

SwitchAR: Perceptual Manipulations in Augmented Reality

Jonas Wombacher
TU Darmstadt
Darmstadt, Germany
jonas.wombacher@tu-darmstadt.de

Zhipeng Li
Department of Computer Science,
ETH Zurich
Zurich, Switzerland
zhipeng.li@inf.ethz.ch

Jan Gugenheimer
TU Darmstadt
Darmstadt, Germany
Institut Polytechnique de Paris,
Télécom Paris - LTCI
Paris, France
jan.gugenheimer@tu-darmstadt.de




Figure 1: Visualization of the main steps of SwitchAR with redirected walking (RDW) as an exemplary perceptual manipulation (PM). 1: Based on an aligned virtual reconstruction of the environment, the user is switched from Augmented Reality (AR) to Virtual Reality (VR) without noticing it. 2: Once the user is in VR, existing PMs like e.g. RDW can be applied. In this example, the RDW can redirect the user to create an offset between their perceived virtual position and their actual position in the real room. The figures in this paper require color to convey their messages. Therefore, please print them in color or, ideally, view them on a screen.

Abstract

Perceptual manipulations (PMs) like redirected walking (RDW) are frequently applied in Virtual Reality (VR) to overcome technological limitations. These PMs manipulate the user's visual perceptions (e.g. through rotational gains), which is currently challenging in Augmented Reality (AR). We propose SwitchAR, a PM for video pass-through AR leveraging change and inattention blindness to imperceptibly switch between the camera stream of the real environment and a 3D reconstruction. This enables perceptual manipulations in what users still perceive as AR. We present our pipeline consisting of (1) Reconstruction, (2) Switch (AR -> VR), (3) PM and (4) Switch (VR -> AR), and discuss its challenges and our solutions.

Author pre-print - UIST '25, Busan, Republic of Korea
2025. ACM ISBN 979-8-4007-2037-6/2025/09
<https://doi.org/10.1145/3746059.3747595>

In a user study (n=20), we found that no participant noticed the switch and only one the PM. Additionally, despite revealing that a manipulation happened, participants could not detect the switch in a consecutive run. SwitchAR is a fundamental basis enabling AR PMs.

CCS Concepts

- Human-centered computing → Mixed / augmented reality.

Keywords

Augmented Reality (AR), Virtual Reality (VR), Perceptual Manipulation (PM), Redirected Walking (RDW)

ACM Reference Format:

Jonas Wombacher, Zhipeng Li, and Jan Gugenheimer. 2025. SwitchAR: Perceptual Manipulations in Augmented Reality. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25), September 28–October 1, 2025, Busan, Republic of Korea*. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3746059.3747595>

1 Introduction

Virtual Reality (VR) applications can make use of a variety of perceptual manipulation (PM) techniques. The subset called Virtual-Physical Perceptual Manipulations (VPPMs)[57] allows developers to influence the users' physical movements. One application for VPPMs is making the users unconsciously change their walking routes through redirected walking (RDW)[41], enabling them to explore virtual environments larger than their real environment. Another application is manipulating hand movements[3, 34, 37], making it possible to increase the availability of passive haptic feedback without the need for additional physical props. This is possible because VR head-mounted displays (HMDs) fully occlude the real world and, therefore, have full control over the users' visual field of view (FOV). This allows them to leverage perceptual thresholds by manipulating the virtual world displayed to the users.

Current Augmented Reality (AR) HMDs only add virtual content to the real world, rather than fully occluding it. This reduced degree of control over the environment, which is still visible to the user, makes it impossible to apply common PMs. Recently, advances in camera quality and size of the FOV allowed AR HMDs based on pass-through camera feeds to gain popularity (e.g., Meta Quest 3, Apple Vision Pro). This has led to pass-through AR becoming a widespread form of AR. Such pass-through AR HMDs are not only able to run AR, but they can also fully occlude the real world in order to run VR instead.

We present SwitchAR, a PM for pass-through AR leveraging change blindness and inattention blindness to imperceptibly switch between the camera stream of the real environment and a virtual 3D reconstruction. By implementing tricks like a distraction task, permanent visual noise and minimization of the visible changes, the switch can be hidden in spite of inaccuracies in the reconstruction. As a result, the user still thinks that they are inside of AR and interacting with the physical environment. At the same time, the system provides full control over the environment reconstruction, allowing developers to use VR PMs.

SwitchAR's pipeline consists of four steps. (1) Creating a 3D reconstruction of the environment. (2) Secretly switching from the pass-through video feed to the reconstruction. (3) Applying a PM. (4) Optionally switching back to the pass-through feed. Across this pipeline, for the whole duration of users wearing the HMD, it is important that they continuously think they are using AR.

In order to validate our approach, we implemented a SwitchAR prototype using RDW as an exemplary PM, since it is the most widely used PM and requires a long and non-trivial manipulation. Our RDW implementation was based on rotation gains applied when users turn around at waypoints, similar to Razaque's approach[41]. This SwitchAR implementation was then evaluated in a user study with 20 participants, where each run of the study consisted of three rounds. We implemented a distraction task, letting participants walk from one virtual painting to the next one

repeatedly in our lab, thinking they were taking part in a memory study. At the end of the second round, the redirection was revealed.

The study's goal was to measure (1) if unsuspecting participants detected the switch, (2) if they detected the PM, and (3) if the PM still worked when the participants knew they are being manipulated. Noticeability of the switch was measured similarly to the original approach of Simons et al.[48] in their inattention blindness study. Participants were asked four increasingly specific and revealing questions, like "Did you notice that the environment changed?" and "Have you been in AR for the whole duration of the study?".

For the open questions, we coded the results and considered all answers, that mentioned a visual change in the environment, as noticing the switch (the raw data of the responses is attached in the supplementary material). We found that none of the participants noticed the switch from AR to VR in any of the experiments' three rounds, and only one consistently noticed the redirection. After revealing to them in round two that they were being manipulated with RDW, 17 of the participants were not able to correctly explain how this manipulation was possible in AR. Even in the third round, at which point the participants already knew about the application of RDW, no participant noticed the switch and only one additional participant noticed the RDW. These findings demonstrate that SwitchAR can be a foundational approach enabling the application of various VR PMs in video pass-through AR.

Our main contributions are (1) the concept and implementation of SwitchAR, most importantly the novel secret switch between AR and VR, (2) the pipeline and a discussion of its challenges with possible solutions and (3) a demonstration of the feasibility of SwitchAR, even after revealing the utilization of a manipulation, in a user study (n=20).

2 Related Work

Our work is related to five main fields, incorporating insights from *Change and Inattention Blindness* and *Cross Realities*, and distinguishing itself from existing research on *Diminished Reality*, *Perceptual Manipulation in Virtual Reality* and *Perceptual Manipulation in Augmented Reality*.

2.1 Change Blindness and Inattention Blindness

Change blindness and inattention blindness are two related effects, which make it possible to leverage shortcomings in human perception to change users' environments without them noticing.

As discussed by Simons et al., change blindness describes "*the inability to detect changes to an object or scene*"[49]. This effect can appear both after longer durations, for example when turning away from an object and looking at it again later[54], and after short disruptions like eye movements or a short occlusion of the whole FOV[49]. Change blindness can be used in combination with VR to achieve different goals, e.g. for optimizing rendering times through changes in the rendering fidelity[6], automatically turning pages without the reader noticing it[59], or redirecting users to overcome space limitations for natural walking[54].

Inattention blindness is about not consciously perceiving parts of the FOV that lie outside the center of attention[48]. This can also include overlooking changes like the appearance of unexpected

objects[30]. In their study, Simons et al.[48] showed that, overall, 46% of their participants did not notice a person carrying an umbrella or wearing a gorilla costume walking across the frame when watching a video of basketball passes. By adding a distraction task, in this case counting the number of passes, the participants' attention was captured continuously, enabling the inattentive blindness. One of the findings was that an increased difficulty of the distraction task led to a decreased detection rate of the unexpected person walking through the video. Utilizing inattentive blindness in VR, Marwecki et al.[33] built a system allowing them to make changes to the virtual scene inside the FOV, without the user noticing it. The presented application scenarios included hidden difficulty adjustments, seamlessly adjusting the experience to the user's interests and decreasing motion sickness.

Change blindness and inattentive blindness are core aspects of SwitchAR, providing a theoretical foundation for realizing an undetected switch from pass-through AR to a virtual reconstruction.

2.2 Cross Realities

The recent years brought advances in the quality and availability of HMDs implementing AR based on pass-through camera feeds. Standalone HMDs like the Meta Quest 3 or the Apple Vision Pro are completely closed off, enabling AR by rendering the pass-through feed on their displays. Because of that, they are also able to switch between different levels of the reality-virtuality continuum[35, 36], including hiding the real world altogether in order to run in VR.

Sun et al.[55] presented a system monitoring the users' emotions and efficiency at their current task. If the system detects, that the user might be more efficient or in a better mood in a different environment, it suggests them to switch from AR to VR or vice versa and continue their task in the new environment.

Gottsacker et al.[12] compared different techniques for visualizing interruptions by outsiders to VR users. While most of the techniques relied on virtual representations of the outsider, therefore staying in VR, one of them utilized transitioning between VR and AR. In this case, the application would hint the users to turn on the HMD's pass-through cameras, allowing them to interact with the other person in AR.

Another direction of research about transitions along the reality-virtuality continuum is the usage of an intermediate replica environment, which is modeled after the physical environment. Such replicas can, e.g., enable more gradual transitions when switching from AR to the replica first and then from the replica to VR, instead of directly switching from AR to VR. Previous work investigated, e.g., the effects of intermediate replica environments on user experience, presence and distance perception[40, 51–53].

The concept of switching between different levels of the reality-virtuality continuum is another central aspect of SwitchAR. In contrast to the presented research on transitions, which users are aware of, SwitchAR is targeted on an imperceptible switch to enable applying PMs while the users still think they are using AR. The most similar works with regard to virtual reconstructions are from Lindlbauer et al.[27], reconstructing the environment with depth cameras, and Kari et al.[17], utilizing an iPhone with the Polycam app instead. Both created virtual reconstructions aligned with the

real environment, but they did not investigate fully replacing the environment with the reconstruction without the user noticing.

2.3 Diminished Reality

Diminished Reality, a concept originally introduced by Mann[31], has been applied in AR to remove content from the user's FOV for a variety of use cases. It can, e.g., enable users to see through solid objects[2, 16] or hide irrelevant objects to reduce distractions[19, 25].

Diminished Reality approaches can be technically similar to SwitchAR, relying on aligned environment reconstructions to remove objects[22]. However, they typically have the goal to hide individual objects, which means that most of the physical environment remains visible. In contrast, SwitchAR is replacing the whole environment with the reconstruction, imposing different requirements on the reconstruction and overall implementation. While this allows to remove individual objects as well, it additionally opens up a new space of PMs in AR.

2.4 Perceptual Manipulation In Virtual Reality

The concepts of perceptual manipulations (PMs) and more specifically virtual-physical perceptual manipulations (VPPMs) were defined by Tseng et al.[57] and Bonnail et al.[4]. PMs "*are mechanisms grounded in limitations of users' cognition and perception with a clear intention to influence users towards a specific outcome*"[4]. VPPMs, being a subset of PMs, "*are perceptual manipulations that are grounded in visual-haptic limitations with the intention to nudge the user's physical movements*"[4].

Beneficial Applications. Various research about VPPMs was conducted with the goal of benefiting the users. Razzaque[41] introduced small, imperceptible scaling factors to the users' rotations. This allowed them to decouple the virtual paths from the real paths walked by the users and manipulate them. The goal of this redirected walking (RDW) was to enable users to move around freely in virtual environments, which are larger than their actual physical environments, while still being able to use real walking as a locomotion. RDW can also be implemented by relying on change blindness instead of visual perceptual thresholds, e.g. by repositioning virtual doorways while the user is not looking at them[54].

Azmandian et al.[3] took advantage of the visual perceptual thresholds in order to enable one physical object to provide haptic feedback for several virtual ones, called haptic retargeting. By manipulating the virtual model of the users' hands, the virtual environment around the users, or both, they were able to direct users to grab the same physical cube each time, when they thought they would be reaching out to different virtual cubes instead.

Hand redirection can also be applied to improve ergonomics, e.g. by allowing airplane passengers to use physical surfaces for haptic feedback while keeping a comfortable posture[34], or by making it easier to reach virtual objects in mid-air to reduce fatigue[37].

As demonstrated by Clarence et al.[8], RDW can also be combined with haptic retargeting. This increased the spatial range of the typical haptic retargeting approach, reducing the need to place many physical props distributed across the user's environment.

Malicious Applications. In contrast to such beneficial applications, there is also a potential for malicious applications. Tseng et al.[57] conducted speculative design workshops to identify potentially harmful application scenarios. One of the resulting type of attacks was called "*puppetry attack*"[57]. The ability to manipulate the users' body movements could be exploited to make the users hurt themselves. For example, RDW could be applied to make them fall down a staircase, while haptic retargeting could be used to make them unknowingly reach out to dangerous objects in their environment, like knives.

Casey et al.[5] investigated potential attack scenarios on the HTC Vive and Oculus Rift VR HMDs. By exploiting software vulnerabilities on the HMDs, they were then able to implement redirection attacks. In a user study, they applied RDW to secretly redirect unsuspecting participants while they were playing VR games.

While there is a variety of these kinds of PMs, either beneficial or malicious, they are based on VR. In contrast to this, SwitchAR is enabling such PMs for video pass-through AR. Therefore, the next section is going to address existing research for PMs in AR.

2.5 Perceptual Manipulation In Augmented Reality

While there is some research about PMs being applied in AR, there is little for VPPMs. As long as users can still see the real world around them, the options for creating sensory conflicts are limited.

Cheng et al.[7] applied virtual overlays to manipulate the users' reaction times to content displayed on a 2D monitor. With a custom approach relying on cropping of the pass-through video feed, Ishii et al.[14] implemented a system for manipulating users' walking directions in AR. It was not able to continuously redirect in the same direction and participants had to walk in a straight line without looking sideways, though.

Just like PMs in VR allow for malicious applications, the same is true for PMs in AR. Roesner et al.[43, 44] and Lebeck et al.[23, 24] considered potential dangers originating from AR HMDs' abilities to cover parts of the user's FOV, e.g. traffic signs being occluded. To address such dangers, they discussed and prototyped safety mechanisms for controlling the visual output of AR applications.

The presented projects applied PMs in AR[7] and redirected knowing users with a custom video cropping approach[14]. However, they did not investigate physically manipulating unknowing users through existing VR VPPMs in AR. In order to apply such VPPMs in perceived AR, SwitchAR is transferring the users into a virtual reconstruction of the environment, without them noticing it. The conceptual steps of this process are going to be explained in the following section.

3 Concept

This section is going to introduce the four steps of SwitchAR's pipeline, as visualized in figure 2. The pipeline combines reconstructing the user's environment, switching from AR to this virtual reconstruction, and then applying the desired PM in order to enable PMs in perceived AR. The optional fourth step of switching back to AR can help to provide a seamless experience.




Figure 2: SwitchAR's pipeline, consisting of four steps with their main goals

We call this state "perceived AR" since the user is constantly thinking they are operating in AR. While we are fully aware that the concept is transitioning users into VR and then applying common PM techniques, we consider this a strength of our approach since we do not need to adapt any of the existing algorithms. However, we acknowledge that the manipulation is working in VR, but still consider it an AR manipulation, due to the fact that users would not be able to tell the difference between constantly being in AR and using SwitchAR.

3.1 Reconstruction

The goal of this first step is to create a virtual reconstruction of the environment that the application is going to be used in. It should be as similar to the real environment as possible, facilitating inattentive blindness and change blindness to make the following switch less noticeable. This similarity includes the reconstruction's aspects of geometry, scale, alignment with the real environment and its texture.

3.2 First Switch

After preparing the reconstruction, the next goal is to switch the HMD from running in AR to running in VR, relying on inattentive blindness and change blindness to prevent the user from noticing it. There is a variety of approaches to this switch from AR to VR to choose from. A straightforward option would be, for example, to instantly activate the virtual reconstruction, fully covering the pass-through video feed. The resulting abrupt changes might draw attention and are more noticeable than gradual changes, though[11, 61, 62]. Fading in the virtual reconstruction by gradually increasing its opacity until it is fully opaque could therefore be a less noticeable and similarly uncomplicated alternative approach.

3.3 Perceptual Manipulation

Once the application is running in VR, PMs already known from VR research can be applied, for example haptic retargeting[3, 8] or RDW[41, 54]. During this phase, the users are not supposed to notice the fact that they are not using AR any more. Therefore, perceptual thresholds for the manipulations should not be crossed.

SwitchAR in combination with PMs manipulating the virtual reconstructions introduces a new side effect, which does not exist when applying the same PMs in VR. Initially, the reconstruction is aligned with the real environment regarding position and orientation, in order to allow for an imperceptible switch (step 1 in figure 3). VR does typically not require alignment between the virtual world and the real environment. When the reconstruction is then manipulated, this alignment will be broken. RDW based on




Figure 3: Exemplary timeline of emerging misalignments (steps 1 to 3) between the physical environment and the reconstruction (colored blue for visualization purposes), and their resolution (steps 4 and 5)

rotational gains, for example, introduces misalignment in the form of offsets in the location and the orientation of the reconstruction (steps 2 and 3 in figure 3). For the manipulation to stay unperceived across the whole usage, this misalignment has to be resolved later on, as described in the next step.

3.4 Optional Second Switch

As a last step, the HMD can switch back from VR to AR without the user noticing it, based on the same principles used for the first switch. For this to be possible, any misalignment created in the previous step has to be resolved first (steps 4 and 5 in figure 3). This can be achieved by applying the previous manipulations of the reconstruction in reverse, or by directly manipulating it to restore the alignment with the real environment. Doing this allows the PMs and the virtual reconstruction to go completely unnoticed.

Depending on the application scenario, this second switch can also be omitted. If the developer intends to manipulate the user only once, or if the manipulation will be obvious to the user anyway, there is no need to resolve misalignments and switch back to AR.

4 Tricks

Based on the foundation of the conceptual pipeline, this section is going to explain the set of tricks, as listed in table 1, which we applied at different steps of the implementation. Figure 13 visualizes the mapping between the pipeline steps and the tricks in effect.

Most of these tricks are based on one of the two following principles. Research on inattention blindness suggests, that unexpected occurrences are less likely to be noticed, if they are more similar to objects, that are currently being ignored on purpose[10, 38, 60]. For SwitchAR, this means that it is desirable to increase the similarity between the virtual reconstruction, whose appearance is not supposed to be noticed, and the real world around the user, which is being ignored while the user is focused on some virtual objects.

Another factor reducing the noticeability of unexpected occurrences is an increased cognitive load, caused by a distraction task with increased difficulty[48]. Therefore, such a distraction task can be applied as part of SwitchAR, strengthening the effect of inattention blindness to reduce the noticeability of the switch from the pass-through feed to the virtual reconstruction.

| Name | Description | Goal |
|---------------------|---|---|
| Style Transfer | Match the reconstruction's color representation to the pass-through feed more closely | Visual similarity |
| Visual Noise | Compensate for missing camera noise on the reconstruction | Visual similarity |
| Virtual Cubes | Let user interact with virtual cubes in AR | Remind user of the AR scene they interacted with in the beginning |
| Body Representation | Cover the missing body in VR with a virtual representation | Visual similarity |
| Distraction Task | Distract the user, increase cognitive load | Distraction, cognitive load |
| FOV Covering | Reduce the amount of visible change by covering part of the FOV with a virtual painting | Reduce visibility of the changes |

Table 1: Overview of the custom tricks applied in our SwitchAR implementation

4.1 Style Transfer

To achieve a high quality reconstruction, we used a rig¹ with three mirrorless cameras to capture images of the room, rather than extracting frames from a pass-through screen capture of the HMD. This created the problem of differences in the color representation between the captured images and the pass-through camera feed, therefore reducing the similarity between the virtual reconstruction and the real AR environment.

In order to increase this similarity, we applied a style transfer based on wavelet transforms[63] to the captured images before using them to create the virtual reconstruction. For the reference style input, we took screenshots from the pass-through video of our Meta Quest 3 HMD. The captured images were paired with the most similar style image before running the style transfer.

The style transfer reduced the differences in color representation, making the resulting reconstruction's texture match the visuals of the pass-through cameras more closely.

Figure 4 shows the effect of the style transfer on a single input image, while figure 5 shows a comparison of screenshots from the pass-through feed, a reconstruction without style transfer, and the final reconstruction with style transfer.

4.2 Visual Noise

Once the virtual reconstruction is available in the HMD, there is another source of dissimilarity between the reconstruction and the real environment. The Quest 3's pass-through AR already contains a small amount of visual grain originating from the camera stream. Because of this, we found that the virtual reconstruction, which

¹<https://rd.nytimes.com/projects/assembling-a-camera-rig-to-capture-complex-spaces-in-3d/>, last visited: 11.09.2024



Figure 4: Comparison of one of the raw input images with the result of applying the style transfer to it



Figure 5: Comparison of pass-through feed with reconstructions before and after the style transfer

did not have any of this visual noise, looked static and artificial in comparison, lacking the visual movement caused by the noise.

To solve this problem and further increase the similarity between reconstruction and real environment, we manually added visual pseudo-random noise, based on an asset from the Unity Asset Store², both to the AR and the VR part of the application. This noise was stronger than the pass-through cameras' noise, and it was displayed on everything, that was supposed to be perceived as part of the real world. Intentionally virtual objects, like paintings or the cubes, were excluded from the noise.

This way, the lack of camera noise on the virtual reconstruction did not stand out as much any more, while the noise also had the benefit of covering up small defects or unevenness in the mesh and the texture. Figure 6 visualizes the noise's effect on the pass-through feed, and figure 7 shows the same for the final reconstruction.



Figure 6: Effect of the visual noise on the pass-through video feed

²<https://assetstore.unity.com/packages/vfx/shaders/fullscreen-camera-effects/old-movie-270021>, last visited: 11.09.2024



Figure 7: Effect of the visual noise on the reconstruction

4.3 Virtual Cubes

Next, the users were instructed to interact with a few virtual cubes in AR before starting the actual task. They could use the Quest 3's hand tracking capabilities to pick up virtual cubes, stack them and throw them around. This allowed the users to accommodate with interacting in a pass-through AR application and create an individual and unique structure with the cubes, which we preserved in the reconstruction after the switch.

The purpose of this was for the users to be able to recognize their own custom cube placement, whenever they looked at the cupboard with the cubes. Our assumption was that this might strengthen their belief that the reconstruction they are seeing is still the AR scene in which they initially interacted with the cubes. Among others, the cubes can be seen in figure 7.

4.4 Body Representation

Whenever the user is inside the virtual reconstruction, the HMD is essentially running VR. This means that the user's body is not visible anymore. When the body is at the center of attention but is invisible, this might break the AR illusion. In case the body is not in the center of attention, it is again a source of dissimilarity between the virtual reconstruction and the real environment.

To solve this problem, we added a virtual avatar for the user in AR and in the reconstruction, increasing the similarity between the reconstruction and real environment. We used an avatar³ slightly larger than the user's body, most importantly at the arms, to occlude most of the real body in AR and therefore make it less apparent that it is not visible in VR.

The initial problem can be seen in figure 8, where the user's hands and arms completely disappear in the virtual reconstruction. First, we added visual representations for the hands, which were included in the Quest 3's hand tracking implementation provided by Meta's SDK (figure 9). Then we extended the virtual representation with a full-body avatar. Provided with a rigged avatar, the Quest 3 offers upper body tracking and lower body prediction based on its inside-out tracking cameras. Figure 10 shows the combination of the hand visualization with the avatar's arms and hands, and the view looking down on the body is depicted in figure 11.

Combining the hand visualization with the avatar, rather than only using the avatar, made it look like there is still something

³<https://www.mixamo.com/#/>, last visited: 11.09.2024

below the avatar's hands, even in the virtual reconstruction. This might make the difference less apparent when the real hands are not visible any more.



Figure 8: Comparison of the pass-through feed and the reconstruction without any virtual body representation



Figure 9: Comparison of the pass-through feed and the reconstruction with a virtual hand representation



Figure 10: Comparison of the pass-through feed and the reconstruction with both the virtual hands and the avatar

4.5 Distraction Task

Without any task, users might explore and closely examine their environment, reducing the opportunities for exploiting inattentive blindness. To address this problem, we introduced a distraction



Figure 11: Comparison of the pass-through feed and the reconstruction when looking down on the feet in the final implementation

task. Users were supposed to look at a sequence of virtual paintings and images, memorizing their content, with limited time for each painting. Having to memorize the paintings' content with time pressure introduced cognitive load, strengthening the effect of inattentive blindness. Actively guiding the users' attention toward virtual paintings also reduced the attention on the environment.

4.6 Covering a Large Part of the FOV During the Switch

Executing the switch while the user is looking at their environment could make potential changes caused by differences between the virtual reconstruction and the real environment visible to them. While it would be possible that the user does not notice it due to inattentive blindness, this still poses a risk of the switch being noticed. In order to reduce this risk, a large part of the user's FOV had to be filled with virtual content during the switch. This strongly reduced the amount of visual changes in the environment visible to the user, therefore increasing the chance of the switch being unnoticed due to inattentive blindness. As the distraction task already contained looking at virtual paintings, we utilized this to make sure the majority of the user's FOV is covered as well. At first, the paintings were hidden through a grey placeholder. The users had to stand close to the picture frame and look at its center in order to display its painting. The timeframe, during which the virtual painting took up a large part of the FOV, could then be used to execute the switch.

Figure 12 shows two screenshots, approximating the user's FOV while looking at a painting, both before ("pass-through") and after ("final reconstruction") the switch.

5 Implementation

Figure 13 shows an overview of the steps required for implementing SwitchAR. The following section is going to explain how we implemented each of these steps in Unity (version 2022.3.10f1). The code of our Unity project will be made available for public access⁴.

⁴<https://github.com/wombacher/UIST25-SwitchAR>



Figure 12: Comparison of looking at a painting in the pass-through feed with in the reconstruction




Figure 13: SwitchAR's pipeline, with the tricks applied at each of the four steps

5.1 Reconstruction

At the time of writing, the Quest 3 did not provide real-time access to its pass-through camera feed, which would be necessary for generating a virtual reconstruction of the environment at runtime. However, with the recent access to the camera stream, in the future the reconstruction could potentially even be created in real-time on the device. For our implementation, we prepared a room that was not changed for the duration of our project. This way, we could use photogrammetry to generate a reconstruction once in advance, and then use that reconstruction in the application later on. We took a few steps to make the photogrammetry reconstruction easier. First, we removed most of the furniture in the room to make its geometry simpler. Then we blocked the natural light coming through the windows to get controlled lighting by the lamps on the ceiling. Finally, we attached ArUco markers to all large, featureless surfaces, like the white walls and grey cupboards. This was necessary to prevent holes in the reconstruction, which occurred whenever the pictures contained too few features to be aligned with each other.

Once the room was set up, we built a camera rig with three Sony Alpha 7R IV cameras, paired with Tamron 28-200mm 1:2.8-5.6 Di III RXD lenses. After taking about 3000 pictures of the room, we applied the style transfer explained in section 4.1 to them.

After the style transfer was done, we used the RealityCapture⁵ photogrammetry software to calculate a virtual reconstruction, including both the mesh and the texture. For performance reasons, we simplified the mesh to 250 000 triangles, allowing the Quest 3 to render it without a drop in frame rate. Because of differences in the rendering between RealityCapture and Unity, combining multiple textures resulted in dark lines along the texture seams in Unity,

⁵<https://www.capturingreality.com/>, last visited: 11.09.2024

which were not present in RealityCapture. To circumvent this, we restricted the reconstruction to use a single texture only, with the maximum resolution of 16384 by 16384.

The final reconstruction we chose still contained some obvious defects. The worst of them were located at positions, where the users would typically not frequently look at them, like some holes in the ceiling, shown in figure 14. One of the lamps in the room was also not reconstructed well. As we presumed that it might occasionally be visible in the users' FOV, we cut out that part of the reconstruction. Instead, we replaced it with a manually modelled version of the lamp. Figure 15 compares both versions of the lamp.

When preparing the implementation for the user study, we also placed a virtual curtain into the environment. The curtain was visible all the time, blocking the view on one end of the room. This allowed us to have an interviewer in the room with the participant, hidden behind the curtain to prevent the risk of breaking the illusion of still being in AR, which would arise when the interviewer is not visible anymore after the secret switch to VR. Figure 16 shows one end of the room being hidden by the virtual curtain.



Figure 14: Example of some of the final reconstruction's defects




Figure 15: Comparison of the lamp from the photogrammetry reconstruction with the manually modelled one

In order to align the virtual reconstruction with the real world, we implemented a manual two-point alignment. By placing two anchor points with the controllers, we could correctly scale, move and rotate the virtual reconstruction. Once we were content with the reconstruction's alignment, we used the spatial anchor feature of the Quest 3 to store it, eliminating the need to do the alignment process again every time the application is started.



Figure 16: A virtual curtain hiding one end of the room

The virtual reconstruction was then finalized with some of the tricks explained earlier, namely the visual noise (section 4.2), virtual cubes (section 4.3) and virtual body representation (section 4.4).



Figure 17: Comparison of the pass-through feed and the reconstruction in the final implementation

5.2 Switching From AR to VR

With the final virtual reconstruction of the room we could toggle between pass-through AR and VR. While the small 2D screenshots do not accurately represent the experience within the HMD, figure 17 shows a side-by-side comparison of the pass-through AR and the final reconstruction for reference. The supplementary materials also contain two 360-degree images, showing the reconstruction before and after the style transfer. These are still flat images, lacking the three-dimensionality of the actual reconstruction used for our implementation, but they provide a good representation, especially when viewed inside a VR HMD.

For the type of switch, there is a variety of possible options to choose from. After some internal testing, we settled on the following setup. When instantly turning on or off the virtual reconstruction to switch between AR and VR, we found that the user perceived a lot of visual movement in their periphery. This can be explained with a higher density of rods in the periphery of the human vision[9], which are sensitive to motion. The perceived motion arose from inaccuracies in the reconstruction and its alignment, for example causing an edge of a cupboard to slightly jump in one direction. By gradually fading the reconstruction's opacity over a duration of three seconds, the visual movement perceived at any one moment

could be reduced, as any offsets and other changes were spread out instead of happening at once.

At this step, we also applied the distraction task, as explained in section 4.5, and the trick of covering a large part of the user's FOV, explained in section 4.6.

5.3 Perceptual Manipulation

Once the users have been brought from AR into VR without them noticing, PMs, that only work in VR, can be applied in what is perceived to still be AR. The manipulation technique we chose for our implementation is RDW[41]. During development, we first tested curvature-based RDW, which worked fine. However, as it was easier to predict where users are going to end up, and it was sufficient for our purposes, we later used rotation-based RDW. In situations, where the users were expected to turn 180 degrees, our implementation applies small rotation gains, rotating the virtual world around the user to amplify or reduce the rotations by up to five percent. This way, the users can be guided along a predictable sawtooth trajectory instead of walking back and forth between the same two fixed positions. In figure 18, the perceived straight path is marked in green, while the red path shows the actual path taken by one of the participants from the user study.

The implementation only applies fixed relative rotation gains without matching the absolute total gain to a predefined target gain. Because of this, differences in the way users turn and walk lead to varying amounts of rotation manipulation and therefore different offsets between the reconstruction and the real environment.

While it was not part of our user study, the control over the users' rotations offered by this RDW implementation would also be sufficient to apply haptic retargeting[3].




Figure 18: Top-down view of the room, including the paths one of the participants walked relative to the virtual reconstruction (green, straight) and relative to the real room (red, sawtooth-pattern)

5.4 Switching Back From VR to AR

Our implementation also includes an option to mostly resolve the reconstruction's misalignment caused by the PM. To achieve this, the negative values of the initial rotation gains are applied in reverse order. In combination with the distraction task (section 4.5) and the paintings covering part of the user's FOV (section 4.6), this then allows switching back to AR without the user noticing it.




Figure 19: Procedure of the user study

6 User Study

In order to validate our implementation of SwitchAR, we designed the following user study. The main focus was to investigate the feasibility of an unnoticed switch from AR to VR. RDW was included in the user study as an exemplary PM to demonstrate the robustness of this switch. Our intention was to choose a PM that was widely used and kept the user inside the reconstruction for a longer duration. We wanted to test the robustness of our illusion and assess if we were able to potentially manipulate the user over a longer duration without them noticing.

6.1 Study Design

The three things we wanted to investigate in the user study were the noticeability of the AR/VR switch, the noticeability of the RDW, and the repeatability of SwitchAR. The repeatability denotes whether the switch and PM still work after the users experienced the misalignment and learned that they are being manipulated.

Overview. As visualized in figure 19, the study was designed as a sequence of three rounds of the participants using the system, each directly followed by an interview. While using the system, the participants walked back and forth between two virtual picture frames on opposing walls. Every time they reached a frame, they had ten seconds to look at a new picture and memorize its content. While they were looking at the first picture, the application switched from AR to VR. From this point on, whenever the participants turned around to walk towards the next picture, the system applied rotational gains to gradually create an offset between the virtual reconstruction and the real world. This way, the picture frames appeared next to their previous position in the real room every time, while preserving their original position in the misaligned virtual reconstruction. In rounds one and three the application then restored the original alignment, as explained in section 5.4, in order to switch back from VR to AR with the user still wearing the HMD.

The participants were supposed to position themselves in front of the wall, looking straight at a logo attached to the wall before taking off the HMD after the final picture. By doing this, any offset along the wall would be easier to spot when taking off the HMD.

Before starting with the first round, the basic concept of AR and the purpose of the study were briefly explained. As the actual intentions about the secret switch to VR could not be disclosed to the participants at this point, they were told the study was supposed to investigate memory abilities in AR. To achieve this, they supposedly were to take part in a picture memorization task in AR. They were debriefed at the end of the experiment, with the interviewer disclosing the actual purpose of the user study and the exact process with the secret switch and the subsequent RDW.

First Round. The first round started out with the participants interacting with the virtual cubes through hand tracking, as explained in section 4.3. The process following this was switching from AR to VR, applying RDW to create an offset, resolving the offset again and then switching back to AR. This way, the first round served as a test of whether the pipeline works under ideal conditions, with completely unsuspecting users. In this round, the participants looked at a total of eleven paintings.

After this first round, the first interview, which was supposed to find out two things, was conducted. First, the participants were openly asked whether they noticed anything in their environment, allowing them to report any strange or unexpected occurrence. Next, we specifically wanted to make sure that the RDW had not been noticed. The participants were asked to mark some of the paintings' positions on a 2D map of the room in top-down view. If they placed all spots on the same two positions, we could tell they did not notice the manipulated painting positions. In case they placed the spots next to each other, they were asked to clarify whether they thought the paintings were actually placed next to each other or not. We suspected some participants might just place them next to each other on the map in order to not make them intersect and cover each other.

Second Round. In the second round, the application did not resolve the offset, which also meant that it did not switch back to AR at the final picture. Therefore, we expected the offset to be noticeable when taking off the HMD. The purpose for this was to investigate participants' reactions and their explanations for suddenly being in a different position in the room. As the offset was not resolved, less paintings were required than in the previous round. Therefore, participants only looked at eight paintings.

The second interview consisted of three parts. It again started with openly asking for anything the participants noticed. Following this, the participants were asked for their explanation of how the offset happened, giving us some insights about which parts of our pipeline they might have noticed. After the participants gave their own explanation, they were told about the RDW manipulating them, causing them to end up in another spot in the room than they expected. Finally, the interviewer let the participants mark painting positions on the map again, this time asking for different paintings than in the first interview.

Third Round. The process of the third round was the same as for the first one, except for the cube interaction. The participants were redirected with the misalignment being resolved afterwards, finishing the round in AR again. However, this round was supposed to test the repeatability of the switch and the RDW, as participants had already experienced the offset in round two.

The final interview contained the open question about anything the participants noticed again, as well as the painting position placement on the map. Next, they were asked for an explanation of why there was an offset in the second round, but not in the others. Following this, a sequence of four questions was asked to determine which steps of our pipeline the participants noticed or understood. (1) Did you notice that the environment changed?, (2) Have you been in AR for the whole duration of the study?, (3) Did you notice a point in time when you changed from AR to VR? and (4) Did you notice that there was a photogrammetry/virtual

reconstruction of the room? This approach to determining whether a participant noticed the unexpected switch was inspired by a set of user studies investigating inattentive blindness with similar procedures [20, 21, 46, 48].

Wrap-Up. After the final interview, the interviewer explained to the participants what actually happened in each of the three rounds. They were told about the switching between AR and VR, the memorization task being a distraction, the RDW and the actual purpose of the experiment. They were also offered to have a look at the virtual reconstruction through the HMD again, this time telling them about some cues and defects that can be used to distinguish it from the real pass-through camera feed.

6.2 Participants

In total, 20 participants were recruited, with a mean age of 24.58 years and a standard deviation of 3.92 years. The ages ranged from a minimum of 20 to a maximum of 35, while one of the participants preferred not to disclose their age. Thirteen of the participants were female and the remaining seven male. On five-point Likert scales, the participants reported a mean VR expertise of 2.8 (standard deviation 1.01) and a mean AR expertise of 2.55 (standard deviation 1.05). The concept of RDW was known to six participants before taking part in the user study, while only three knew the concept of photogrammetry. None of the participants had used photogrammetry themselves before. The experiments had an average duration of 49 minutes, and participants were compensated with €15.

6.3 Data Analysis

For the open questions about anything the participants noticed, the answers were coded to identify those that mentioned the switch from AR to VR. For the final four questions, two of the authors discussed the responses to code the participants into four groups, based on their level of understanding of the SwitchAR pipeline. The transcribed participant responses with the codings and groupings are included in the supplementary materials for reference.

When placing the painting positions on the map, most of the participants made an explicit statement about whether they perceived the paintings to be at the same two positions every time, or whether they were actually placed on different positions on the wall. Across the three rounds, there were five instances, where a participant did not give such an explicit statement. However, in all of these cases the positions were placed in clusters with a maximum distance between them of less than four millimetres (converted to actual space on the tablet that was used). Therefore, they were considered as representing the same positions, contributing the small offsets to inaccuracies in using the trackpad.

6.4 Results

Across all three rounds, no participant reported having noticed the switch. Overall, the mean offset between the physical room and the virtual reconstruction caused by the RDW was 44.87cm, with a standard deviation of 19.41cm. Using the data of round two only, the round in which the offset was revealed on purpose, the offset was 55.46cm, with a standard deviation of 21.31cm. Despite that, only one of the twenty participants consistently reported varying painting positions.




Figure 20: Number of participants noticing the switch and varying painting positions

Noticing the switch or RDW during usage. As visualized in figure 20, no participant reported noticing the switch in any of the three rounds (e.g. "No. It's always the same"(P6), "I just followed the arrow, I didn't notice anything"(P12), "Ah, no"(P13), "No, nothing"(P19), or "Nope. Should anything be noticeable?"(P20)). Participants that mentioned noticing something mentioned only things unrelated to the switch (e.g. the time available for each painting being too short (P3) or a change in the number of paintings (P10)). The rest of the answers can be looked up in the supplementary materials.

In the first round, eighteen participants did not report that the paintings' positions varied, with P11 and P14 being the only two participants to do so. In the second and third round, P11 was the only participant reporting varying positions, making them the only participant to do so across all three rounds.

Noticing the offset at the end of round two. The mean offset revealed on purpose when the participants took off the HMD was about 55cm, with a standard deviation of about 21cm. 16 of the participants did not seem to notice or react to this offset. In that case, they were instructed to look at the logo on the wall again, first through the HMD and then without the HMD, giving them another chance to notice the offset. Following this, three of them still did not notice it. Among them, two had relatively small offsets, about 17cm for P13 and 34cm for P19. However, P3 had an offset of roughly 84cm, which they still did not notice. This might have been enabled through change blindness, based on the short interruption caused by the HMD blocking the view when taking it off.

Post-rationalization. We used the question after each exposure ("Did you notice anything?") as the metric to measure if participants were able to notice the switch to VR. In our final four questions, we gradually revealed more about the underlying mechanism, e.g. by asking about VR, and therefore consider the responses as a post-rationalization of participants. We were interested if, once we revealed parts of our concept, participants would be able to deduce in hindsight how and where the switch could have happened.

Based on these final questions, the participants were coded into four groups. Group 1 mentioned nothing of the things that were




Figure 21: Post-rationalization: Level of understanding of SwitchAR's pipeline

asked for, showing the lowest level of understanding and noticing anything (13 participants). Group 2 only mentioned the virtual reconstruction once it was brought up in our question (3 participants). Group 3 mentioned that, in hindsight, they were probably not using AR all the time, but they did not identify the moment of the switch (2 participants). Group 4 was able to deduce the probable moment of the switch but did not notice when it happened (2 participants).

Thirteen participants did not gain any insights (e.g., "Nope. I thought that was the room the whole time." (P3)) into the SwitchAR pipeline, placing them in group 1 (figure 21). The three participants in group 2 (P5, P8, P19) only mentioned that there could have been a virtual reconstruction, e.g. "That was my theory about the shift. But I thought it wouldn't look that realistic."(P5). They did not report anything related to the switch or the reconstruction at the previous questions. When asked if they were constantly in AR, P7 and P15 mentioned they could have not been using AR exclusively, e.g. "I'm not sure about round two. Because if I had seen the whole real world, I would have seen the logo in the right place. And I didn't."(P15). However, they were not able to identify the moment of the switch, placing them in group 3. Finally, the two participants in group 4 (P12, P17) were able to correctly deduce the moment in which the switch must have happened, e.g. "Mh, probably when I looked at the pictures. And then focused on... the picture frames."(P17). However, this was only done after we explicitly asked about a switch to VR, and when asked before if they were constantly in AR, P17 responded with "Mmm, yes". Overall, all participants in each group were not able to notice the switch during the experiment.

7 Discussion

In the following, we discuss our findings along three main points: (1) *Feasibility*, (2) *Generalizability* and (3) *Misuse*.

Feasibility. As no participant was able to notice the actual switch in any of the three rounds, our experiment demonstrated the feasibility of SwitchAR to transition users from an AR environment into a VR environment without them noticing the change and still believing they were interacting in the real world (AR). Additionally, with

only one participant perceiving the rotational gains, it is possible to apply VPPMs like RDW in an AR context using SwitchAR.

Following the reveal of the RDW after the second round, the third round allowed us to investigate the repeatability and robustness of the system. For 18 of the 20 participants, the knowledge of the RDW did not reduce the effectiveness of SwitchAR or allow them to better understand the system. This is in line with the findings of Simons[47] in a follow-up project of their original inattentional blindness experiment[48]. They found that the expectation of a specific unexpected occurrence, in our case the RDW, did not improve the chance to notice other unexpected occurrences, like in our case the switch between AR and VR.

After the offset at the end of round two, the participants were told about the RDW, and the third round then ended without an apparent offset. When asked about this discrepancy afterwards, only three participants (P7, P11, P19) stated that the RDW might have led them along a different path. The other 17 participants either had no explanation, thought that the reason was a difference in their own behaviour, or even thought that there was no RDW in the third round. However, RDW was applied with the same magnitude of rotation gains in all three rounds. Two participants (P12, P15), which in the interviews deduced that there might have been VR in hindsight, explicitly said that there was no VR used in the third round any more. P15 even specifically paid attention to potential changes in the environment in the third round, but still did not notice the switch or the reconstruction. Based on this, it seems to be feasible to apply the concept of SwitchAR in applications that are being used repeatedly by the same user.

Generalizability. We consider SwitchAR to be a fundamental enabler for perceptual manipulations in Augmented Reality. Since we are able to unnoticeably transition users into a fully virtual environment, we argue that most PMs that are currently used in VR would also work in AR using SwitchAR.

Admittedly, PMs such as haptic retargeting[3] and manipulating the perception of weight[42, 45] will require a more precise tracking of the user's hands and arms. However, once this technological progress happens, there would be no conceptual problems using the pipeline we present in this work to enable these types of hand-arm redirections in AR. Therefore, the main focus of our research was to demonstrate that we are able to switch from AR to VR without the user noticing, and the following RDW was mainly an example for demonstration and evaluation purposes. This unlocks new application scenarios for these PMs, transferring their benefits from VR applications to AR applications. Aside from the known use cases of PMs in VR, the different context of AR might also enable new, unknown use cases for PMs with unique benefits.

The presented prototype was restricted to a highly controlled indoor environment, based on the goal of demonstrating the general feasibility of SwitchAR as a first step. The concept of SwitchAR can also be extended to more complex and dynamic environments, though. A more demanding environment would then require a more complex reconstruction as well, for example based on neural radiance fields[32] or 3D Gaussian Splatting[18].

Misuse. We are aware that the ability to secretly transition a user from AR to VR enables a large set of potential misuse and can be considered a new security (and even safety) vulnerability of video

pass-through AR HMDs. The same scenarios identified for VR by Tseng et al.[57] are still valid for the perceived AR created by our pipeline. Attackers could use RDW to let users fall down the stairs or make them harm themselves with dangerous objects. We suspect the potential for harm could even be larger in perceived AR, than it is in VR applications, in which the users know they are perceiving an artificial environment. Users might be less careful and more daring, because they think they are still able to see the real world around them, supposedly making it easy to recognize potentially dangerous situations.

We believe that one of the reasons SwitchAR has strong potential for misuse is its ability to undermine the "device-gap", a concept introduced by Slater et al. [50]. When discussing potential dangers of realism, Slater et al. argue that *"the fact of putting on the HMD itself demarcates reality from virtual reality—so that unless participants are induced to somehow forget that they are wearing the HMD, they will not believe that the virtual scenario is a real one."* While SwitchAR does not make users forget that they are wearing an HMD, it does hide the fact that at one point during the experience the users are not perceiving the real environment anymore, but a fully virtual reconstruction.

This creates a scenario in which participants are likely not able to distinguish what of the displayed content is currently really happening in their environment. As discussed in prior work[4, 39], SwitchAR is another example that raises the question about "Perceptual Integrity" and "Perceptual Human Rights". Extended Reality (XR) has the ability to fully control what users perceive in their physical environment and how they perceive it, and it is still not clear to what degree this has to be restricted through legal mechanisms. Nevertheless, in this work our goal was to demonstrate how known concepts such as change and inattention blindness can be leveraged as a mechanism to enable PMs in AR. We hope that raising awareness of these capabilities will encourage the community to explore new positive application scenarios, and the legislation together with the companies distributing these devices to start implementing more safety mechanisms to protect users and preserve their perceptual integrity. We can imagine a future in which XR applications will have similar permission agreements as current smartphone apps. Users would therefore have to explicitly agree to being manipulated, resulting in a scenario where you would have consensual PMs.

The specific manipulation used in our SwitchAR prototype could be potentially prevented by a system similar to Arya, proposed by Lebeck et al.[24]. The policy of only allowing AR applications to occupy a certain fraction of the HMD's displays could inhibit an unnoticed switch from AR to VR, as the physical world could no longer be fully occluded by virtual content. However, such intervention techniques would follow a similar pattern as we currently see in the field of security, where attackers adapt to newly developed security mechanisms, which again, as a result, will adapt to the new attack, triggering the cycle again. Future research will have to focus on more fundamental questions about what it is that we are protecting. Bonnail et al.[4] discuss the concept of "perceptual integrity," which they explain as being the "agency and trust in what we are perceiving to be unaltered by technology"[4]. At its core, SwitchAR is undermining this concept of perceptual integrity,

and future safety mechanisms could have "upholding the user's perceptual integrity" as a core goal throughout the experience.




Figure 22: Basic transformations of the model manipulation space (the blue room represents the virtual model and the grey/brown one the physical environment)

8 Design Space for Future Applications

To present the possibilities arising from a technique like SwitchAR we present a resulting design space grounded in the premise that many PMs can be seen as a transformation of the virtual environment (see Figure 22 and video figure at the timestamp 04:32). In the following we will present the type of experiences that can be created in perceived AR when the virtual environment is translated, rotated or scaled relative to the physical environment of the user. Additionally, these three types of transformations can be arbitrarily combined to achieve more manipulations. As mentioned in section 3.3, such model transformations introduce the side effect of creating a misalignment between the physical environment and the virtual reconstruction.

We implemented each of these transformations in our prototype to get an impression of how they function. We have anecdotal insights from self-experimentation, but further research is required to fully understand how they might impact user behaviour during and after exposure. We consider the following as a set of potential future research directions we encourage the community to explore.

Model Translation. Manipulating the translational relationship between the physical and virtual world enables, e.g., the ability to amplify or reduce movements made by a user in the physical

environment. For each step the user takes, the virtual world can be moved in the opposite or same direction as the movement. The user would then perceive that each step taken in the environment is either larger or smaller, creating the illusion of moving faster or slower. This type of manipulation was presented in prior work, such as the seven league boots technique in VR[1, 13], while SwitchAR allows to transfer it to AR experiences.

Model Rotation. Rotating the virtual environment could be applied to manipulate the user's sense of balance. Tilting the virtual floor could manipulate the user's visual perceptions regarding balance, potentially creating conflicts with the vestibular perceptions and conveying an incorrect sense of balance. This could be applied to user studies investigating balance and posture[15, 28, 56, 58], where the ability to use perceived AR instead of VR might elicit more natural behaviour from the participants. One potential application scenario could be to tilt the virtual floor in a way that the user gets the perception of walking uphill or downhill to impact the perceived effort required for walking in AR.

When testing the floor tilt with colleagues in our lab, they relied on the floor shown in the HMD to plan their steps. At a point in the lab where the tilted virtual floor was located above the physical floor, one of them tried to place their foot in midair. As the physical floor was lower than expected, they stumbled and continued to walk very cautiously afterwards. Additionally, we had the impression that even small degrees of floor tilting could be perceived easier than model scaling or translations. This might be connected to the vestibular system being more sensitive to conflicts, but future work can explore what exact values are perceivable by the participants and if they are different from doing these types of manipulations in a purely virtual environment.

Model Scaling. Finally, scaling the virtual model allows, e.g., to manipulate the perceived body size of the user. By uniformly scaling the model up or down on all three axes, the user's perceived body size relative to the environment would decrease or increase accordingly. As shown in prior work, a larger body height in VR can improve the user's confidence[26, 29]. SwitchAR could be applied to investigate if this effect is stronger when the user believes they are using AR.

When experimenting with ourselves and colleagues in our prototype, e.g. uniformly shrinking the environment by 1% per second, we were surprised that the ongoing shrinking or growing is not directly perceived. Only once a certain threshold is reached, e.g., you are smaller than the door handle, users got the impression that something was wrong. When leaving the reconstruction after interacting in a different size, the perception of the physical environment was changed. Reaching for things or walking around started to feel different and in our future work we aim to explore how we could measure this "aftereffect" and how it could impact the users' motor control abilities.

9 Research Platform and Aftereffects

Research Platform. One area of application we are envisioning for SwitchAR is as a research platform. What we presented in this paper is that we are able to create an easy to reproduce study design that can make a participant believe they are seeing the real world despite

interacting inside a reconstruction we have full control over. This illusion is so robust that we were able to apply RDW without the participants noticing. None of our 20 participants noticed that they interacted inside a reconstruction. This means we could also use this setup to ask new questions in the field of psychology to better understand human cognition and perception. SwitchAR is able to create illusions that were (to the best of our knowledge) not possible or extremely difficult to create beforehand. Using SwitchAR, we can create scenarios that are physically impossible, but the user believes that they are happening. We can start asking questions such as: What happens if during the VR-Phase we start "shrinking" the user with every step taken? At what point does the user notice the manipulation? How does this manipulation impact the user's physical actions, their trajectory or gait? While shrinking a user in VR is something we are already able to do in a lab, the user always knows that they are entering a digital environment that is able to create impossible environments. SwitchAR is able to research impossible environments while the user believes they are interacting with the real world, not knowing we control the full environment. This can be seen as a VR environment in which a user is experiencing "maximum" presence. If we define presence as "the subjective experience of being in one place", we argue that once the user is transitioned from AR to VR, they are having the subjective experience of being in exactly that virtual environment.

Aftereffects. Apart from applying PMs with the goal of achieving an effect during usage, this "maximum" presence might also allow SwitchAR to elicit strong aftereffects. We encountered an indicator for this when testing our prototype. One of the developers was exploring the virtual room model, while the HMD was applying a PM that was gradually lowering their eye height. The developer was aware of the virtual model and the PM, but the eye height changes were too slow to actually perceive them. After taking off the HMD, the developer experienced hesitation when interacting with the physical world. When grabbing the door handle, the developer saw and felt that they were moving their arm correctly to reach the handle. However, at the same time they had the feeling that the door handle's position was wrong, that it should have been higher. Despite the short exposure time of about 5 minutes and the developer being aware of the virtual model and the PM, the system could still affect their proprioception. For users that constantly believe they are interacting with the physical world in AR, SwitchAR might be able to cause stronger aftereffects, especially with longer exposure.

10 Limitations and Future Work

The concept of SwitchAR can only apply the PMs in what is perceived as AR, but is technically VR, not AR. However, we argue that this distinction is not actually relevant for the user experience. The system starts in AR, ends in AR and only switches to VR temporarily in between, whenever there is a need for PMs. As the user does not notice these switches and therefore equally perceives the whole usage as AR, their experience does not differ from a system running in actual AR all the time.

The restriction to pass-through AR is a technical limitation of our implementation of SwitchAR. The reason for this restriction is the reliance on being able to fully occlude the real world in the

complete FOV of the user. This is needed for replacing the real environment with a manipulable virtual reconstruction. As soon as technological improvements allow optical see-through AR HMDs to cover the whole FOV with pixels that can block the light from the physical environment, SwitchAR is able to work on the same principles as described in this paper on optical AR HMDs.

As mentioned in the discussion, our prototype was restricted to a highly controlled indoor environment, based on the goal of demonstrating the general feasibility of SwitchAR as a first step. Therefore, future work can extend the concept of SwitchAR to more complex and dynamic environments, potentially by utilizing neural radiance fields[26] or 3D Gaussian Splatting[16] instead of photogrammetry.

In theory, the switch would also be possible without any distraction task or perceptual tricks. Provided a virtual reconstruction, which is very similar to the real environment with very accurate alignment, the amount of change and visual movement happening during the switch would be very low. This would make it possible to potentially switch from AR to VR at any time without the users noticing it, even if they are currently looking at the surrounding environment. The more defects in the reconstruction and inaccuracies in the alignment there are, the more distraction and perceptual tricks are needed to hide the switch, though. The set of tricks we applied in our implementation seemed to achieve the desired effect in the user study. We argue that until technology provides a perfect reconstruction, our demonstration of SwitchAR provides an easy to reproduce system allowing researchers to explore its concept already. However, we do not know how much each of the tricks contributed, to which degree each of them had to be applied and if some of them could even have been omitted completely. Quantifying the impact of each of these factors is something that future work can focus on.

Another option for future work is to run user studies comparing RDW in VR with RDW in AR. They could investigate, e.g., detection thresholds, workload and fatigue to determine if the implementation with SwitchAR differs from traditional VR implementations.

Finally, using AR HMDs that provide access to their cameras to the developers⁶ could enable the application of real-time reconstruction techniques. This way, SwitchAR would not have to rely on a controlled environment with a pre-built reconstruction any more. Instead, the system could create a custom reconstruction of the user's environment at the beginning of the session.

11 Conclusion

In this work, we introduced SwitchAR, a PM for video pass-through AR that leverages change blindness and inattention blindness to imperceptibly switch between the camera stream and a 3D reconstruction of the real environment, enabling VR redirection techniques in what users still perceive as AR. We presented its concept, based on (1) Reconstruction, (2) Switch from AR to the reconstruction, (3) PM and (4) Switch back to AR, as well as a prototype implementation with RDW as an exemplary PM. We then evaluated this implementation in a user study (n=20) with three rounds of switching from AR to VR followed by redirections. The results

confirmed the feasibility of SwitchAR, with no participant noticing the switch, and only one consistently noticing the RDW. Even after revealing the manipulation in the second round, no participant noticed the switch in the third round, demonstrating the repeatability and robustness of the system. This unlocks the transfer of both the benefits and the dangers of VR PMs to AR. Therefore, SwitchAR is a foundational approach to enable perceptual manipulations in Augmented Reality.

Acknowledgments

This research work has been co-funded by the National Research Center for Applied Cybersecurity ATHENE, and the LOEWE initiative (Hesse, Germany) within the emergenCITY center.

References

- [1] Parastoo Abtahi, Mar Gonzalez-Franco, Eyal Ofek, and Anthony Steed. 2019. I'm a Giant: Walking in Large Virtual Environments at High Speed Gains. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300752
- [2] Benjamin Avery, Christian Sandor, and Bruce H. Thomas. 2009. Improving Spatial Perception for Augmented Reality X-Ray Vision. In *2009 IEEE Virtual Reality Conference*. 79–82. doi:10.1109/VR.2009.4811002
- [3] Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. 2016. Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 1968–1979. doi:10.1145/2858036.2858226
- [4] Elise Bonnal, Wen-Jie Tseng, Mark McGill, Eric Lecolinet, Samuel Huron, and Jan Gugenheimer. 2023. Memory Manipulations in Extended Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 875, 20 pages. doi:10.1145/3544548.3580988
- [5] Peter Casey, Ibrahim Baggili, and Ananya Yarramreddy. 2021. Immersive Virtual Reality Attacks and the Human Joystick. *IEEE Transactions on Dependable and Secure Computing* 18, 2 (2021), 550–562. doi:10.1109/TDSC.2019.2907942
- [6] Kirsten Cater, Alan Chalmers, and Colin Dalton. 2003. Varying rendering fidelity by exploiting human change blindness. In *Proceedings of the 1st International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia* (Melbourne, Australia) (GRAPHITE '03). Association for Computing Machinery, New York, NY, USA, 39–46. doi:10.1145/604471.604483
- [7] Kaiming Cheng, Jeffery F. Tian, Tadayoshi Kohno, and Franziska Roesner. 2023. Exploring User Reactions and Mental Models Towards Perceptual Manipulation Attacks in Mixed Reality. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 911–928. <https://www.usenix.org/conference/usenixsecurity23/presentation/cheng-kaiming>
- [8] Aldrich Clarence, Jarrod Knibbe, Maxime Cordeil, and Michael Wybrow. 2024. Stacked Retargeting: Combining Redirected Walking and Hand Redirection to Expand Haptic Retargeting's Coverage. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 684, 13 pages. doi:10.1145/3613904.3642228
- [9] Christine A. Curcio, Kenneth R. Sloan, Robert E. Kalina, and Anita E. Hendrickson. 1990. Human photoreceptor topography. *Journal of Comparative Neurology* 292, 4 (1990), 497–523. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.902920402> doi:10.1002/cne.902920402
- [10] Yifan Ding, Connor Hulst, Rishi Raja, and Daniel Simons. 2023. Similarity of an unexpected object to the attended and ignored objects affects noticing in a sustained inattention blindness task. *Attention, Perception, & Psychophysics* 85 (10 2023). doi:10.3758/s13414-023-02794-2
- [11] Leah Findlater, Karyn Moffatt, Joanna McGrenere, and Jessica Dawson. 2009. Ephemeral adaptation: the use of gradual onset to improve menu selection performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 1655–1664. doi:10.1145/1518701.1518956
- [12] Matt Gottsacker, Nahal Norouzi, Kangsoo Kim, Gerd Bruder, and Greg Welch. 2021. Diegetic Representations for Seamless Cross-Reality Interruptions. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 310–319. doi:10.1109/ISMAR52148.2021.00047

⁶<https://developer.apple.com/documentation/visionos/accessing-the-main-camera>, last visited: 08.12.2024

- [13] Victoria Interrante, Brian Ries, and Lee Anderson. 2007. Seven League Boots: A New Metaphor for Augmented Locomotion through Moderately Large Scale Immersive Virtual Environments. In *2007 IEEE Symposium on 3D User Interfaces*. doi:10.1109/3DUI.2007.340791
- [14] Akira Ishii, Ipezi Suzuki, Shinji Sakamoto, Keita Kanai, Kazuki Takazawa, Hiraku Doi, and Yoichi Ochiai. 2016. Optical Marionette: Graphical Manipulation of Human's Walking Direction. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 705–716. doi:10.1145/2984511.2984545
- [15] Jeffrey Jacobson, Mark S. Redfern, Joseph M. Furman, Susan L. Whitney, Patrick J. Sparto, Jeffrey B. Wilson, and Larry F. Hodges. 2001. Balance NAVE: a virtual reality facility for research and rehabilitation of balance disorders. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (Baniff, Alberta, Canada) (VRST '01). Association for Computing Machinery, New York, NY, USA, 103–109. doi:10.1145/505008.505027
- [16] Denis Kalkofen, Eduardo Veas, Stefanie Zollmann, Markus Steinberger, and Dieter Schmalstieg. 2013. Adaptive ghosted views for Augmented Reality. In *2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 1–9. doi:10.1109/ISMAR.2013.6671758
- [17] Mohamed Kari, Reinhard Schütte, and Raj Sodhi. 2023. Scene Responsiveness for Visuotactile Illusions in Mixed Reality. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 84, 15 pages. doi:10.1145/3586183.3606825
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. arXiv:2308.04079 [cs.GR] <https://arxiv.org/abs/2308.04079>
- [19] Hanseob Kim, TaeHyung Kim, MyungHo Lee, Gerard Jounghyun Kim, and Jae-In Hwang. 2020. Don't Bother Me: How to Handle Content-Irrelevant Objects in Handheld Augmented Reality. In *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology* (Virtual Event, Canada) (VRST '20). Association for Computing Machinery, New York, NY, USA, Article 32, 5 pages. doi:10.1145/3385956.3418948
- [20] Carina Kreitz, Robert Schnuerch, Henning Gibbons, and Daniel Memmert. 2015. Some See It, Some Don't: Exploring the Relation between Inattentional Blindness and Personality Factors. *PLOS ONE* 10, 5 (05 2015), 1–16. doi:10.1371/journal.pone.0128158
- [21] Hannah Korrel Kristen Pammer and Jason Bell. 2014. Visual distraction increases the detection of an unexpected object in inattentional blindness. *Visual Cognition* 22, 9–10 (2014), 1173–1183. arXiv:<https://doi.org/10.1080/13506285.2014.987859>
- [22] Christian Kunert, Tobias Schwandt, and Wolfgang Broll. 2019. An Efficient Diminished Reality Approach Using Real-Time Surface Reconstruction. In *2019 International Conference on Cyberworlds (CW)*, 9–16. doi:10.1109/CW.2019.00010
- [23] Kiron Lebeck, Tadayoshi Kohno, and Franziska Roesner. 2016. How to Safely Augment Reality: Challenges and Directions (*HotMobile '16*). Association for Computing Machinery, New York, NY, USA, 45–50. doi:10.1145/2873587.2873595
- [24] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2017. Securing Augmented Reality Output. In *2017 IEEE Symposium on Security and Privacy (SP)*, 320–337. doi:10.1109/SP.2017.13
- [25] JangHyeon Lee and Lawrence H. Kim. 2025. DiminishAR: Diminishing Visual Distractions via Holographic AR Displays. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 30, 16 pages. doi:10.1145/3706598.3713415
- [26] Grace Y. S. Leung, Adrian K. T. Ng, and Henry Y. K. Lau. 2021. Effect of Height Perception on State Self-Esteem and Cognitive Performance in Virtual Reality. In *Engineering Psychology and Cognitive Ergonomics*, Don Harris and Wen-Chin Li (Eds.). Springer International Publishing, Cham, 172–184.
- [27] David Lindlbauer and Andy D. Wilson. 2018. Remixed Reality: Manipulating Space and Time in Augmented Reality (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3173703
- [28] Anat V. Lubetzky and Bryan D. Hujsak. 2019. A virtual reality head stability test for patients with vestibular dysfunction. *Journal of Vestibular Research* 28, 5–6 (2019), 393–400. arXiv:<https://journals.sagepub.com/doi/pdf/10.3233/VES-190650> doi:10.3233/VES-190650
- [29] Anna-Leena Macey, Simo Järvelä, Daniel Fernández Galeote, and Juho Hamari. 2023. Feeling Small or Standing Tall? Height Manipulation Affects Speech Anxiety and Arousal in Virtual Reality. *Cyberpsychology, Behavior, and Social Networking* 26, 4 (2023), 246–254. arXiv:<https://doi.org/10.1089/cyber.2022.0251> doi:10.1089/cyber.2022.0251 PMID: 36989502.
- [30] Arien Mack and Irvin Rock. 1998. Inattentional Blindness: Perception Without Attention. *Visual Attention* 8 (01 1998).
- [31] Steve Mann. 1999. Mediated Reality. *Linux J.* 1999, 59es (March 1999), 5–es.
- [32] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. 2021. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. arXiv:2008.02268 [cs.CV] <https://arxiv.org/abs/2008.02268>
- [33] Sebastian Marwecki, Andrew D. Wilson, Eyal Ofek, Mar Gonzalez Franco, and Christian Holz. 2019. Mise-Unseen: Using Eye Tracking to Hide Virtual Reality Scene Changes in Plain Sight. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 777–789. doi:10.1145/3332165.3347919
- [34] Daniel Medeiros, Graham Wilson, Mark McGill, and Stephen Anthony Brewster. 2023. The Benefits of Passive Haptics and Perceptual Manipulation for Extended Reality Interactions in Constrained Passenger Spaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 232, 19 pages. doi:10.1145/3544548.3581079
- [35] Paul Milgram and Fumio Kishino. 1994. A Taxonomy of Mixed Reality Visual Displays. *IEICE Trans. Information Systems* vol. E77-D, no. 12 (12 1994), 1321–1329.
- [36] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. 1994. Augmented reality: A class of displays on the reality-virtuality continuum. *Telematics and Telepresence Technologies* 2351 (01 1994). doi:10.1117/12.197321
- [37] Roberto A. Montano Murillo, Sriram Subramanian, and Diego Martinez Plasencia. 2017. Erg-O: Ergonomic Optimization of Immersive Virtual Environments (UIST '17). Association for Computing Machinery, New York, NY, USA, 759–771. doi:10.1145/3126594.3126605
- [38] Steven B. Most, Daniel J. Simons, Brian J. Scholl, Rachel Jimenez, Erin Clifford, and Christopher F. Chabris. 2001. How not to be Seen: The Contribution of Similarity and Selective Ignoring to Sustained Inattentional Blindness. *Psychological Science* 12, 1 (2001), 9–17. arXiv:<https://doi.org/10.1111/1467-9280.00303> doi:10.1111/1467-9280.00303 PMID: 11294235.
- [39] Joseph O'Hagan, Jan Gugenheimer, Florian Mathis, Jolie Bonner, Richard Jones, and Mark McGill. 2024. A Viewpoint on the Societal Impact of Everyday Augmented Reality and the Need for Perceptual Human Rights. *IEEE Security and Privacy* 22, 1 (jan 2024), 64–68. doi:10.1109/MSEC.2023.3333988
- [40] Fabian Pointecker, Judith Friedl-Knirsch, Hans-Christian Jetter, and Christoph Anthes. 2024. From Real to Virtual: Exploring Replica-Enhanced Environment Transitions along the Reality-Virtuality Continuum. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 799, 13 pages. doi:10.1145/3613904.3642844
- [41] Sharif Razzaque. 2005. *Redirected walking*. The University of North Carolina at Chapel Hill.
- [42] Michael Rietzler, Florian Geiselhart, Jan Gugenheimer, and Enrico Rukzio. 2018. Breaking the Tracking: Enabling Weight Perception using Perceivable Tracking Offsets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173702
- [43] Franziska Roesner and Tadayoshi Kohno. 2021. Security and Privacy for Augmented Reality: Our 10-Year Retrospective. <https://api.semanticscholar.org/CorpusID:236167221>
- [44] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (apr 2014), 88–96. doi:10.1145/2580723.2580730
- [45] Majed Samad, Elia Gatti, Anne Hermes, Hrvoje Benko, and Cesare Parise. 2019. Pseudo-Haptic Weight: Changing the Perceived Weight of Virtual Objects By Manipulating Control-Display Ratio. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300550
- [46] Benjamin Schöne, Rebecca Sophia Sylvester, Elise Leila Radtke, and Thomas Gruber. 2021. Sustained inattentional blindness in virtual reality and under conventional laboratory conditions. *Virtual Real.* 25, 1 (March 2021), 209–216. doi:10.1007/s10055-020-00450-w
- [47] Daniel J Simons. 2010. Monkeying around with the Gorillas in Our Midst: Familiarity with an Inattentional-Blindness Task Does Not Improve the Detection of Unexpected Events. *i-Perception* 1, 1 (2010), 3–6. arXiv:<https://doi.org/10.1068/i0386> doi:10.1068/i0386 PMID: 23397479.
- [48] Daniel J Simons and Christopher F Chabris. 1999. Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception* 28, 9 (1999), 1059–1074. arXiv:<https://doi.org/10.1068/p281059> doi:10.1068/p281059 PMID: 10694957.
- [49] Daniel J. Simons and Daniel T. Levin. 1997. Change blindness. *Trends in Cognitive Sciences* 1, 7 (1997), 261–267. doi:10.1016/S1364-6613(97)01080-2
- [50] Mel Slater, Cristina Gonzalez-Liencre, Patrick Haggard, Charlotte Vinkers, Rebecca Gregory-Clarke, Steve Jelley, Zillah Watson, Graham Breen, Raz Schwarz, William Steptoe, Dalila Szostak, Shivashankar Halan, Deborah Fox, and Jeremy Silver. 2020. The Ethics of Realism in Virtual and Augmented Reality. *Frontiers in Virtual Reality* 1 (2020). doi:10.3389/frvir.2020.00001
- [51] Mel Slater, Anthony Steed, John McCarthy, and Francesco Marinelli. 1998. The Virtual Ante-Room: Assessing Presence Through Expectation and Surprise. In *Virtual Environments '98, Eurographics Workshop Proceedings Series*.

- [52] Frank Steinicke, Gerd Bruder, Klaus Hinrichs, Markus Lappe, Brian Ries, and Victoria Interrante. 2009. Transitional environments enhance distance perception in immersive virtual reality systems. In *Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization* (Chania, Crete, Greece) (APGV '09). Association for Computing Machinery, New York, NY, USA, 19–26. doi:10.1145/1620993.1620998
- [53] Frank Steinicke, Gerd Bruder, Klaus Hinrichs, Anthony Steed, and Alexander L. Gerlach. 2009. Does a Gradual Transition to the Virtual World increase Presence?. In *2009 IEEE Virtual Reality Conference*. 203–210. doi:10.1109/VR.2009.4811024
- [54] Evan Suma Rosenberg, Seth Clark, David Krum, Samantha Finkelstein, Mark Bolas, and Zachary Warte. 2011. Leveraging change blindness for redirection in virtual environments. *Proceedings - IEEE Virtual Reality*, 159 – 166. doi:10.1109/VR.2011.5759455
- [55] Lin Sun, Liam S Harpur, Matthew E Broomhall, and Paul R Bastide. 2020. Switching realities for better task efficiency. <https://patents.google.com/patent/US10783711B2/en> US Patent 10,783,711.
- [56] Tsubasa Tashiro, Noriaki Maeda, Takeru Abekura, Rami Mizuta, Yui Terao, Satoshi Arima, Satoshi Onoue, and Yukio Urabe. 2024. Adaptation of Postural Sway in a Standing Position during Tilted Video Viewing Using Virtual Reality: A Comparison between Younger and Older Adults. *Sensors* 24, 9 (2024). doi:10.3390/s24092718
- [57] Wen-Jie Tseng, Elise Bonnal, Mark McGill, Mohamed Khamis, Eric Lecolinet, Samuel Huron, and Jan Gugenheimer. 2022. The Dark Side of Perceptual Manipulations in Virtual Reality. In *CHI Conference on Human Factors in Computing Systems* (CHI '22). ACM. doi:10.1145/3491102.3517728
- [58] Yukio Urabe, Kazuki Fukui, Keita Harada, Tsubasa Tashiro, Makoto Komiya, and Noriaki Maeda. 2022. The Application of Balance Exercise Using Virtual Reality for Rehabilitation. *Healthcare* 10, 4 (2022). doi:10.3390/healthcare10040680
- [59] Andrew D. Wilson and Shane Williams. 2018. Autopager: exploiting change blindness for gaze-assisted reading. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications* (Warsaw, Poland) (ETRA '18). Association for Computing Machinery, New York, NY, USA, Article 46, 5 pages. doi:10.1145/3204493.3204556
- [60] Katherine Wood and Daniel J. Simons. 2017. The role of similarity in inattentional blindness: Selective enhancement, selective suppression, or both?†. *Visual Cognition* 25, 9-10 (2017), 972–980. arXiv:<https://doi.org/10.1080/13506285.2017.1365791> doi:10.1080/13506285.2017.1365791
- [61] Steven Yantis and John Jonides. 1984. Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of experimental psychology. Human perception and performance* 10 (10 1984), 601–21. doi:10.1037/0096-1523.10.5.601
- [62] Steven Yantis and John Jonides. 1990. Abrupt Visual Onsets and Selective Attention: Voluntary Versus Automatic Allocation. *Journal of experimental psychology: Human perception and performance* 16 (02 1990), 121–34. doi:10.1037/0096-1523.16.1.121
- [63] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. 2019. Photorealistic Style Transfer via Wavelet Transforms. arXiv:1903.09760 [cs.CV] <https://arxiv.org/abs/1903.09760>