Arguably augmented reality : relationships between the virtual and the real

Schraffenberger, H.K.

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

6 *The Invisible Cube: Introducing Sonically Tangible Objects*

As the previous chapter has shown, augmented reality often mimics reality. Researchers strive for photorealism and aim for scenarios where virtual objects cause the same occlusions and shadows as physical objects (see, e.g., Gibson and Chalmers, 2003; Madsen, Jensen, and Andersen, 2006). Similarly, scientists include physics simulations to make virtual objects adhere to physical laws and move like real objects (e.g., Chae and Ko, 2008; S. Kim, Kim, and Lee, 2011). In line with this, many AR projects allow users to interact with virtual content in the same way as they would interact with real physical objects (e.g., Corbett-Davies, Dünser, Green, et al., 2013; Ohshima et al., 1998). Simply put, much AR research and development focuses on making virtual objects as real as possible.

There is nothing wrong with this aim. Many contexts, such as training environments or exposure therapy sessions require a realistic setting to be effective. However, the imitation of reality is not the only path towards creating meaningful AR experiences. A unique power of virtual objects is that they do not have to look, feel or behave like real objects. In our opinion, AR research can harness this power and create imaginative forms of realities that offer new and unique experiences. Accordingly, our research follows another direction: Instead of imitating reality, we want to create new experiences that have no equivalent in a purely physical world. We are interested in how augmented reality scenarios can differ from strictly physical, 'unaugmented' environments.

In the previous chapter, we have explored this idea with respect to interaction between the virtual and the real. However, the virtual objects used in this preceding exploration still have mimicked physical objects—both with respect to appearance and behavior. For instance, we have used virtual spheres that look like physical spheres and that seemingly are affected by gravity (as well as other imaginative forces). In this chapter, we want to focus in on the idea that virtual objects do not have to mimic physical objects and explore alternative manifestations of the virtual. To do so, we create a virtual object that does not look, feel or behave like any real object. Although we are not interested in imitating reality, our project shares important goals with existing AR research: it tries to convey the presence of a virtual object in an otherwise real environment.

Like in the previous chapter, we build on the idea that virtual objects can differ from real objects. However, in this chapter, we take this concept one step further. If virtual objects have different qualities than real objects, they might also be perceived differently from how real objects are perceived. For instance, while most real objects can be seen, we might not sense a virtual object simply by looking at it.

Building on the ideas that virtual objects (1) can differ from real objects and (2) can be perceived differently from how we perceive real objects, we have developed a new kind of virtual object—the so-called sonically tangible cube. Unlike real objects, this cube is invisible and does not provide tactile feedback. However, touching the virtual cube triggers binaural sounds that appear to originate from the exact spot where it is touched. Our initial experiments show that through this sonic feedback, virtual objects can gain an almost-tactile quality and appear as if they were actually present in real space. It is this idea of making virtual objects both tangible and present through spatial sonic feedback that makes "sonically tangible objects" unique.

Several questions have fueled the development of the invisible cube and our research into sonically tangible objects. For instance, we were wondering if it is possible to leave out the tactile component in tangible perception. If there is no tactile stimulation, would the virtual object still be perceived as part of real space—and if so, would it be experienced as an object with a tactile or physical component? We were intrigued by how one experiences an object that provides no tactile sensations. Most importantly, however, we were eager to learn more about how virtual objects in an AR environment can differ from real objects.

While we provide preliminary answers to these questions, the main contribution of this chapter lies in the proposed concept of sonically tangible objects. This concept, to the best of our knowledge, is new. So far, inferences regarding the perceptual qualities of the invisible cube are based on informal testing and on our subjective experience with the cube.

The central idea—that a virtual object can be tangible but not tactile—calls for a distinction between the terms tangible and tactile. In this chapter, things are called tangible, if they can be perceived by touching (being in contact with) them.[1] Only objects that also stimulate the tactile receptors (as found in the skin and tissue) are referred to as tactile. This understanding creates room for objects that are tangible but not tactile.

[1] We use the word "touch" to refer to gestures where the body is (brought) in contact with an object.

Unlike in previous chapters, we will discuss issues related to implementation and technology in more detail. This is because this knowledge is necessary to reproduce and experience sonically tangible objects. Another reason for this is that different concepts, such as making sounds appear from specific positions in the real environment, are strongly interlinked with the chosen implementation, such as the use of binaural recordings (this recording technique will be explained in more detail later on).

One primary goal of this exploration is to create a new type of object and with it, a new type of experience. In line with this, we focus on the practical exploration of sonically tangible objects rather than on theory. However, by pursuing this idea, we also hope to learn more about the possible manifestations of AR. Furthermore, we hope to better understand the concept and context of sonically tangible objects. This is why we will have a look at sonically tangible objects from different perspectives and place it in the context of related research.

The chapter is structured as follows: In the following section 6.1, we focus on the practical side of sonically tangible objects. We share choices made and insights gained during the development of the invisible cube, describe the used setup and implementation and discuss our experience with it. In section 6.2, we place the cube in the context of pertinent research and compare our project with related work. Because our project is multi-disciplinary, we consider research from various fields, such as augmented reality, tangible interaction and perception. The chapter ends with a reflection on the project and possible directions for future research (section 6.3).

6.1 *The Sonically Tangible Cube*

The sonically tangible cube is a virtual object. It is unlike any real object in the sense that it is non-tactile, invisible and lacks many physical properties, such as weight and temperature. It does, however, have sonic and spatial properties such as a shape, loudness and sound texture. Although the cube has no tactile component, its presence can be perceived through touch. When fingers enter the cube, sound appears to originate from the spot where the virtual object is touched. The resulting sonic feedback not only corresponds to the fingers' positions but also fits the movement of the fingers. Fast finger movements result in more agitated soundscapes while slower movements cause less dense, more distinct feedback. As the cube is virtual, fingers can move through it and explore its inner texture.

6.1.1 *Implementation and Setup*

The virtual cube is 20 cm x 20 cm x 20 cm of size and it floats 10 cm above the otherwise empty work desk of the author. The technical setup consists of a Leap Motion Controller (see `www.leapmotion.com`), which detects the position of the participant's fingertips in real space. The Leap Motion is placed on the desk and senses hand movement

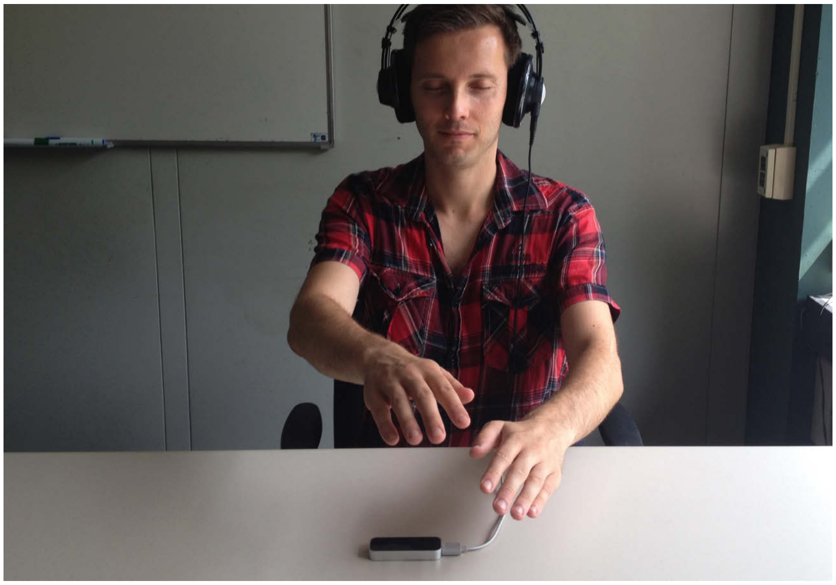

Figure 6.1: A colleague explores the virtual, invisible and non-tactile cube. A Leap Motion Controller is used to track the position of his fingertips.

above the device (see figure 6.1). A custom *Max* (2014) patch, which runs on an Apple Mac mini, interprets the data provided by the Leap Motion. Interfacing with the Leap Motion device is realized with a Max external object 'aka.leapmotion' by Akamatsu (2014). In our current setup, the frame rate of the Leap Motion device is around 57 fps when the office is naturally lighted and slightly above 200 fps when the amount of interfering infrared light is reduced by darkening the room. The Max patch evaluates whether and where the participant is touching the cube on the basis of the fingers' coordinates. If the fingers are located within the 20 cm x 20 cm x 20 cm area that has been defined as the cube, their movement triggers pre-recorded binaural sounds. This interpretation of the finger position works for every finger independently and allows the participant to explore the cube with up to ten fingers at a time.

A constraint of the current setup is that the sound only matches the fingers' position if the participant is sitting at the right spot and directly facing the cube.[2] Also, due to the frame rate of the Leap Motion device, very fast hand-movement can cause a mismatch between the hand-position and the spatial information of the triggered sound. Moreover, finger movement is sensed best, if the hands are held horizontally. Fingers that are not in a direct line of sight from the Leap Motion device cannot be sensed.

[2] This is the case because the pre-recorded sounds are positioned in relation to the listener. If a sound originated from right in front of a listener while being recorded, it also appears right in front of the listener when it is played back. If one moves while listening to a recording, the recorded sounds move along. Hence, for the cube to remain at the intended position, the participant can not move. The underlying binaural recording technique will be explained in more detail in section 6.1.2.

6.1.2 Development and Choices

The sonically tangible cube was developed iteratively during the course of several months. In the course of the project, the author acted as a researcher, developer and participant. Additionally, colleagues were occasionally asked to provide feedback and describe their experience with the cube.

From the beginning, we have explored the idea of making virtual objects tangible and present through sonic feedback. The topic of (in)visibility was left aside for future research and hence, many evaluations have been conducted with closed eyes. This choice was made because in a very early stage, it became clear that it was easier to concentrate on the sonic aspects with closed eyes. Furthermore, not seeing an object in the space seemed to potentially interfere with the audible information. We thus decided to first only focus on the tangible and audible experience. Aside from this, four determining observations and decisions were made concurrently in the early stages of the development process. These important choices concerned (1) the shape of the object, (2) the recording technique, (3) the recorded material and (4) the sound design (including the mapping between movement and sound)—together, these choices determine the properties of the invisible object.

The Shape: A Cube

One of the most fundamental early decisions regards the shape of the object. We have started out with several simple geometric shapes. For instance, we have used a laser pointer and triggered a sound every time the laser was interrupted to see if this would evoke the experience of touching a virtual string or line. Furthermore, we have tested triggering sounds whenever the fingers crossed a predefined plane in the real environment. We expected that this might create the feeling of 'crossing a virtual border'. However, our initial experimentation indicated that it is very difficult to experience a plane or a line. Running one's hands freely through a three-dimensional object and exploring both its borders and inner texture offered the most intriguing, tactile-like experience and promised to convey an object's presence in space best. Given that we deemed it best to begin the exploration with a *simple* 3D object with a very clear geometry, we decided to focus on a cube-shaped virtual object.

The Recording Technique: Binaural Audio

Another crucial decision concerns the sonic aspect of the project. In the beginning, simple synthesized clicks were played back in mono (feeding the identical signal to both the left and the right channel) through closed Beyerdynamics DT 770 Pro headphones whenever a virtual object was touched. This was done in order to learn about the effects of linking movement in a certain area in the environment to a very basic sonic response. However, our initial trials showed that the resulting experience was closer to being informed or 'being told' that one's hand had entered a predefined space and there was no direct sensory experience of an object in space. While there was a clear link between the movement and the sound, it only felt as of one triggered the sound through movement and not as if the sound originated from the virtual object.

This experience did not come as a complete surprise. After all, interacting with real objects and materials—crumbling paper, scratching on a surface, typing on a keyboard or moving the mouse—causes sounds that originate from the objects themselves and from the position where the objects are touched. Based on this, we assumed that in order to convey the presence of an object at a certain position in space, it could help to make it seem as if the object's sounds originate from its position.

Here the idea of using binaural audio to achieve this effect came into play. We had previously encountered this technique when reading about and listening to the work of Cardiff (n.d.), who uses binaural recordings for her audio walks (see also section 3.1.3). Binaural audio is based on the notion that hearing makes use of two signals: the sound pressure at each eardrum (Møller, 1992). If these two signals are recorded in the ears of a listener (or a dummy head), the complete auditive experience—including the three-dimensional spatial information of the sounds—can be reproduced by playing the signals back at the ears. By making binaural sound recordings in which the sounds originate from the spatial position of the virtual cube, and playing these recordings back via headphones when the fingers enter this area, the resulting sounds will sound as if they originate from the location of the fingers/object.

Based on our theoretical considerations, the use of binaural audio seemed promising. However, as the author had little practical experience with this recording technique, many questions remained open: Would binaural recordings indeed allow us to make sounds appear from any position within the room? Would these sounds be distinguishable from 'real' sounds. Would the experience really be that different from listening to mono-sounds or stereo recordings? In order to get a better idea about the actual potential and limitations of binaural recordings, we decided to first investigate its qualities with some simple experiments. For example, we recorded the sound of someone knocking on the closed office door and the sound of the ringing phone while working in the office. We intentionally chose sounds that originate from objects that are (already) physically present in the space. Likewise, we made sure to use sounds that are not accompanied by visible changes (it is not possible to see whether the phone is ringing or whether someone is standing behind the door).

From these initial experiments, it became clear that binaural audio indeed can convey the desired experience of 'something happening' in the real space and 'something being present in the space'. Listening back to the recordings while working, the sounds seemed to originate from those exact spots where they originally had happened and sounded equally real as the original sound. Consequently, the author often was in a state of doubt: had someone knocked on the door, was someone passing in front of the office door, did someone open the door to the toilet across the hallway, was someone actually calling? The virtual ringing of the phone was practically indistinguishable from a real call. Similarly, it was almost impossible to tell whether someone actually was waiting behind the door, or whether a knocking sound was played back from the recordings. In fact, a simple virtual knock proved powerful enough to communicate the presence of something or someone in real space.

Although the sounds were extremely realistic, there was also a simple and safe way for the author to distinguish between real and virtual events: Whereas real sounds remained unaffected by head movement, virtual sounds would move along! This happened because binaural recordings are positioned in relation to the listener. If a sound originated from right behind a listener while being recorded, it also appears right behind a listener when it is played back. If one moves while listening to a recording, the recorded sounds move along.

From our initial exploration, we concluded that binaural recordings could create a strong sense of something happening in the real environment. However, we also realized that in order to make a sound originate from a certain spot in space, we would have to determine the position and orientation of the listener. This is why we decided that for this initial realization of sonically tangible objects, participants would sit in the author's office chair looking straight ahead without moving their head.

The choice for binaural audio went hand in hand with a switch to the open AKG K702 headphone. Due to the open nature of the headphones, the recorded sounds mix in with the sounds naturally present in the environment. This additionally supports the experience that the virtual object inhabits our real physical space rather than a virtual or separate space.

The Recordings: Crumbles of Aluminum-Foil in a Plastic Bag

What should the virtual sonic object sound like? The choice of using binaural recordings introduced the question of what to record. We were searching for sounds that (1) are abstract (do not invoke the idea of a specific real object), (2) have a tactile quality (indicate touching) and (3) support the idea of a non-solid object/material that allows the fingers to move through it. Several different sound sources were tested: for example, foils, paper, plastics, packaging materials from everyday objects, rattles and empty bottles. All sounds were produced by interacting with the materials with the hands and fingers. This choice was based on the assumption that sounds that actually are created by hand/finger movement are more likely to fit the exploratory hand gestures of the participant and are more likely to create a tactile-like experience (similarly to how the sound of squeaking nails on a chalkboard can be an almost-tactile, physical experience, even if someone else is scratching the board). For the current implementation of the sonically tangible cube, we have settled on the sound of aluminum foil, produced by squashing a tiny plastic bag filled with small crumbles of the foil (see figure 6.2).

Figure 6.2: We recorded the sounds of squashing a tiny plastic bag filled with small crumbles of the aluminum foil (the image shows a recreation).

It was clear that this 'crumbling sound' has to be recorded at the intended position of the virtual cube. However, it soon became apparent that recording only one sample at the location of the cube was not convincing because the sound did not seem to move when the fingers were moving. We thus determined that when, e.g., moving one finger from the left side of the object to the right side of the object, the sound should move along with the finger. More generally, we decided sounds should seemingly originate from the exact position where the cube is touched (at the fingertips of the participant). To achieve this, we divided the cube into 64 sub-cubes of 5 cm x 5 cm x 5 cm (see figure 6.3) and recorded five-second samples of aluminum foil sounds at all 64 positions within the cube.

The 64 recordings were made with a ZOOM H4 audio interface and two DPA 4060 microphones. The microphones were placed slightly above the ear-entrance of the author and the sound was recorded with a basic Max patch. For the recordings, the author successively produced the desired sound by squashing the little plastic bag and rubbing the aluminum crumbles against each other at each of the 64 sub-

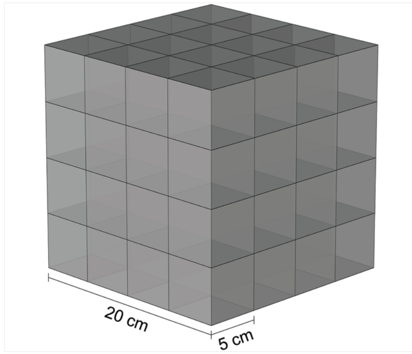

Figure 6.3: The sonically tangible cube was divided into 64 sub-cubes. A binaural recording was made at all 64 positions. Image of the cube contributed by Wim van Eck.

areas. Aside from this, the author was sitting motionlessly in front of the desk, facing the cube just like participants do during the experience (see figure 6.1).

Sound Design and Mapping: Movement Causes Sound

When a participant interacts with the cube, the positions of his/her fingers determine which of the 64 recorded audio samples are played back. If a finger is placed in a sub-cube, the corresponding recording is activated. However, first tests showed that simply playing back the recordings resulted in a sound that only matched the fingers' positions, but not the different variations in hand and finger movement (slow, fast, no movement, etc.). Because this felt not convincing yet, we experimented with more complex settings that map the movement of the fingers to parameters in the sound design.

Our current implementation knows two sound design settings. The first setting makes use of granular synthesis. Granular synthesis makes use of very short snippets of audio; so-called grains. These grains typically are between 1 and 50 ms long and can be layered (Roads, 1988). Playback parameters such as speed or volume can be varied for each grain individually. In our sound design, all grains are taken from the binaural recordings and between 10 ms and 20 ms long. A random offset is used to vary the position in the binaural recording from where each grain is taken. When a grain is played back, it is randomly varied slightly in pitch/playback speed.[3] As a result, every grain sounds differently. (This appeared to be crucial for the believability of the experience.) The key concept in this first setting is that the change of a finger's position triggers the playback of an audio grain taken from the corresponding 5-second recording. This means that fast finger movements trigger a lot of grains successively, whereas no movement does not trigger any grains.

[3] By varying the speed, the pitch automatically changes as well. Furthermore, the changes in playback speed also in influence the spatial characteristics of the sounds. However, as those variations were minimal this effect was perceptually negligible.

The second setting follows a similar underlying idea. However, instead of working with short grains, the entire 5-second recording is layered, varied and looped if there is movement in a sub-cube. In this setting, a faster movement activates more layers. For instance, moving slowly, one would only listen to one version of the recording. Moving a bit faster, a second (varied) version of the recording would join in, starting from a different position and playing in a slightly different speed.[4]

[4] Here, too, the changes in playback speed influence the spatial characteristics of the sounds. However, these variations were again small and the effect negligible.

The settings have many similarities: Both result in a louder, more complex and dense soundscape if the finger moves fast and in a softer, less dense but more distinct soundscape if the movement is slow. As this happens for each finger individually, the number of fingers used by the participant has a similar effect: The more fingers are involved in the exploration, the denser the resulting sound. For either setting, some movement is necessary to 'excite' the virtual cube and to elicit its sounds. No movement results in silence, even if the hand is placed in the cube. However, as it is impossible to keep one's fingers completely still, occasional slight trembling of the digits will cause corresponding sound output.

There are also differences between the settings. In particular, the two settings differ with respect to the textural nature of the sound. Whereas the granular synthesis results in a more gritty and rough soundscape, the layered loops produce a thinner, airier sound texture. Furthermore, the granular synthesis sounds a bit more abstract and less like a real-world recording. (This fits well with our intention not to mimic the real world.)

Ultimately, the choice of the recording technique, the chosen material we recorded and the chosen sound design and mapping all determine aspects of the sonic qualities of the cube—together, they determine how it sounds to touch the object.

6.1.3 Experiencing the Cube

Ideally, the previous section has provided some insight into how the invisible cube sounds. However, how does the cube *feel*—does touching the cube really feel different from simply moving one's hands through thin air? Do we experience the cube as present in space, do we perceive it as tangible? It is important to systematically investigate this by performing experiments with a group of unbiased participants in the future. In the following, we compare experiences with the cube to experiences we know from the everyday world. These comparisons are objective in the sense that some similarities between the cube and existing real-world phenomena are simple facts. However, in addition, the author also provides her own subjective account of how it feels to interact with the cube.

On some level, experiencing the cube has indisputable similarities with moving one's hand through a beam of light. When touching a beam of light, one can clearly see the beam's presence in space but one cannot feel it. Similarly, in the case of the cube, one can hear the cube's presence in space, but one cannot feel it. In both cases, there is no tactile stimulation on our fingertips. At the same time, the cube can also be considered an opposite to a beam of light. Whereas we can see a light beam but cannot experience it with touch gestures, we cannot see the cube but we can experience it with touch gestures.

Given those similarities and differences, does touching the cube actually feel like touching a beam of light? Judging from the experience of the author, similarities between the two indeed exist. In particular, both seem to provide a quite fascinating 'dissonance between the senses'. In other words, both create an intriguing experience where one sense tells us 'something is there', whereas another sense tells us 'nothing is there'.

In addition to touching light, experiencing the cube also has some undeniable similarities to feeling out a physical object blindly with one's hands. After all, it is only through the physical act of touching that we can perceive the cube in the first place. There is no notion of the object unless one is in contact with it. Also, like in typical haptic perception, the experience of the object takes time and happens through exploratory gestures with one's fingers. Furthermore, touching a real object can cause sounds at the position where the object is touched. The same happens when touching the sonically tangible cube.

So does touching the cube feel like touching a physical object blindly? When exploring the cube, the author indeed often was reminded of blindly interacting with a physical object. However, this might have been caused by the simple fact that the cube was primarily explored with closed eyes. Although exploring the cube with closed eyes was a conscious decision, the author noticed that closing the eyes also came naturally when exploring the cube. This likely happened because the presence of the virtual cube was experienced as stronger and more convincing when the eyes were closed. Possibly, the author tried to avoid the conflict between what she saw and what she heard. (However, as mentioned above, the conflict between senses also was experienced as fascinating.)

According to the author's experience, exploring the cube also feels similar to exploring a physical object blindly because a model of the object's shape emerges in one's mind over time. Of course, when interacting with the cube, the author was always aware that she was dealing with a cube of a certain size. However, she would still repeatedly reconstruct its shape and form an internal representation of the object over time. When exploring the cube, a mental image of a cubic cloud floating in the space before the researcher regularly emerged. In some way, the cube seemed to be 'just air'. However, at the same time, it also seemed clear that 'something was there'.

Although there are obvious similarities between the way we blindly interact with physical objects and the way we perceive the cube, there are also quite some key differences. One can, for example, not hold, move and turn the cube. Instead, it is possible to move right through the cube and explore its inner texture and structure. Unlike with physical objects, it is impossible to simply follow the contour of a sonically tangible object and to explore its shape that way (cf. Lederman and Klatzky, 1987). Rather, the contour can be perceived by repeatedly crossing (zigzagging around) the border of the object and moving in between the sonic space of the cube and the silent space surrounding it.

In the author's personal view, these differences between the cube and real objects are the most fascinating aspects of the project. According to her experience, perceiving and interacting with the cube indeed feels different than interacting with any real-world object. In particular, zigzagging around the borders of the cube feels fascinating. In addition, drawing three-dimensional shapes with one single finger inside of the cube is intriguing (see figure 6.4). This might be the case because here the relationship between movement and sound is experienced most clearly. However, the fact that the virtual cube is different from any real object also has a downside. The cube (at times) can evoke the frustrating feeling that it is impossible to really 'get hold of it' or to grasp it. In line with this, the author often wondered how she could move the cube around.

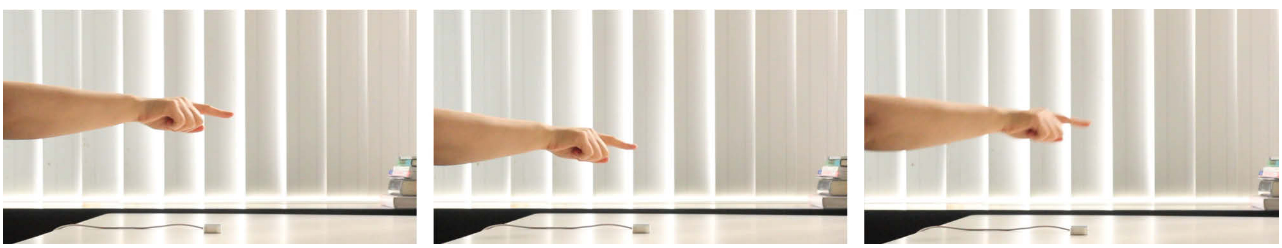

Figure 6.4: The author interacts with the cube.

Interacting with the cube is not only similar to touching light, or interacting with an object blindly—it also has undeniable similarities with playing gesture controlled open-air instruments such as the Theremin. The Theremin is played by moving one's hands in the space between two antennas. The position of the hands determines the sound. When interacting with the cube, movement of the hands in space similarly results in sonic output that corresponds to the position of the hands. The author has never played the Theremin or other gesture-based instruments. This makes it impossible to compare the experiences subjectively. However, in the author's experience, playing an instrument and exploring the cube are comparable; likely due to the cube's "sonic expressiveness". In fact, the cube seems to allow for some (but only very limited and basic) forms of musical expression. These become particularly apparent when holding the hand completely still inside the cube and then shaking the hands to a variable extent (see figure 6.5). By doing so, one can create a variety of rustling noises and exert a high amount of control over the resulting sound textures and volume. However, when compared to many existing instruments, there is little possibility for sonic variation and no possibility to control the pitch. (In this regard, the cube seems more comparable to the family of shaker instruments.) However, it is easy to imagine a similar Leap Motion setup that focuses on music generation rather than on communicating the presence of a virtual object in space.[5].

[5] Not surprisingly, the potential of the Leap Motion device in the context of new digital musical instruments has explored by several researchers in the past (e.g., Han and Gold, 2014; E. S. Silva et al., 2013)

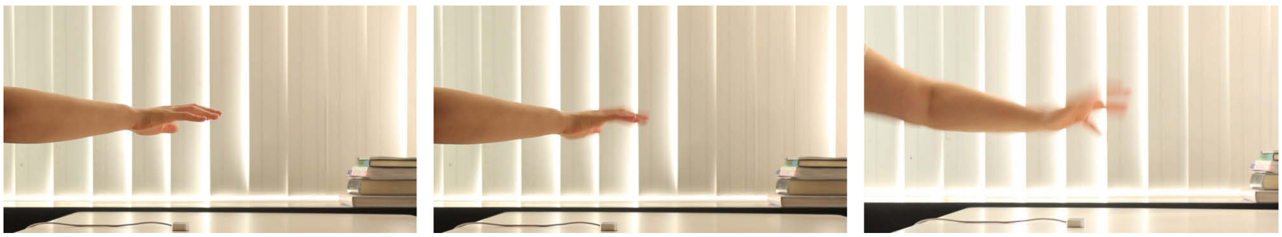

Figure 6.5: The author shakes the hand to a variable extent to produce different types of rustling sounds.

In the course of this study, the cube was mostly explored with closed eyes. This raises the question in how far the cube is experienced as part of (and in relation to) the real environment. When the author interacted with the sonically tangible object, the cube was certainly experienced as present in 'her environment'. In other words, the cube seemed to exist in the same space and environment as the author. Because the author felt present in her office (even with closed eyes), the cube also was experienced as existing in this office. However, the relationship between the cube and other elements in the office space was often experienced as rather weak. For instance, the fact that the cube was floating over the office table felt more like conscious knowledge rather than like a perceptual experience. At times, the link between the cube and the real environment seemed stronger when exploring the cube with open eyes. However, opening the eyes, in turn, seemed to weaken the experienced presence of the cube in space.

When it comes to experiencing the cube as a part of the real environment, an interesting observation was made while fine-tuning the sound design. At this stage, the author often interacted with the cube but also had a computer monitor placed on the desk to make changes to the mapping. The used patch (program) included a visual representation of the cube, as well as visual representation of the author's fingers (see figure 6.6). When this visualization was shown on the screen, the cube was no longer experienced as part of the real environment. Rather, it immediately felt as if one were reaching into the space depicted on the monitor and touching the virtual object shown on the screen. This happened although the sound was recorded at (and therefore indicated) a position in front of the screen.

A question that needs to be addressed is whether the tactile sensation is missing when exploring the cube. In the author's opinion,

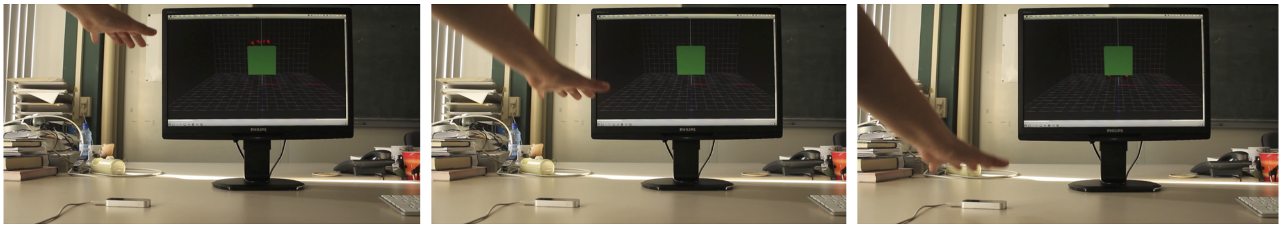

Figure 6.6: During the development of the sound-design, a monitor was placed on the desk. As soon as a visual representation of the cube and the fingertips was presented, the cube no longer felt like a part of the real environment. Instead, it felt as if one were reaching into the space depicted on the monitor and as if one were touching the displayed cube in this space.

this is not the case. Just like we do not miss a tactile sensation when listening to music or looking at an object, the author was not missing a tactile sensation when experiencing the cube. However, as mentioned above, hearing oneself touch the cube without feeling a tactile sensation on the fingertips felt somewhat like a contradiction between senses.

During this project, the author often wondered what bystanders would experience if they witnessed the author's (admittedly awkward looking) interaction with the cube. With respect to this, the author was presented with some surprising initial insights after presenting a video of her interaction with the cube at a colloquium. Naturally, the video (in combination with stereo sound) was not able to convey the experience of interacting with sonically tangible objects directly. Yet, an interesting observation was made: After watching the author interact with the cube and listening to the sonic results on speakers (rather than headphones), some audience members were under the impression that something was present in the space above the Leap Motion and that the author was interacting with this 'thing'. Other audience members, however, indicated that in their opinion, the space was clearly empty. Of course, these reactions are based on a video and were shared in a very informal context. As such, they do not yet answer how the project is experienced by bystanders that are present in the same space as the cube. However, some feedback from the audience suggests that the cube can not just be experienced directly, but also indirectly—by witnessing someone else's interaction with it.

To conclude: It remains difficult to put the experience of the cube in words. However, to the author one thing seems clear: Touching the cube is different from simply moving one's hands through thin air. When we move our hands through air, we feel nothing but empty space. The cube, however, is experienced as something that is present and as something that can be touched and that invites playful exploration. It seems to inhabit the space, albeit in a non-physical way. Although the experience is not tactile in the traditional sense, it definitely has tactile-like aspects. According to the author's experience, running once fingers through the object feels like 'something is here that can be touched'.

6.2 *The Cube in Context*

Our project is multi-disciplinary; it draws from and contributes to various fields of research, such as augmented reality, tangible interaction and perception. In this section, we take a second look at the cube from various different perspectives and discuss the project in the light of related research.

6.2.1 *Augmented Reality*

We have arrived at the concept of sonically tangible objects as part of our trajectory researching augmented reality. The invisible, sonically tangible cube can be seen as an AR project because virtual content is presented in and relates to the real environment. More specifically, the project is concerned with the presence of virtual content in real space, and like many AR projects, aims at making it seem as if additional virtual objects were part of the otherwise real environment. (We have termed this form of AR *presence-based AR* in section 3.4.) In this general sense, the sonically tangible cube relates to all projects, that aim to convey the presence of virtual objects in real space.

In existing research, the presence of virtual objects in the environment is typically conveyed by visually displaying the object in the space. Our cube is similar to traditional visual virtual objects because both do not have a tactile component. A difference between the sonically tangible cube and traditional visual virtual content is that visual content can be experienced without touching it. When experiencing visual virtual content, the absence of tactile stimuli only plays a role if and when one actually touches the object. When experiencing the cube, the absence of tactile stimuli is always apparent when it is perceived.

Our project explores whether the presence of virtual objects can be experienced through a combination of touch gestures and spatial sound. Hence, our project specifically relates to those AR projects that use sound and/or tangible interaction to convey the presence of (invisible) virtual objects in real space. A project where the presence of something virtual is perceived tangibly is Sekiguchi et al.'s (2005) so-called Ubiquitous Haptic Device. The little box, when shaken, conveys a feeling of a virtual object being inside the device. Similarly, a wearable haptic device by Minamizawa, Fukamachi, Kajimoto, Kawakami, and Tachi (2007), called the Gravity Grabber, allows participants to perceive the ruffle of the water in a glass, although he/she actually is holding an empty glass. More examples of objects that convey the presence of virtual objects in real space by means of tangible cues can be found in subsection 4.4.3.

Projects that let a participant experience the presence of "something that is not really there" by means of sound have been discussed in detail in subsection 4.4.2. An example is the SoundPacman game by

Chatzidimitris et al. (2016). This game makes use of spatialized sound in order to give game elements a position in the real environment and to communicate the location of virtual ghosts to the player. Other examples are Cilia Erens' (see Erens, n.d.) and Janet Cardiff's (see Erens, n.d.) sound walks. Both artists use binaural recordings of everyday sounds that blend in with the sounds present in the real environment when the participant navigates the space and listens to the composition on headphones.

Given that our implementation was inspired by Cardiff's approach, it is maybe not surprising that a discussion of Janet Cardiff's work by Féral (2012) also helps our understanding of sonically tangible objects. Féral defines "presence effects" as the feeling that an object (or body) is really there, even when one knows that it is not. This not only plays a role in Cardiff's works but also relates to the experience of the sonically tangible cube. While the ears make it feel as if the cube is present, the lack of tactile (and visual) stimuli informs us that nothing is there.

Our project shares important goals with much existing AR research. Most importantly, it shares the common goal of making it seem like a virtual object existed in real space. However, at the same time, our project also differs from most AR research in one important regard: We do not try to make it seem as if the virtual object was a real, physical object. In fact, the cube aims to be different from any real object. In contrast, many existing AR projects aim at creating a scenario where the user cannot distinguish between what is real and what is virtual. For instance, Azuma (1997) writes "[a]fter the basic problems with AR are solved, the ultimate goal will be to generate virtual objects that are so realistic that they are virtually indistinguishable from the real environment" (p. 380). In line with this, Vallino (1998) suggests that ideally, "[v]irtual and real objects are visually indistinguishable" (p. 20) and R. Silva et al. (2003) point out that "[a]lthough many AR applications only need simple graphics such as wireframe outlines and text labels, the ultimate goal is to render the virtual objects to be indistinguishable from the real ones" (p. 10).

We clearly do not share this 'ultimate' goal. Instead, our goal is to create an object that is so different from everything we know from the real world, that it evokes a new kind of experience. We thus believe, AR research has to distinguish between two goals: (1) making it seem as if a virtual object was a real, physical, material object and (2) making it seem as if a virtual object were really *present* in the space. Whereas those two goals typically go hand in hand in existing AR research, we only share the second goal. We do not mind if the object is experienced as 'virtual'. Ideally, when touching the cube, one would think "something is here, but it is not like any real object I know".

6.2.2 *Tangible and Embodied Interaction*

The cube deals with (in)tangibility, requires active bodily engagement and it explores the possibilities of a tangible experience without tactile stimuli. As such, our research relates to the field of tangible and embodied interaction.

The term 'embodied interaction' has been coined by Dourish (2004), who defines embodied interaction as "the creation, manipulation and sharing of meaning through engaged interaction with artefacts" (p. 126). Research into tangible and embodied interaction is an entire research field in itself, which in turn draws from and contributes to many other fields. However, one key aspect that links the sonically tangible cube to this field is the explicit link between interaction and experience. As Dourish (2004) points out, tangible interaction builds on the idea that we experience the world through directly interacting with it, and that acting in the world happens through exploring the opportunities it offers for action (see, e.g., p. 18).[6] This idea is also at the basis of the sonically tangible cube. One can only experience the cube through interaction with it, and perceiving the cube happens through action in the space and by exploring the possibilities for interaction. There is no way to perceive or experience the cube passively or without a body.

[6] In its original context, this point is part of a larger argument, namely that this is true for both tangible interaction and social computing. Accordingly, one of the main points of Dourish's book is that tangible and social computing share embodiment as a core element.

On the one hand, one could argue that this focus on perception through interactions sets the invisible cube apart from those common virtual objects that we can only perceive by looking at them. On the other hand, one could argue that all forms of perception—also seeing an object—entails interaction with the world. In fact, more recent theories of perception suggest that all perception is active. For instance, Noë (2004) states that "perceiving is a way of acting. Perception is not something that happens to us, or in us. It is something we do" (p. 1). Our point with respect to this is that perception is not always *experienced* this way. Seeing or hearing something often feels like it 'happens to us', whereas exploring the shape of an object blindly typically involves conscious active exploration. With the sonically tangible cube, tangible interaction and tangible perception are explicitly designed to be experienced as one and the same thing.

In our opinion, the sonically tangible cube can not only be placed in the context of AR, but also in the context of the growing body of embodied interaction research that is concerned with "bodily action, human experiences, and physicality, in the context of interaction with and through a world comprised of computationally mediated artifacts and environments" (Antle et al., 2011, p. 9). A specific contribution of our project in this context is a new form of tangibility that is evoked through a combination of spatialized sound and touch gestures and which does not involve any tactile stimulation.

We believe that an important link to explore further is the relationship between presence and tangibility. In our experience, tangibility and presence are closely linked. More specifically, presence is a necessary condition for tangibility. We can only touch an object if it is present. If we touch an object, we and the object are both present in the same space—at least in a mediated way.

6.2.3 Human-Computer Interaction

One possible area of application for sonically tangible objects is the field of Human-Computer Interaction (which overlaps with the above-mentioned field of tangible and embodied interaction), and in particular what Chan et al. (2010) call *intangible displays*. Intangible displays are visual virtual interfaces that appear in mid-air, in front of a user's eyes. Aside from simply displaying information they also allow for interaction: Users can touch virtual objects, such as buttons, with their physical hands. However, intangible displays do not provide tactile feedback when they are touched.[7] Chan, Kao, Chen, Lee, Hsu, and Hung (2010) address this lack of tactile feedback by providing visual and audio feedback. In their experiments, they played short sounds whenever participants touched the surface of the intangible display. While similar, their project differs from ours in the sense that sound is used to *inform* the user about the fact that they have successfully touched the object. The sound serves as feedback and not as an integral part of the object.

[7] As mentioned, when referring to the cube, we use the term tangible rather than intangible. We dot so because although there is no tactile feedback, the cube is perceived by touching (being in contact with) it.

Although originally not intended this way, the concept of sonically tangible objects could be used to improve the interaction with intangible displays. It could increase the spatial presence of the display, provide better feedback about the user's hand position and movement through the display and is likely to make the "the awkward feeling of 'touching' a mid-air display" (Chan et al., 2010, p. 2626) less awkward and more tactile-like.

Another related HCI project is the so-called BoomRoom (Müller et al., 2014). In this room, sounds seem to originate from certain spots in real space (this is realized with a circular array of 56 loudspeakers and Wave Field Synthesis). These sounds can be 'touched' in order to grab, move and modify them. Although related, their project differs in the sense that it focuses on the localization and direct manipulation of sound rather than on the presence and tangibility of virtual objects.

6.2.4 Perception

Our project relates to perception research, and in particular research into haptics, tactile illusions and cross-modal interactions as well as sensory substitution.

Haptics

The sonically tangible cube is perceived by explorative hand gestures. This links it to the field of haptics. Haptic perception typically involves active exploration (Lederman and Klatzky, 2009). Haptics is commonly understood as a perceptual system that derives and combines information from two main channels: kinesthetic perception and cutaneous sensation (Lederman and Klatzky, 2009). Cutaneous sensation is derived from the receptors that are found across the body surface and that allows us to feel, for example, pressure or temperature. The kinesthetic channel refers to perception of limb position and movement in space, which is derived from the receptors embedded in muscles, tendons and joints.

Kinesthetic perception also plays a key role in the perception of the virtual cube—it provides the participants with the information about where and how fast their fingers are moving in space. This awareness is crucial in order to link what one hears to one's movement in space. What makes the perception of the sonically tangible cube different from common haptics is the lack of cutaneous feedback (including tactile sensations). Rather than 'feeling something at the position where they touch an object' the participants 'hear something at the position where they touch the object'.

An aspect that would be interesting to explore systematically in the future is the use of exploratory gestures that one applies to explore sonically tangible objects. In "Hand movements: A window into haptic object recognition", Lederman and Klatzky (1987) identify six typical movement patterns that are used to explore the properties of real objects and link them to the type of knowledge they reveal about an object. For instance, Lederman and Klatzky (1987) show that we learn about the exact shape of an object by following its contour. As our project shows, other types of gestures apply when it comes to sonically tangible objects. For instance, we experience the shape by zigzagging in and out of the object. Furthermore, it is possible to explore the inner texture of an object, but there is no mass, weight or temperature to explore. In the future, it would be interesting to learn more about hand gestures used to explore the properties of *virtual* objects, and the role sound can play in conveying their properties. Based on our experience with the sonically tangible cube, we believe the sound caused by an exploratory gesture can play an important role, especially when it comes to textural properties of virtual objects.

Tactile Illusions and Cross-Modal Interactions

The sonically tangible cube aims to create a tactile-like experience. There are several studies that indicate that sound can influence actual tactile experiences. The "Parchment-skin illusion" (Jousmäki and Hari, 1998) shows that modifying the sounds that accompany hand-rubbing can influence the tactile sensation of the skin. It was found that accentuating the high frequencies can lead to the experience of a higher level of skin roughness. Hötting and Röder (2004) have discovered another auditory-tactile illusion. In their experiment, one tactile stimulus was accompanied by several tones. As a result, participants reported that they perceived more than one tactile stimulus. What sets these illusions apart from our cube is that in both cases, the participants were presented with a tactile stimulus.

A study that suggests that a tactile experience can be evoked without presenting any tactile stimuli has been reported in the context of Virtual Reality. Biocca, Kim, and Choi (2001) suggest that visual cues can cause haptic illusions. In their experiment, participants reported that they felt physical resistance when manipulating virtual objects in a virtual environment although the interface contained no haptic displays and the environment provided no direct stimulation to the haptic channel. However, it has to be noted that the participants were wearing gloves that allowed them to manipulate objects by pinching their fingers together. Hence, also here, tactile stimuli (from pinching the fingers together and from the gloves) were present. Thus, it remains unclear whether the visual cues evoked the tactile experience or simply altered present tactile sensations.

Sensory Substitution

One could argue that sonically tangible objects allow us to hear sounds instead of feeling a tactile sensation. In this sense, the cube relates to projects that use sound to substitute touch. One such sensory substitution system is F-Glove (Hafidh, Osman, Alowaidi, El-Saddik, and Liu, 2013). This haptic substitution system aims at helping patients that suffer from the symptoms of Diabetic Peripheral Sensory Neuropathy, such as sensory loss at the fingertips and resulting difficulties with manipulating objects. F-Glove uses audio feedback to inform the patients of the pressure they apply to objects. The volume of the sound is mapped linearly to the applied pressure. Unfortunately, it is not clear whether the system simply informs the patients of the pressure they use via sound or whether they start experiencing pressure directly, via the auditory sense. Naturally, the experience of the cube is quite different from not having a sense of touch, as your hand can simply reach through the virtual sonic object.

6.2.5 *Open Air Instruments and Sound Installations*

Our project relates to the field of sonic interaction. In particular, it relates instruments and installations that use hand or body gestures in free space to produce sound, such as the above-mentioned Theremin. Like our research, such gesture instruments and installations are based on a mapping between body movement and sound.

The artwork *Very Nervous System* (1986-1990) by David Rokeby (1986-1990) is an early example of an interactive sound installation where body movement in open space generates sound. However, the sound of such artworks and instruments like the Theremin usually does not appear to originate from the location of the movement, which is a key difference from sonically tangible objects. Furthermore, with few exceptions, they do not (try to) express the presence of virtual objects in space.

One exception—an instrument that actually does convey the presence of virtual objects in space—is the invisible drumkit by Demian Kappenstein and Marc Bangert (*The Invisible Drums of Demian Kappenstein and Marc Bangert* 2011). In their invisible setup, each virtual drum is placed at its regular position in space. Hitting the invisible virtual drums triggers pre-recorded samples of a real drum set. The position of the sticks and the speed of the movement determine which sample is triggered. Similarly to the cube, the virtual drum kit becomes perceivable through the interaction. Furthermore, based on our own experience with a video of the performance, the presence of the virtual drum-kit in space also becomes apparent through witnessing (perceiving) this interaction. However, to the best of our knowledge, sounds do not seemingly originate from the location of the drums. This makes the project fundamentally different from the invisible cube.

6.2.6 Science Fiction

The sonically tangible cube is unlike any real object and it is experienced differently from how we perceive real objects. We can find similar ideas of objects that differ from physical objects in fiction and science fiction. For example, the film Ghostbusters (Reitman, 2004) features ghosts whose presence can be sensed with the help of custom devices. A key difference between any fictive objects we know from movies, stories and books is the fact that one experiences our cube's presence in space directly, through a new sensory combination of touch and sound. We are not aware of any such object being described or depicted in stories. However, we believe there is much to learn about how virtual objects could look or behave from the domain of fiction.

6.3 Reflection and Outlook

With the sonically tangible cube, we have introduced a prototype of a sonically tangible object and a new, sound-based form of augmented reality. The proposed cube is invisible and non-tactile. According to our experience, it is nonetheless perceived as spatially present in our real, physical environment. This suggests that virtual objects do not have to look or feel like real objects in order to be a believable part of our real, physical space.

The virtual cube is non-tactile and yet tangible. The experience of the cube can be seen as one possible answer to the question of how it could feel to touch an object that provides no tactile feedback. According to our impression, the virtual sonic object offers an almost-tactile experience that has no equivalent in a purely physical world. However, this still has to be confirmed by experiments with unbiased participants.

The current implementation of the cube primarily serves as a proof of concept. While we are happy with its current state, we have many ideas on how to improve the cube and explore the concept of sonically tangible objects further. For instance, it would be interesting to allow the participant to manipulate the cube in space. During presentations of the project, we have repeatedly encountered the question whether one could, e.g., move the cube in space or resize the object.

Concerning the sonic qualities, future experiments can reveal which sounds are most suitable for creating tactile-like experiences and possibly test whether sounds that are created with the hands work best. It would be interesting to find out more about how to sonically represent imaginary material and communicate different densities, textures and shapes with sound.

So far, we have chosen to work with binaural recordings. In the future, it will be valuable to explore computational methods for simulating the sounds' origins in space. If this is successful, it will be much easier to allow participants to move through space freely and experience the cube from different positions. Furthermore, it will be simpler to create polymorphic sonically tangible objects of different shapes and sizes and to place them at various positions and in different spaces.

One aspect that was left aside so far is the topic of (in)visibility. This offers several intriguing directions for future research. For example, we are eager to learn how participants interpret the absence of visual clues. On the one hand, it might lead to a contradiction between senses: "I can hear it, but I see that nothing is there". On the other hand, the lack of visual stimuli could be interpreted as a property of the object: "Something is there, it is invisible". Furthermore, it would be interesting to compare the experience of the cube with open and closed eyes, and, as an additional condition, also add a visual dimension to the cube (e.g., by means of a head-mounted display) to learn more about the influence of (in)visibility on the experience.

One limitation of this research is that so far, our inferences are based on informal tryouts and our own subjective experience with the cube. Our experience might not fully represent how others perceive the cube and we cannot entirely rule out the possibility that it is influenced by the expectations and hopes we have for the project. We hope to extend the presented research and conduct experiments with unbiased participants in the future.

Whereas this study focuses on the experience of the author when interacting with the cube, another interesting direction to pursue on the future is how the cube is experienced by a bystander who witnesses the interaction. Based on our experience when presenting the video of the author's interaction with the cube, we expect that seeing the interaction and hearing the sonic result (even if it is played back on speakers) can create the impression of an invisible object being part of the space. In this context, it would also be interesting to know if bystanders imagine some kind of tactile stimulation when witnessing the interaction.

A limitation of the current setup is that the participant cannot freely move their head while experiencing the cube. This is due to the use of binaural recordings. Another constraint that stems from the fact that recordings were used, is that the cube can only be experienced in the particular office of the researcher and at the particular spot where the sounds have been recorded. Playing the sounds back in another room or at another position would likely sound weird. This is because the qualities of the room (such as the fact that the room is small and windows reflect the sounds from the left) have shaped the recordings. (If the sounds were, e.g., played back in a big room, one would immediately notice the lack of reverb.)

In many ways, sonically tangible objects are the culmination of our preceding research into augmented reality. The project builds on various ideas discussed in earlier chapters. For instance, it takes up our claim that AR is not just something we see, and that virtual content can take non-visual and multimodal forms (cf. chapter 3). It furthermore builds on spatial relationships between the virtual and the real, which have been identified as an important aspect of existing AR projects in chapter 3. In addition, it focuses on the presence of an added, virtual object in real space—a common form of AR that we have discussed in detail in section 4.2. What is more, the project takes up the issue of interaction between participants and virtual content, which has been a topic of interest in the preceding chapter. Similarly, it builds on the idea that virtual objects do not have to follow physical laws and can differ from real objects, that likewise has been a major point of interest in chapter 5.

Although the project has emerged within the context of AR research, it also raises questions that go beyond the field of AR and that fall outside our own area of expertise. For instance, it would be interesting to learn more about what happens on a perceptual level. Are sound and kinesthetic information combined, similarly to how cutaneous information and kinesthetic information are integrated in traditional haptic perception? Can the combination of spatial sound and kinesthetic information lead to cross-modal interactions? What happens if the spatial information of the audio does not match the position of the fingers? Do we perceive the lack of tactile stimuli as "something missing" or do we fill in this information? We have put much emphasis on describing the concept in a way that allows others to reproduce it and we want to invite researchers to join our investigation of sonically tangible objects.

Subject: Interview

------- Forwarded message -------
From: Hanna Schraffenberger
Date: 2013/4/3 15:53
Subject: Interview
To: Lev Manovich

Dear Lev,

Maybe you remember me from Facebook. I work at the Augmented Reality Lab in The Hague and I am one of the editors of the AR[t] magazine. When I read your article “The Poetics of Augmented Space”, I realized that I would like to interview you about augmented reality for our magazine. A short time ago, I finally also read your book “The Language of New Media”. As a consequence, I’d like to interview you even more. So I hope you’ll agree to an interview?

Best regards,
Hanna

P.S. The questions are in the attachment.

What is Augmented Reality?

To begin with, I would like to ask you what you consider augmented reality (AR) to be. In "The Poetics of Augmented Space" (Manovich, 2006) you describe AR as "the laying of dynamic and context-specific information over the visual field of a user" (p. 222). It would be great if you'd address the topic once more. Firstly, because our readers might not have read your article. And secondly, because I think that this point of view unnecessarily limits AR to the visual sense.

In "The Poetics of Augmented Space", you mention Janet Cardiff's audio walks as great examples of laying information over physical space. These walks are designed for specific walking routes. While navigating the environment, one gets to listen to a mix of edited sounds that blend in with the sounds of the surroundings, as well as spoken narrative elements and instructions such as where to go and what to look at (see Cardiff, n.d., 1991). In contrast to 'typical' visual AR, the user is presented with auditory information that relates to the immediate surrounding space. Personally, I would call this augmented reality. Wouldn't you?

Augmented Space

What is special about AR compared to other forms of Augmented Space? In your article "The Poetics of Augmented Space" you discuss the concept of Augmented Space. Augmented Space refers to all those physical spaces that are overlaid with dynamic information such as shopping malls and entertainment centers that are filled with electronic screens and all those places where one can access information wirelessly on phones, tablets or laptops. Besides AR, you mention several other technological developments in the context of Augmented Space, among which, for example, monitoring, ubiquitous computing, tangible interfaces and smart objects. Is AR just one of many related recent phenomena that play a role in overlaying the physical space with information? What's special about AR compared to other forms of Augmented Space?

What else can be augmented?

Something I really like about your article is that you see augmentation as an idea and a practice rather than a collection of technologies. However, so far, you have only discussed the augmentation of space. I was wondering whether you have considered other manifestations of augmentation as well. I don't think augmentation is limited to a space or an environment. I'd even say that often it's not the space that is augmented, but something else.

For example, you mention software that performs tasks according to the mood, pattern of work, focus of attention or interests of their user. However, I am doubtful whether our experience of a space is affected by this kind of information. Let's imagine that my phone registered that I have been sitting still for a long time and reminds me to take a short break to stretch my legs. This information relates to one individual in the space (me), to the activity the person is performing (sitting still), but I don't think it has anything to do with the surrounding space. Hence, I might consider it an augmentation of the activity (not moving, sitting still) or an augmentation of the user (me), but I don't consider it an augmentation of space.

Edwin van der Heide (my colleague and supervisor) and I have recently given this topic a lot of thought, and we were fascinated by the questions: "What is actually augmented in augmented reality? What else can (we imagine to) be augmented?" We came up with the answer, that in AR, something virtual augments something real. More specifically, the virtual augments that to which it relates. In our view, space is one of the possibilities, but likewise, we have considered things like augmented objects, augmented humans, augmented perception, augmented content and augmented activities. What is augmented depends on what the additional content relates to. I am curious whether you'd agree. Do you think that all forms of augmentation bring along an augmentation of space or influence our experience of the immediate surrounding space?

Information and space—one coherent gestalt?

In "The Poetics of Augmented Space" you raise a question that intrigues me a lot. Do the real space and the dynamically presented information add up to one single coherent phenomenological gestalt or are they processed as separate layers?

I am a bit of a sound-person and it has always fascinated me that sometimes the sounds of a radio seem to mix in with environmental sounds. For example, the ticking of a red streetlight might perfectly mix in with the rhythm of the song that is currently playing. Listening to a radio play, an event could sound so real and so nearby, that I'd turn around, just to find, that nothing is happening there. But of course, most often, the sound of the radio just exists as a separate, independent layer of content. The voice of the newsreader doesn't mix with the voice of my colleague, nor does it relate to my environment. Most of the time, a song is just a song and has nothing to do with the surrounding space— until someone starts dancing or tapping their foot. So judging from my experience of listening to the radio, information and the surrounding space can be perceived as one single mixed thing as well as independently. But besides these two options, there are more possibilities. For example, the newsreader might tell me about a traffic jam and thereby inform me about my immediate physical space. Here, the information and my spatial surroundings aren't perceived as a single gestalt, but nevertheless, there is a relationship between both. I think the same is true for Augmented Space. Often, information and space might be related, even when they don't add up to one phenomenological gestalt. So some questions I'd like you to answer with respect to Augmented Space are: When information and space are perceived independently from each other—would you still call these occurrences Augmented Space? What if information and space are perceived as separate but related layers? And more fundamentally: When and why do information and space add up to one single gestalt?

New Media

One of the main questions I want to ask you is: What makes augmented reality special? I have posed that question with respect to other forms of augmented space. I'd like to ask it again with respect to the history of new media.

Personally, I don't think of AR as a recent phenomenon. Of course, there are more and more so-called AR applications, AR technologies and new media works that work with AR. However, when we consider the concept of AR, we find examples that date back centuries. An example of ancient AR is the Pepper's Ghost trick (Lamers, 2013). It uses a second room, glass and special lighting in order to let objects seem to appear, disappear or morph into each other in an otherwise real, physical environment.

But even if the concept isn't new, current manifestations of AR might still bring something new and special to the table. If we look at contemporary AR and compare that with other forms of new media, what's special about it and what isn't?

AR and the second space

From *The Language of New Media* (Manovich, 2001), I understood that throughout media history, the screen was used to separate two absolutely different spaces. For example, this function of the screen applies equally to renaissance paintings and to modern computer displays. When we imagine a typical AR scenario in which virtual objects are integrated into a real scene (e.g. a virtual bird is sitting on a real tree) there is no second space. It's the same physical space, which appears to contain both virtual and real elements. Is this a fundamental change in visual culture?

AR and the quest for realism

The quest for realism in computer graphics is something that has always bored me. You note that new technological developments illustrate how unrealistic the previous existing images were. At the same time, they remind us that current images will also be superseded. I was wondering: How does AR fit in the widespread aspiration towards realism? On the one hand, visual AR could be considered a huge step back. The 3D models that are usually integrated into real space don't come close to the level of photorealism we know from cinema. On the other hand, the virtual leaves the realm of virtual space and enters our real physical environments—with respect to that the images might be experienced as more realistic than ever...

Will AR take the quest for realism to a new level? I can imagine, when striving for realism, the virtual things that appear to exist in our physical space should not only look like real things—ideally, they also feel like them, smell like them, taste like them and behave like them. Will photorealism be traded in for a form of realism that encompasses all senses? Do you think new media will develop towards a more multimodal form?

AR and cinema

In *The Language of New Media*, you relate different forms of new media—e.g. Virtual Reality, websites and CD-ROMs—to cinema. How about the relation between AR and cinema?

I'm certainly not a cinema expert, but I guess most of what we see in visual AR has been present in cinema for a long time. For example, AR research is very concerned with registering virtual objects in real space. As far as I understand it, this can be seen as an analogy to compositing in films: an attempt to blend the virtual and the real into a seamless whole 'augmented' reality. Do you agree?

You oppose compositing to montage: while compositing aims to blend different elements into a single gestalt, montage aims to create visual, stylistic, semantic, and emotional dissonance between them. Do we have montage in AR as well? For instance, you give the example of montage within a shot, where an image of a dream appears over a man's sleeping head. The same could easily be done in AR. So I would think, AR can learn from cinema both with respect to compositing and with respect to montage. However, I also wonder: Does cinema use other techniques to create fictional realities that are not (yet) used in AR? Does AR use techniques that might be adapted by cinema in the future?

AR as spatialized databases

One of the main claims in *The Language of New Media* is that at their basis, all new media works are databases. You argue that what artists or designers do when creating a new media work, is constructing an interface to such a database. More specifically, you write about the elements of a database:

"If the elements exist in one dimension (time of a film, list on a page), they will be inevitably ordered. So the only way to create a pure database is to spatialize it, distributing the elements in space." (p. 238)

In AR, virtual and real elements are distributed in real space. Can we understand this as a pure database? What are the consequences of working with spatialized elements? What are the inherent limitations and possibilities when working with this form? (I can imagine it has consequences, e.g. for storytelling? As you point out, we cannot assume that elements will form a narrative when they are accessed in an arbitrary order.)

AR and future research

With *The Language of New Media*, you did not only provide a theory of new media; you also pointed your readers towards aspects of new media that were still relatively unexplored at that time and you suggested directions for practical experimentation. Are there certain aspects of augmented reality you consider especially interesting for future experiments and explorations?