt. Mittlerweile existiert eine Vielzahl von sogenannten AR-Applikationen. In vielerlei Hinsicht könnten diese AR-Applikationen nicht unterschiedlicher sein. Sie nutzen eine breite Palette verschiedener Technologien wie Headsets, Projektoren, Kopfhörer und taktile Displays. Sie bieten verschiedene Sinnesreize, wie Bilder, Klänge und Düfte und stellen verschiedene virtuelle Inhalte bereit, wie zum Beispiel 3D-Modelle, Textinformationen, Fotografien und Tonaufnahmen. Sie verändern unsere Erfahrung der realen Welt auf verschiedene Arten, indem sie zum Beispiel scheinbar Objekte aus unserer Umgebung entfernen oder zusätzliche Elemente darin integrieren. Sie sind für viele verschiedene Kontexte, wie Arbeit, Unterhaltung und Bildung konzipiert. Dementsprechend dienen sie einer Vielzahl von Zwecken. Zum Beispiel wollen uns einige AR-Applikationen über unsere Umgebung informieren, während andere dazu da sind, uns zu unterhalten.

In ihrer Gesamtheit vermitteln bestehende Anwendungen nur ein undeutliches Bild von AR und werfen somit die Frage auf: *"Was ist Augmented Reality"* und *"Welche Formen kann AR annehmen?"*. In dieser Arbeit gehen wir diesen Fragen nach und erforschen die grundlegenden Eigenschaften und möglichen Erscheinungsformen von AR.

In Kapitel 1 werfen wir anhand von Beispielen einen ersten Blick auf AR. Wir illustrieren die Vielfalt der AR-Landschaft und stellen die

Frage "Was ist Augmented Reality?".

In Kapitel 2 untersuchen wir, wie diese Frage in anderen Forschungsarbeiten beantwortet wird. Wir betrachten existierende Definitionen und Beschreibungen von AR und identifizieren drei gängige Auffassungen von Augmented Reality: Erstens wird AR generell als eine Technologie betrachtet. Zweitens wird AR weithin im Sinne von visuellen virtuellen Schichten (Overlays) verstanden, die über unsere Sicht der realen Welt gelegt werden. Drittens wird angenommen, dass AR virtuelle Objekte räumlich und dreidimensional in der echten Welt positioniert (virtuelle Elemente und die echte Umgebung werden miteinander registriert).

Zugleich bringt unsere umfassende Literaturrecherche auch einige andere und *breitere* Auffassungen von AR zutage. Zum Beispiel stoßen wir auf Beschreibungen von AR, die auch nicht-visuelle virtuelle Inhalte (wie Klang) berücksichtigen, und begegnen wir Forschern, die sich explizit dagegen aussprechen, AR als eine Technologie zu betrachten. In ihrer Gesamtheit legen die verschiedenen Positionen nahe, dass AR Technologien involvieren *kann*, die virtuelle Bilder über unsere Sicht auf die Welt legen und diese Bilder mit der realen Welt registrieren. Gleichzeitig bekommen wir den starken Eindruck, dass sich mehr hinter AR verbirgt als solche Technologien. Wir fragen uns daher, was, wenn nicht nur eine Technologie, AR ist oder sein kann.

In Kapitel 3 beziehen wir Stellung zu bestehenden Definitionen, stellen wir allgemein akzeptierte Ansichten infrage und plädieren wir für neue (oder zumindest andere) Sichtweisen auf AR. Zu allererst lassen wir das Verständnis von AR als Technologie hinter uns. Stattdessen schlagen wir vor, dass AR-Technologie Augmented Reality lediglich ermöglicht. Demgemäß konzentrieren wir uns auf die resultierenden AR-Umgebungen und -Erfahrungen statt auf die Technologien, die diese ermöglichen. Zweitens fassen wir AR als eine multimodale und interaktive Umgebung auf und vertreten den Standpunkt, dass AR alle Sinne anspricht. Anstatt uns auf das zu konzentrieren, was ein Benutzer oder Teilnehmer in AR sieht, konzentrieren wir uns auf nicht-visuelle, multimodale und interaktive Aspekte der realen Welt wie auch der virtuellen Inhalte. Drittens sehen wir AR als Ergebnis der Beziehungen zwischen dem Virtuellen und dem Realen. Während allgemein angenommen wird, dass AR die räumliche Registrierung von virtuellem Inhalt und der realen Welt in 3D voraussetzt, legen wir nahe, dass andere Arten von Beziehungen zwischen dem Virtuellen und dem Realen vorstellbar sind, die möglicherweise zu anderen und neuen Formen von AR führen. Diese drei Ideen werden zusammengeführt und münden in unsere Definition von AR als einer interaktiven und multimodalen Umgebung, in der ein Teilnehmer einen *Zusammenhang* zwischen virtuellem Inhalt und der realen Umgebung erfährt.

In Kapitel 4 untersuchen und illustrieren wir die verschiedenen Arten, auf die sich das Virtuelle und das Wirkliche aufeinander beziehen

(und sich dadurch gegenseitig ergänzen) können. Damit befassen wir uns mit der zweiten Schlüsselfrage dieser Forschungsarbeit: "Welche Formen kann AR annehmen?". Auf einer fundamentalen Ebene unterscheiden wir AR von Szenarien, in denen Teilnehmer keinen Zusammenhang zwischen dem Virtuellen und dem Realen erfahren. Wir identifizieren dann zwei Kernbeziehungen, die AR ermöglichen: *(a) räumliche Beziehungen* zwischen dem Virtuellen und dem Realen (hier scheint virtueller Inhalt im realen Raum zu existieren) und *(b) inhaltliche Beziehungen* zwischen dem Virtuellen und dem Realen. Danach hinterfragen wir, auf welche Weise virtueller Inhalt seine reale Umgebung beeinflussen kann. Basierend auf der Rolle, die der virtuelle Inhalt im echten Raum spielt, unterscheiden wir zwischen fünf Formen von AR:

1. *Extended reality:* Szenarien, in denen das Virtuelle die reale Umgebung ergänzt.

2. *Diminished reality:* Fälle, in denen virtueller Inhalt reale Elemente scheinbar aus der echten Umwelt entfernt.

3. *Altered reality:* Umgebungen, in denen die virtuellen Informationen die augenscheinlichen Eigenschaften der realen Welt verändern.

4. *Hybrid reality:* Szenarien, in denen das Virtuelle eine echte Umgebung, die ohne die virtuellen Ergänzungen als unvollständig betrachtet würde, vervollständigt.

5. *Extended perception:* Fälle, in denen nicht wahrnehmbare, aber reale Aspekte der realen Welt in virtuelle Informationen übersetzt werden, die wir mit unseren Sinnen wahrnehmen können.

Anschließend gehen wir erneut auf Szenarien ein, in denen virtuelle Objekte scheinbar in der realen Welt existieren und diese erweitern. Wir stellen fest, dass das Vorhandensein virtueller Objekte im realen Raum Möglichkeiten für Beeinflussungen und Wechselwirkungen zwischen dem Virtuellen und dem Realen eröffnet. Auf dieser Ebene unterscheiden wir zwischen zwei Kernbeziehungen zwischen dem Virtuellen und dem Realen: *(1) physische Beziehungen* (das Virtuelle und das Reale scheinen sich gegenseitig physisch/physikalisch zu beeinflussen) und *(2) Verhaltensbeziehungen* (das Virtuelle und das Reale nehmen sich gegenseitig wahr und reagieren auf einer sozialen Ebene oder Verhaltensebene aufeinander). Danach besprechen wir kurz andere mögliche Beziehungen zwischen dem Virtuellen und Realen, wie zeitliche Beziehungen zwischen den beiden oder musikalische Beziehungen zwischen virtuellen und realen Instrumenten.

Kapitel 5 geht näher auf eine bestimmte Beziehung zwischen dem Virtuellen und dem Realen ein, und widmet sich der Interaktion zwischen beiden. Aufgrund der Tatsache, dass sich virtuelle Objekte nicht an physikalische Gesetze halten müssen und nicht direkt Kräfte auf

reale Objekte ausüben können, stellen wir uns die folgenden Fragen: Welche Arten der Interaktion zwischen dem Virtuellen und dem Realen sind sowohl möglich als auch glaubwürdig? Wir untersuchen (1), ob virtuelle Objekte auf eine realistische Art mit echten Objekten interagieren können und (2) ob sie auf eine phantasievolle, aber glaubhafte Weise interagieren können. Um diese Fragen zu beantworten, verfolgen wir sowohl einen theoretischen als auch einen praktischen Ansatz. Wir ziehen bestehende Forschungsarbeiten und AR-Beispiele zu Rate, führen eigene erste Experimente durch und reflektieren die Ergebnisse dieser Experimente. Dies führt uns zu der Schlussfolgerung, (1) dass virtuelle und reale Objekte Aktionen und Reaktionen, die wir aus der realen Welt kennen, glaubhaft simulieren können und (2) dass sie sich auch auf phantasievolle Arten, die in der realen Welt kein Äquivalent haben, gegenseitig beeinflussen können.

Kapitel 6 baut auf dem Gedanken auf, dass virtuelle Objekte sich von echten Objekten unterscheiden können. Wir untersuchen, ob virtuelle Objekte auch anders wahrgenommen werden können als reale Objekte. Um diese Möglichkeit zu untersuchen, entwickeln und präsentieren wir einen Prototyp dessen, was wir "Sonically Tangible Object" nennen. Konkret präsentieren wir einen virtuellen, unsichtbaren und nicht-taktilen Würfel, der in einem realen Raum platziert ist. Dieser Würfel kann durch explorative Handgesten erfahren werden. Wenn man den Würfel mit den Fingern berührt, werden sogenannte binaurale Klänge ausgelöst, die scheinbar von genau der Stelle ausgehen, an der das Objekt berührt wird. Unsere ersten Experimente legen nahe, dass dieser klang- und bewegungsbasierte Ansatz taktil-artige Empfindungen auslösen und die Anwesenheit virtueller Objekte im realen Raum vermitteln kann. Wir besprechen das Konzept, die Umsetzung und unsere Erfahrungen mit dem "Sonically Tangible Object" und stellen unsere Forschung in einen breiteren Kontext.

Kapitel 7 schließt die Doktorarbeit ab. Wir kehren zu unseren Hauptfragen zurück (*'"Was ist Augmented Reality?"* und *"Welche Formen kann AR annehmen?"*) und reflektieren die Antworten, die wir gefunden haben. Darüber hinaus gehen wir auf offene Fragen ein, die sich während dieser Forschungsarbeit ergeben haben und die wir jetzt, nachdem wir eine umfassende Literarturstudie durchgeführt haben, praktische Erfahrungen mit AR gesammelt haben und eine eigene umfassende Theorie der AR entwickelt haben, beantworten können. Des weiteren fassen wir Erkenntnisse zusammen, die das Designen von AR-Erlebnissen unterstützen können. Auch besprechen wir methodologische und technologische Einschränkungen unserer Forschung und weisen mögliche Richtungen für zukünftige AR-Forschung und -Entwicklung auf. Unter anderem schlagen wir vor, sich weniger auf die Nachahmung unserer bestehenden Realität zu konzentrieren und stattdessen neue, fantasievolle und kreative Formen von AR zu schaffen, die in einer rein physikalischen Welt kein

Äquivalent haben.

Acknowledgements

This PhD position would never have existed, weren't it for Edwin van der Heide and Yolande Kolstee. Thank you for creating this position in a joint effort, and for providing an inspiring and interdisciplinary environment for my research.

This endeavor was not always easy, and I would have not been able to finish it without the support of my supervisors, colleagues, friends, family and Edwin van der Heide in particular. During this enterprise, he has acted as a daily supervisor and has done the work of an inofficial co-promotor. This thesis is the result of our dialogues, discussions and debates. Edwin, I cannot thank you enough for your support. Thanks for thinking along, questioning my thoughts and providing feedback at every step of the way.

Bas, thank you for your refreshing view on science, sharp thoughts and for introducing me to Kasteel Slangenburg. Fons, thank you for being there to help 'drive' this thesis over the finish line and for the opportunities you provided me with since I started studying in the Netherlands.

This trajectory was long, yet time flew by because I have been in great company. First, a big thank-you to my friends and colleagues in Leiden: Amalia, thanks for your friendship, for listening and giving advice at every road bump. You are an amazing friend. Wim, thank you for always being around when I needed a friend who's enthusiastic about AR. Lu, thanks for being your cheerful self and for introducing me to hip hop dancing. Zane and Christoph, thanks for keeping me company during the difficult parts, and for the great hours we spent working together! Staas, Joris, Ionica, Katy, Nanda, Jan, Marcello - thanks for inspiring me, everyone in their own way. Thanks to the binaries—you are the best crew a PhD student can wish for. My gratitude also goes out to my friends and colleagues at the AR Lab. In particular, I want to thank Esmé, Jouke, Mariana, Reba and Yolande (again). You made working on the AR[t] magazine a blast. My colleagues in The Hague: Thanks for helping me gain "teaching muscles and endurance" and for making this intense process fun! My colleagues and friends at Tilburg University: Thank you for being so supportive. Alwin, thanks for your friendship and for introducing me to new types of research. Thank you, my new colleagues at Radboud

University. I am glad to complete the last little details of my thesis in such a inspiring and motivating environment. Team IRMA, thanks for reminding me that teamwork is the best kind of work!

My friends in and around Leiden: Maarten, Jessy, Joey, Danica, Vincent, Yvette. Thank you for making me feel home in the Netherlands. Thanks to my friends in Germany: No matter how long we haven't talked, it will always feel like it was just yesterday!

My dear family-in-law: Hanna en Cees, Pytrik, Jennifer, Elina*, Cian, and Jesper, Josefine, Edwin, Jaro and Fabian. Thanks for being my "Dutch family".

Many thanks to my aunt Heidi, my great aunt Helene, my grandma Sonja* and my cousins Ronja, Andreas and Daniel. Thanks for reminding me that the world is bigger and more diverse than the academic world.

A million thanks to my parents, Edith and Willi. You have both supported and inspired me for more than 35 years. You have always fostered my curiosity and encouraged me, even with endeavors that many other parents would have considered crazy. This thesis is for you.

Since day one, I have worried about writing these acknowledgements and forgetting someone who should be mentioned. If you expected to find your name here but could not find it, this *thank you* is yours.

Last but not least, Jurriaan. You augment my reality. Thank you for all the little and big things. But most of all, thank you for being you.

Curriculum Vitae

Hanna Schraffenberger is a creative researcher working in the intersection between science, art and technology. Her research interests include human-computer interaction, augmented reality, interactive art as well as privacy and security.

Hanna Schraffenberger was born in Stuttgart (Germany) in 1983 and studied Audiovisual Media (BEng) from 2005 until 2008 at Stuttgart Media University. Subsequently, she has been a graduate student at Leiden University, The Netherlands. She completed the Media Technology MSc program with honors in 2011. In 2011, Hanna Schraffenberger has started her PhD research at the Media Technology research group at the Leiden Institute of Advanced Computer Science (Leiden University). Her research examines the fundamental characteristics and potential manifestations of augmented reality (AR). During the first three years of her PhD research, she has been working as a researcher in residence at the Augmented Reality Lab (formerly based at the Royal Academy of the Arts, The Hague, The Netherlands).

During the second half her PhD research, Hanna Schraffenberger has worked as a lecturer at the "Communication & Multimedia Design" program at The Hague University of Applied Sciences (2015-2016) and at the "Cognitive Science and Artificial Intelligence" group at Tilburg University (2016-2017). She now works as an Assistant Professor at the Artificial Intelligence department of Radboud University (Nijmegen) as well as at the Donders Institute for Brain, Cognition and Behaviour. In addition, she volunteers as a front-end developer and UX engineer for the Privacy by Design Foundation.

Hanna Schraffenberger has has developed and taught courses on human-computer interaction, augmented and virtual reality, data visualization, app development, interaction design, usability and programming. In addition to conducting research and teaching students, she is interested in communicating science to a broader audience with both conventional and unconventional kinds of media. She has acted as the editor-in-chief of the AR[t] magazine, in which researchers from all over the world share their interest in augmented reality, discuss its applications in the arts and provide insight into the underlying technology.