



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



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

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 new-found 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 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., [Billingham, 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 such-like. 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*.

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

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.

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 ([Green 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](#).

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

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.

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).

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

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.



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.*).

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,

⁶ It is rather ambivalent whether music and television 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.



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 allow 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 (Piekariski 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

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.

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.

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 Lary 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

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

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

⁸Of 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.

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

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.

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-

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.

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.)

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 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

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.

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 object 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 world. 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](#)).

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

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 \(2010\)](#).

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).

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

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.

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}

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

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.

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 re-compose 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 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.

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

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.

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.

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

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

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.

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.

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 CO₂ 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 lighting. (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.

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

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](#).

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 Espresso bar 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

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

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.

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 CO₂ 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

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

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

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).

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.

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, Encarnacao, 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. On 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-

relationships 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 “out-door” 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.

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 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

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.

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 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 transform the real environment or real objects in the environment. We propose the term *altered reality* to describe this sub-form 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 drink's 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.

