Combating VR Sickness through Subtle Dynamic Field-Of-View Modification

Ajoy S Fernandes*
Columbia University

Steven K. Feiner†
Columbia University

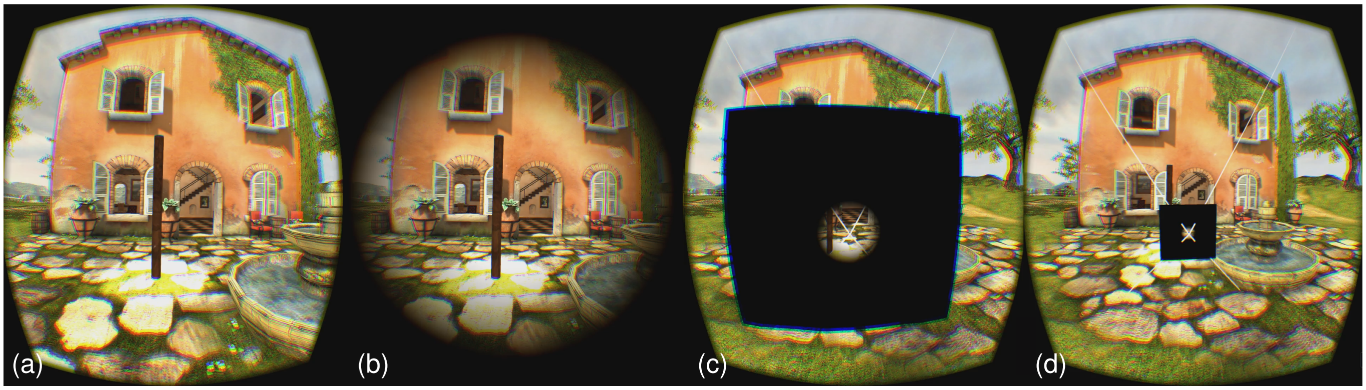

Figure 1: (a) View of VE from one eye with no FOV modification. Study waypoint (lit vertical post) appears at center. (b) Same view restricted by 90° FOV soft-edged circular cutout. While clearly visible at small size shown here, slow restriction to this FOV was imperceptible to majority of study participants when viewed in HWD. (c) Geometry of FOV restrictor and view frustum used to create subfigure a. Cutout is opened wide enough to not affect FOV. (d) Geometry of FOV restrictor and view frustum used to create subfigure b, viewed from same location as in subfigure c. Note that FOV restrictor is scaled about its center in its plane and does not translate relative to center of projection.

Abstract

Virtual Reality (VR) sickness can cause intense discomfort, shorten the duration of a VR experience, and create an aversion to further use of VR. High-quality tracking systems can minimize the mismatch between a user's visual perception of the virtual environment (VE) and the response of their vestibular system, diminishing VR sickness for moving users. However, this does not help users who do not or cannot move physically the way they move virtually, because of preference or physical limitations such as a disability.

It has been noted that decreasing field of view (FOV) tends to decrease VR sickness, though at the expense of sense of presence. To address this tradeoff, we explore the effect of dynamically, yet subtly, changing a physically stationary person's FOV in response to visually perceived motion as they virtually traverse a VE. We report the results of a two-session, multi-day study with 30 participants. Each participant was seated in a stationary chair, wearing a stereoscopic head-worn display, and used control and FOV-modifying conditions in the same VE. Our data suggests that by strategically and automatically manipulating FOV during a VR session, we can reduce the degree of VR sickness perceived by participants and help them adapt to VR, without decreasing their subjective level of presence, and minimizing their awareness of the intervention.

Keywords: VR sickness; cybersickness; virtual reality; head-worn display; field of view.

Index Terms: H.1.2 [Models and Principles]: User/Machine Systems—-Human Factors; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—-Artificial, augmented, and virtual realities; H.5.2 [Information Interfaces and Presentation]: User Interface—-Ergonomics

*e-mail: asf2168@columbia.edu
†e-mail: feiner@cs.columbia.edu

IEEE Symposium on 3D User Interfaces 2016
19–23 March, Greenville, SC, USA
978-1-5090-0842-1/16/$31.00 ©2016 IEEE

1 Introduction

Virtual reality (VR) head-worn displays (HWDs) are on their way to becoming mass-market products. However, one potential barrier to adoption is *VR sickness*, which can cause symptoms similar to those of motion sickness [12]. These symptoms include headaches, stomach awareness, nausea, vomiting, pallor, sweating, fatigue, drowsiness, and disorientation [11], and are explicitly listed in the "Health and Safety Warnings" accompanying current VR platforms (e.g., [15]). In earlier work on vehicle simulators using a variety of display technologies, researchers noted that people develop a tolerance to the related experience of *simulator sickness* [9] over multiple sessions, and that by having users undergo an adaptation program, such as through increasing exposure time by session, users might more easily adapt to the experience [14]. However, given the unpleasantness of some of the symptoms, having a bad first experience can potentially deter users from trying a system again.

High-precision low-latency tracking, high–frame-rate rendering, and short-persistence displays have sometimes been claimed to eliminate or drastically reduce VR sickness, insofar as they can minimize the mismatch between a user's visual perception of the virtual environment (VE) and the response of their vestibular system [5, 18, 11, 33]. While this can help moving users, it does not address users who do not or cannot move physically the same way they move virtually. This can be the case when the user's tracked environment is significantly smaller than the VE they wish to explore, when the user prefers to remain relatively stationary physically when moving virtually, or when the user is simply unable to move physically because of a disability. In these scenarios in which actual physical and intended virtual motion are significantly and inescapably mismatched, VR sickness cannot be eliminated by tracking and responding to physical motion with greater accuracy.

How might we address these situations? We begin by noting the well known relationship between display Field of View (FOV) and VR/simulator sickness —Decreasing FOV, in general, can decrease sickness (e.g., [5, 13]). However, there is also a well known relationship between FOV and *presence*, the subjective experience of being in one environment, even when one is physically situated in another [32]—Decreasing FOV can reduce the user's sense of presence (e.g., [4, 23, 34]).

To reconcile these two effects, we decided to investigate the utility of dynamically decreasing FOV in situations in which a larger FOV would be likely to cause VR sickness (when the mismatch between physical and virtual motion increases), and dynamically restoring it in situations in which VR sickness would be less likely to occur (when the mismatch decreases). Could this increase a user's comfort in a VE? Further, would it be possible to change the FOV in a sufficiently subtle way that users might not perceive that the change was occurring, yet could still benefit from it while not experiencing a noticeably decreased sense of presence?

To explore this possibility, we implemented a software testbed for dynamic FOV modification, example views from which are shown in Figure 1(a–b). Our testbed uses the Oculus Rift DK2 stereoscopic HWD and is built by modifying the player controller of the Tuscany demo that ships with it. We first ran pilot studies to determine suitable parameters for the virtual geometry used for FOV modification (Figure 1c–d) and the rate with which it should change, to achieve an encouraging tradeoff between VR sickness and presence. We then performed a formal study with seated participants (Figure 2) to determine how well our intervention works.

This paper thus makes the following contributions:

1. We present the first study of which we are aware that explores the effects of automatically and continuously changing a participant's FOV throughout the duration of a VR session.

2. Our study shows that dynamic FOV modification can reduce perceived VR sickness.

3. We demonstrate that FOV modification can be made sufficiently subtle that most users do not notice when it occurs and do not experience a noticeable change in presence in sessions in which they experience it.

In the remainder of this paper, we first review previous related work. Then, we introduce the design and implementation of dynamic FOV modification. Next, we present a user study that explores the effectiveness of our approach compared with an unmodified FOV control condition, for physically stationary users virtually exploring a VE. Following this, we discuss the results of the study. Finally, we present our conclusions and directions for future work.

2 Previous Work

Researchers have long studied the relationship between FOV, presence, and VR/simulator sickness. For example, Prothero and Hoffman used physical eye masks to reduce the FOV of a wide-FOV HWD [19]. Their study supports the idea that a wider FOV is necessary for a higher degree of presence.

While simulator sickness and VR sickness are closely related, symptoms tend to be more severe with HWDs than with screen-based simulators [28, 17, 24]. Lin et al. explored the effects of FOV on presence, simulator sickness and various other variables using a screen-based simulator [13]. Their study employed a within-subject design, with data from 10 participants at four FOVs (60°, 100°, 140°, and 180°). Participants showed higher simulator sickness and presence scores as FOV increased. Seay et al. [23] performed a multi-screen driving simulator study, and found that higher FOV correlated with some aspects of higher presence and simulator sickness, suggesting that a large FOV is a "double-edged sword." Other researchers have explored the effects of placing objects near the

Figure 2: Seated study participant wearing HWD and holding gamepad used for virtual navigation.

face that affect a user's FOV. For example, Whittinghill et al. describe how the presence of a virtual nose viewed in an HWD reduces VR sickness [31].

In addition to decreasing presence and VR sickness, smaller FOVs have also been shown to reduce performance in tasks in VR. Wells and Venturino explored the question of how much reducing the FOV of an HWD affects performance in a visual search/monitoring task, and showed that smaller FOVs hamper task performance [30]. These findings were mirrored by Arthur [2], who noted that FOV was significant in predicting performance of two tasks: searching for targets through head turns, and walking through a simple maze while avoiding walls.

While these studies all use a single FOV throughout a VR session, some attempts have been made to change the FOV in discrete steps during a session. Kim et al. [10] used a real-time cybersickness detection system that monitored biosignals from participants in a multiple-projector virtual environment (VE). Their system temporarily narrowed the FOV in discrete steps and presented an audio message telling the participant to slow down and take a deep breath when their electrophysiological inputs indicated they were experiencing cybersickness. More recently, one contributor to an online forum devoted to the Oculus Rift placed in front of each virtual camera FOV restrictors, which were simple textures with a transparent circle through which to look [22]. The FOV could be changed manually through keyboard presses. The developer of the application intended it to be used to acclimate users to large FOV VR without experiencing as much VR sickness at the beginning, by increasing FOV over a series of sessions.

The work cited thus far involves no changes or discrete detectable changes to the FOV during a session. In contrast, research in perception has explored how large features in a scene can be changed discretely [26] or slowly [25], without being detected by a majority of observers. Intentionally imperceptible change was later used in the development of techniques for redirected walking in VR [20], in which the mapping of physical rotation to virtual rotation [20] or the layout of the VE itself [29] are dynamically modified during a session. This earlier research inspired us to explore whether imperceptible modifications to FOV could be effective in reducing VR sickness.

Finally, the closest work to our own is a patent application we found while writing up the results of our study, which proposes dynamically modifying FOV in response to physiological and motion signals from a user, in order to decrease or prevent VR sickness [3]. However, it does not describe a user study testing the concept, and does not suggest the potential to vary FOV in such a way that the

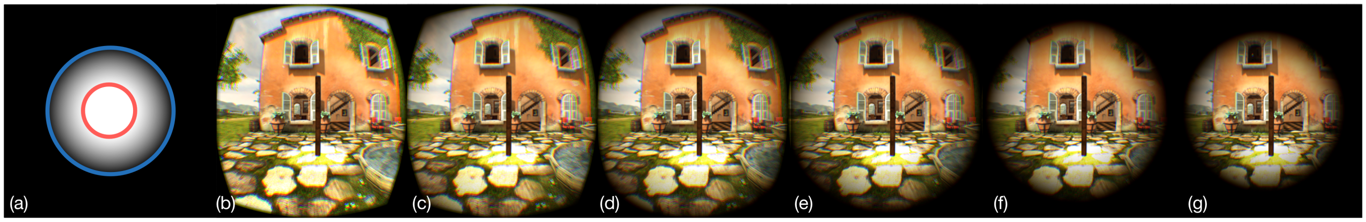

Figure 3: (a) FOV restrictor texture with IFOV (red circle) and OFOV (blue circle). IFOV–OFOV pairs generated by preferred cutout texture (see Section 4.1): (b) 120°–155°, (c) 68°–120°, (d) 58°–110°, (e) 50°–100°, (f) 43°–90°, (g) 36°–80°.

user is unaware of it.

3 Equipment

We use an Oculus Rift DK2 HWD with integrated 6DOF position and orientation tracking, driven by Oculus SDK 0.4.4 on an AMD Phenom II X4 965 Black Edition Quad Core Processor (3.4 GHz), 8GB RAM, with Nvidia GeForce GTX 680 running Windows 8.1. 6DOF head tracking allows the seated user to translate and rotate their head within the tracking volume of the DK2, as shown in Figure 2, with the DK2 cable loosely guided by a Manfrotto 143 Magic Arm and 275 Mini Spring Clamp attached to the back of the chair. In addition to 6DOF-head-tracked control of the view, a Logitech Gamepad F310 controller is used to translate along the ground and rotate around the up axis. The application was developed by modifying the Oculus Rift Tuscany demo using Unity 4, and runs at an average of 75 frames per second, with a measured latency of 15–30ms.

4 Design and Implementation of Dynamic FOV Restrictors

To manipulate the perceived FOV of the DK2, we created two rectangles, each of which was placed close to and in front of the center of projection of one of the two view frusta, and parallel to its base, one for the left eye, and one for the right eye. We call these rectangular structures *FOV restrictors* (Figure 1c–d). Each FOV restrictor was mapped with a texture containing a black opaque rectangle with a variable transparency circular hole in the center, which forms a see-through cutout. Each FOV restrictor is placed at the same fixed distance from its center of projection, and when scaled up or down, respectively, about its center, increases or decreases the perceived FOV, up to the maximum of the display.

The variable transparency circular hole is defined by an inner radius and an outer radius. These two radii specify an annulus, shown in Figure 3(a), that linearly increases in opacity from completely transparent within the inner radius (red circle), corresponding to an inner FOV that we refer to as *IFOV*, to completely opaque beyond the outer radius (blue circle), corresponding to an outer FOV, *OFOV*. To implement a hard edge cutout, we set $IFOV = OFOV$. To create a soft edged cutout, we set $IFOV < OFOV$, causing transparency to decrease linearly from *IFOV* to *OFOV*, as shown in Figure 3(b–g).

We ensure that a FOV restrictor is scaled no smaller than the planar cross section of the camera frustum in which it resides, to prevent the scene from being viewed around it. However, if a FOV restrictor is scaled much larger than that cross section it could extend into and partially occlude the other camera frustum. To prevent this, each FOV restrictor is rendered only by its own camera, allowing it to be arbitrarily large. Although only one person experienced our study at a time, restricting rendering this way in a multi-person VE would ensure that a user's FOV restrictors were invisible to others. Finally, to ensure that the FOV restrictors are rendered on top of all other scene content in our Unity testbed, each FOV restrictor is assigned to a copy of the scene camera for its eye. This special camera, which has a high "depth" value, renders the FOV restrictor last after clearing the depth buffer.

4.1 FOV Restrictor Appearance Parameters

To determine appearance parameters for the FOV restrictors, we designed a set of formative pilot experiments, which we ran with eight participants. To find a suitable minimum FOV ($minFOV$), a participant wearing the DK2 was placed at a set location in the modified Tuscany VE (see Section 5.2). While physical head movement was tracked with the DK2 6DOF tracker, participants remained seated and moved relatively little. The participant was given a gamepad with which to translate on the ground and rotate about the vertical axis in the VE. We used cutouts with a hard circular edge ($IFOV = OFOV$), where the view through the cutout is completely transparent, and the rest of the FOV restrictor is opaque. Participants were not told about the existence of the FOV restrictors.

With the gamepad initially disabled and starting from an initial FOV (120°) greater than that of the DK2 FOV, the experimenter scaled down the FOV quickly in steps of 2.5° relative to the eye (ca. 1s per step), while the participant was told to report if they noticed any visual effect. At the point at which the participant reported noticing something, they were asked what they noticed. Participants first noticed a change at an FOV of 95° (mode, max 107.5°, min 75°), and all then noted that the FOV was decreasing. At that point, the participant was allowed to move using the gamepad. The experimenter continued to decrement the FOV, while the participant was asked to report when they thought the incrementally decreased FOV detracted from the experience. This point was 80° (mode, max 100°, min 75°). Then, the participant was told to report when the FOV no longer detracted from the experience, and the experimenter started incrementing the FOV until the participant indicated this was the case, which occurred at an FOV of 90° (mode, max 100°, min 75°). We will refer to this as the participant's *preferred minFOV*.

In earlier experimentation, we had noticed that hard-edged cutouts are relatively easy to see and distracting, especially as they change in scale. To address this issue, we created FOV restrictors with soft-edged cutouts, as described above. We used a set of 11 textures, each defined by the same outer radius and a different inner radius. Assume that each of these 11 textures is mapped to a FOV restrictor scaled such that $OFOV = 100°$. The first texture we used up to this point (texture 1) was generated so that its inner and outer radii were equal, and thus $IFOV = 100°$, forming a hard-edged cutout. At this scale, each successive texture has an *IFOV* 10° smaller than the *IFOV* of the previous texture, such that the second texture (texture 2) would have $IFOV = 90°$, texture 3 would have $IFOV = 80°$, and so on, down to texture 11, with $IFOV = 0°$.

At this point, our pilot-study participants were using FOV restrictors with $IFOV = OFOV$ *preferred minFOV* for that participant. Our next step was to determine the type of soft-edged cutout participants prefer. That is, while maintaining $OFOV =$ *preferred minFOV* (to ensure that anything beyond the participant's preferred minFOV remained completely occluded), we wished to determine a preferred *IFOV* that traded off the distraction caused by a harder-edged cutout against the brightness reduction caused by a smaller *IFOV*.

To determine an appropriate *IFOV*, each participant used the gamepad to travel in the VE, during which we successively replaced the cutouts, by switching among our 11 textures. This way, we changed the *IFOV*, while maintaining $OFOV = preferred\ minFOV$. The participant alternated between pairs of *IFOV*s, similar to how an ophthalmologist's patient compares lenses, until the participant decided on an *IFOV*.

Participants confirmed through verbal feedback that hard-edged cutouts (created by texture 1) were distracting, as were the Mach bands caused by an *IFOV* close to the *OFOV*. Assuming FOV restrictors scaled such that $OFOV = 100°$, strong Mach bands were visible for *IFOV* between $90° - 70°$ (textures 2–4), but were greatly reduced for *IFOV* between $60° - 40°$ (textures 5–7). Decreasing *IFOV* further (textures 8–11) completely eliminated visible Mach bands, but at the expense of a level of brightness reduction that participants also found undesirable. $IFOV = 50°$ (texture 6) was unanimously preferred by participants, and was what we finally chose.

We then initialized this texture by scaling it such that $IFOV = 120°$ while resting, placing it outside the DK2 FOV. Noting that the *IFOV* and *OFOV* of any texture are scaled differentially, since each is determined by scaling the tangent of half its angle, this resulted in $OFOV \approx 155°$. *IFOV* and *OFOV* change in tandem as the FOV restrictor geometry is scaled, producing the images shown in Figure 3(b–g)

4.2 FOV Restrictor Scaling-Rate Parameters

After determining appearance parameters for our soft-edged cutouts, we had to determine an appropriate restrictor scaling rate. The tradeoff here was between scaling fast enough to reduce VR sickness, but slow enough to minimize distraction. To accomplish this, we ran an additional pilot study, in which we changed the FOV restrictor scale as a function of gamepad speed and angular velocity, ignoring physical head motion. We noted that participants were less aware of the FOV restrictors converging during gamepad-controlled rotation. They also felt more uncomfortable making these rotations; therefore, we decided to scale down the restrictors at a much quicker rate as angular velocity increased. Based on these informal pilot experiments, the FOV restrictors were intially set to contract at rate *CRate*, defined as:

$$CRate = Abs(angularVelocity/20) + (overallSpeed * 4),$$

where *angularVelocity* is the rate at which the participant rotates using the gamepad in °/s, and *overallSpeed* is the gamepad translation speed, which is controlled by the gamepad joystick position, and is capped at a maximum of 1.5 m/s.

FOV restrictors contracting at the same rate were more noticeable at a small FOV than at a large FOV. This suggested changing the contraction rate based in part on the FOV. Recall the 80° (mode) FOV at which pilot-study participants considered hard cutouts to detract from the experience. In the following description, we assume that cutouts stop contracting at a minimum FOV (*minFOV*) of 80°. Based on further informal experimentation, we decided on the following approach.

For a virtually moving participant, when $OFOV > 130°$, then $IFOV > 80°$. Here, *OFOV* is not in the participant's FOV, and rapid contraction or expansion of the transparent part of the cutout is not as noticeable. Thus, we set $CRate = CRate * 3$. When $120° < OFOV \leq 130°$, we converge the FOV restrictors at *CRate*. When $80° < OFOV \leq 120°$, then $CRate = CRate * 0.5$, as cutout movements are most noticeable here. When $OFOV = 80°$, we maintain the FOV restrictor's current state.

For a virtually stationary participant, when $IFOV < 80°$ and $OFOV < 130°$, FOV change is more noticeable. Here, we use $CRate = -3°/s$. When $80° \leq IFOV < 120°$, constant changes in FOV are not as noticeable. Here, we use $CRate = -9°/s$. When $IFOV = 120°$, as seen in Figure 3(b), no part of the circular cutout occludes any of the VE. Here, there would be no further enlargement or contraction of the restrictors when the participant is stationary.

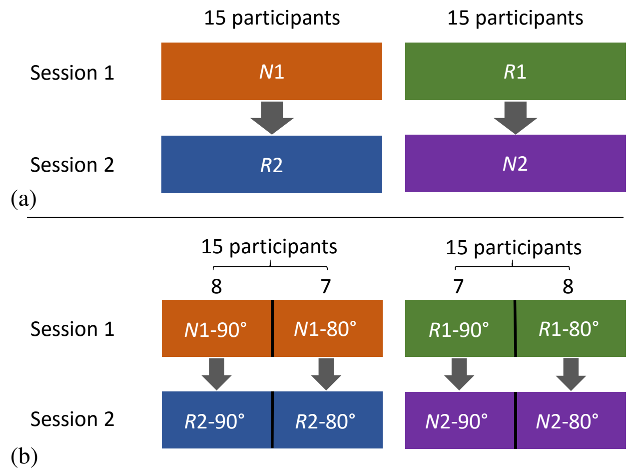

Figure 4: Design of two-session experiment. (a) Division into two groups, counterbalanced as to condition order. (b) Further subdivision of groups by minimum FOV of restrictors used in *R* session.

Finally, note that while 6DOF head-tracking also controls the view, the FOV restrictors respond only to gamepad translation and rotation.

Figure 5: Overhead view of modified Tuscany VE villa, with red dots representing waypoint positions and yellow lines delineating study path. (Roof was removed for image, but not for study.)

5 User Study

5.1 Experiment Design

We employed a two-session mixed experimental design, testing the effectiveness of a dynamically changing FOV vs. the unmodified FOV of the HWD. A total of 32 affiliates of our institution were recruited through email and posters to participate in the experiment; of these, 30 were present for both sessions, and only their data was analyzed. Each participant's two sessions were held on different days, counterbalanced by condition.

When experiencing the intervention, the dynamically changing FOV was set to converge to $OFOV = 90°$ at minimum for 15 participants, and $OFOV = 80°$ at minimum for the other 15. We chose these values because our pilot study participants determined 90° (mode) to be the *preferred minFOV* and 80° (mode) to be the largest FOV that detracted from their experience. All participants were numbered by the order in which they performed their first session;

Figure 6: Bird's-eye view of modified Tuscany VE, with waypoint posts joined by yellow lines delineating the study path. Villa is at bottom right.

for example, the first participant was 1, and the tenth participant was 10. In their first session, odd-numbered participants experienced the control condition (N = No Restriction), while even-numbered participants experienced the intervention condition (R = Restriction). When a participant returned for their second session, they experienced the other condition.

We refer to a group of participants with the notation $Restrictor[Session][\text{-}FOV]$, where $Restrictor$ is N or R, optional $Session$ is 1 or 2, and optional FOV is 80° or 90°. As shown in Figure 4(a), of the 30 participants analyzed, the 15 who started with the control condition are $N1$, and the 15 who started with the restrictors are $R1$. $N1$ participants used restrictors in their second session (referred to as $R2$), while $R1$ participants used the control in their second session (referred to as $N2$). As another example, Figure 4(b) shows that the eight participants in $N1$-90° experienced their second session as $R2$-90°.

5.2 Virtual Environment

The VE used was a modification of the original Tuscany demo from Oculus SDK 0.4.4, as shown in Figures 5–6. A set of 168 sequential waypoints, only one visible at a time, were added to guide the participant's movement in the VE. Each waypoint is a vertical post with a surrounding particle effect animation and lit base (Figures 1, 3, and 7). Beginning with the first waypoint, as a participant approached within 1.3m of a waypoint, it would disappear, and the next would appear in an easily detectable location.

The scene was also modified to make the land outside the walled villa accessible. Several objects were added to make the outskirts more interesting, including an additional house with a smoking chimney. All textured 2D billboard objects within the outer grassy environment were removed.

5.3 Procedures

Upon arriving for the first session, each participant was given the Stereo Optical Co. Inc. Stereo Fly Test to screen for stereo vision. The participant then completed a questionnaire about their demographics (gender, susceptibility to motion sickness, prior use of VR devices and typical computer use and frequency) and a pre-exposure Simulator Sickness Questionnaire (SSQ) [8]. Following this, they watched an instructional video that explained how to traverse the VE using the gamepad. They were instructed to follow the set of waypoints at their own pace, as each appeared. After every five waypoints, gamepad translation and rotation were disabled (although 6DOF head tracking still controlled the view) and the participant was asked the question, "On a scale of 0–10, 0 being how you felt coming in, 10 is that you want to stop, where are you now?" (the same question used by Rebenitsch and Owen) [21]), as shown in Figure 7. The participant was instructed to use a set of dedicated buttons on the gamepad to enter their response, which we refer to as their *discomfort score*,

The 6DOF position and rotation of the participant's avatar in the VE was tracked every frame. Participants were free to do what they wanted as they progressed through the VE, as long as they followed the waypoints. Participants were instructed in advance that if they reached *discomfort score* 10 the experiment would terminate. They were also told that they could terminate the experiment immediately if they hit a predefined combination of buttons or notified the experimenter. Otherwise, the experiment would end when the participant reached the last waypoint. Upon completing or terminating the experiment, participants were asked to fill out a post-experiment questionnaire, including a Witmer–Singer presence questionnaire [32] and the second half of the SSQ. In addition, participants were explicitly asked if they noticed anything strange happening in their virtual experience (as described below).

During the second session, the participant completed the pre-experiment questionnaire, performed the experiment using the condition they had not yet tried, and ended by answering the same post-experiment questionnaire.

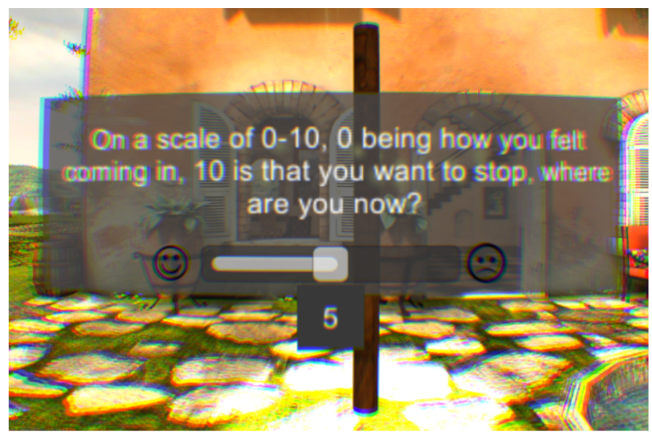

Figure 7: Question asked to determine participant *discomfort score*.

5.4 Information Gathered

Information was collected through our instrumented VE and pre- and post-session questionnaires. All questionnaire data was obtained using Google Forms. At each frame, the VE recorded the position and orientation of the participant in the VE, the position and orientation of their head in the physical world, and their most recently entered discomfort score. We assume that each recorded discomfort score is representative of the participant's level of discomfort up until their next recorded discomfort score. This made it possible to compute the participant's time-weighted *average discomfort score* (*ADS*) over the duration of their session. We also computed an *ending discomfort score*—the value at which the participant ended the experiment (10 if they terminated the experiment; lower otherwise).

Each post-session questionnaire also included a set of questions to gauge whether the participant noticed the FOV restrictors in action, and if they did, whether the restrictors hampered their experience. These questions were based on those used by Suma et al. to determine the effectiveness of their application of change blindness to redirected walking [29]. Participants were asked to rate the following seven questions (of which only italicized questions 3 and 5 relate to the FOV restrictors) on a scale of 1–7, where 1 was "Did not notice or did not happen" and 7 was "Very obvious":

1. I saw the virtual environment get smaller or larger.
2. I saw the virtual environment flicker.
3. *I saw the virtual environment get brighter or dimmer.*
4. I saw that something in the virtual environment had changed color.
5. *I felt like my field of view was changing in size.*
6. I felt like I was getting bigger or smaller.
7. I saw that something in the virtual environment had changed size.

Following these questions, participants were asked to answer *yes* or *no* to the statement, "I noticed at least one of the things listed in this section or something else unusual." They were given an additional set of five questions if they answered yes:

1. Based on your last answer, what did you notice the most?
2. Rate on a scale of 1–7 how confident you were about noticing this, where 1 was "Did not notice or did not happen" and 7 was "Very obvious."
3. Do you think this made the virtual environment more or less comfortable. If so, how?
4. Do you think this made the virtual environment more or less enjoyable? If so, how?
5. Rate on a scale of 1–7 whether you would want this to be included in future virtual experiences, where 1 was "Do not want in future virtual experiences" and 7 was "Definitely want in future virtual experiences."

Participants were then asked to respond with either *yes* or *no* to the question "I noticed one more of the things listed earlier or something else unusual." If they responded with *yes*, they were instructed to answer the same five questions about what they noticed the second most. Finally, participants were asked, "Is there anything else you would like us to know?"

5.5 Hypotheses

Based on results from a pilot study, we formulated four hypotheses, where the first three involve the user's level of discomfort and the last one addresses their awareness of the FOV restrictors. Recall that participants who did not use FOV restrictors in their first session ($N1$), used them in their second session ($R2$). Similarly, participants who used FOV restrictors in their first session ($R1$), did not use them in their second session ($N2$).

H1 *$N1$ participants will report more discomfort than $R1$ participants.* This relationship is expected if the FOV restrictors reduce VR sickness. Note that the comparison is between participants and involves no order effects.

H2 *$N1$ participants will report more discomfort than $R2$ participants.* We expect this if the FOV restrictors reduce VR sickness. However, we note the possibility that $R2$ participants might experience less discomfort because it was their second session, and in this within-subject comparison, they might be more inclined to preemptively terminate their second session if they had a bad experience in $N1$.

H3 *R-90° participants will report more discomfort than R-80° participants.* We expect that the wider FOV will not be as effective in decreasing the level of VR sickness.

H4 *A majority of R-90° participants will not notice the FOV change, while a majority of R-80° participants will notice the FOV change.* Note that our pilot experiment identified an 80° FOV as detracting from the experience.

6 Participants and Groups

Of the 32 participants recruited, 30 (11 female, age 20–27, average 22.1) attended both sessions, and are the only participants whose data was considered in our data analysis. The remaining two attended only the first session. Of these two, one used N, and terminated the session, communicating that he did not want to attend the next session because of an unpleasant initial experience. The other participant who quit also experienced N, but finished the session with an *ADS* under one. However, this participant could not attend his next session due to an unforeseen schedule conflict. No participants failed the Stereo Fly Test.

Stanney and Kennedy [27] note that 30–40% of participants in flight-simulator studies do not experience simulator sickness, but only 5–10% of participants in early VR studies are asymptomatic. We also expected some of our participants would not experience VR sickness, noting further that there have been many changes in the technology being used since those early studies. Inspecting our data, we observed that 6 out of the 30 participants had ADS< 1 (all < 0.65) in both their R and N sessions.

As we are interested in users who *do* experience VR sickness symptoms, we decided to remove from our data and analysis the six participants who were essentially unaffected by VR sickness, leaving 24 participants. Upon making these adjustments, we found that our data was normally distributed, as verified by an Anderson–Darling normality test with a confidence of 83.07%

Of the 24 participants whose data we used to evaluate H1–H3, 12 started in $N1$, and the other 12 started in $R1$. Generally, males are less prone to VR sickness [6, 16]. The $N1$ group had 9 male and 3 female participants, while $R1$ had 7 male and 5 female participants. There are good correlations between motion sickness and VR sickness [7]. $N1$ had two people susceptible to motion sickness, while group $R1$ had none. As H4 addresses detection of the FOV restrictors, we used all 30 participants to analyze it.

We used a Bonferroni-corrected $\alpha = 0.025$ $(0.05/2)$ as a threshold for significance.

Table 1: Participant Discomfort By Session

	Completed (12)	Mean ADS (SD)	Mean SSQ (SD)	Mean RDS (SD)
$N1$	1	4.95 (1.35)	61.1 (30.4)	7.33 (1.92)
$N2$	9	3.16 (1.94)	48.4 (44.4)	4.63 (3.07)
$R1$	6	2.97 (1.24)	57.7 (17.2)	5.18 (3.01)
$R2$	6	3.93 (2.25)	48.1 (30.2)	5.05 (2.86)

7 Results

H1: $N1$ vs. $R1$

In Figure 8(a), $N1$ vs. $R1$, each thin line provides the progression of discomfort scores of a single participant over time. The thin dark orange lines represent $N1$ participants, and the thin green lines represent $R1$ participants. The circle or diamond at the end of each line represents the point at which a participant either terminated (finished early) or completed the session. One of 12 $N1$ and 6 of 12 $R1$ participants completed the first session. The lower right portion of the graph, which corresponds to spending a longer time in the VE with lower discomfort scores, consists mainly of $R1$ participants, while the upper left portion, which corresponds to those who terminated early, contains a higher proportion of $N1$ participants.

The discomfort scores of all $N1$ participants were averaged at each point in time to produce the thick dark orange line, representing the mean discomfort score at each point in time. If a participant terminated early, a score of 10 was used for them from that point on, until the last user finished. If a participant finished without terminating (i.e., completed), their ending score was used from that point on, until the last user finished. This was also done for the

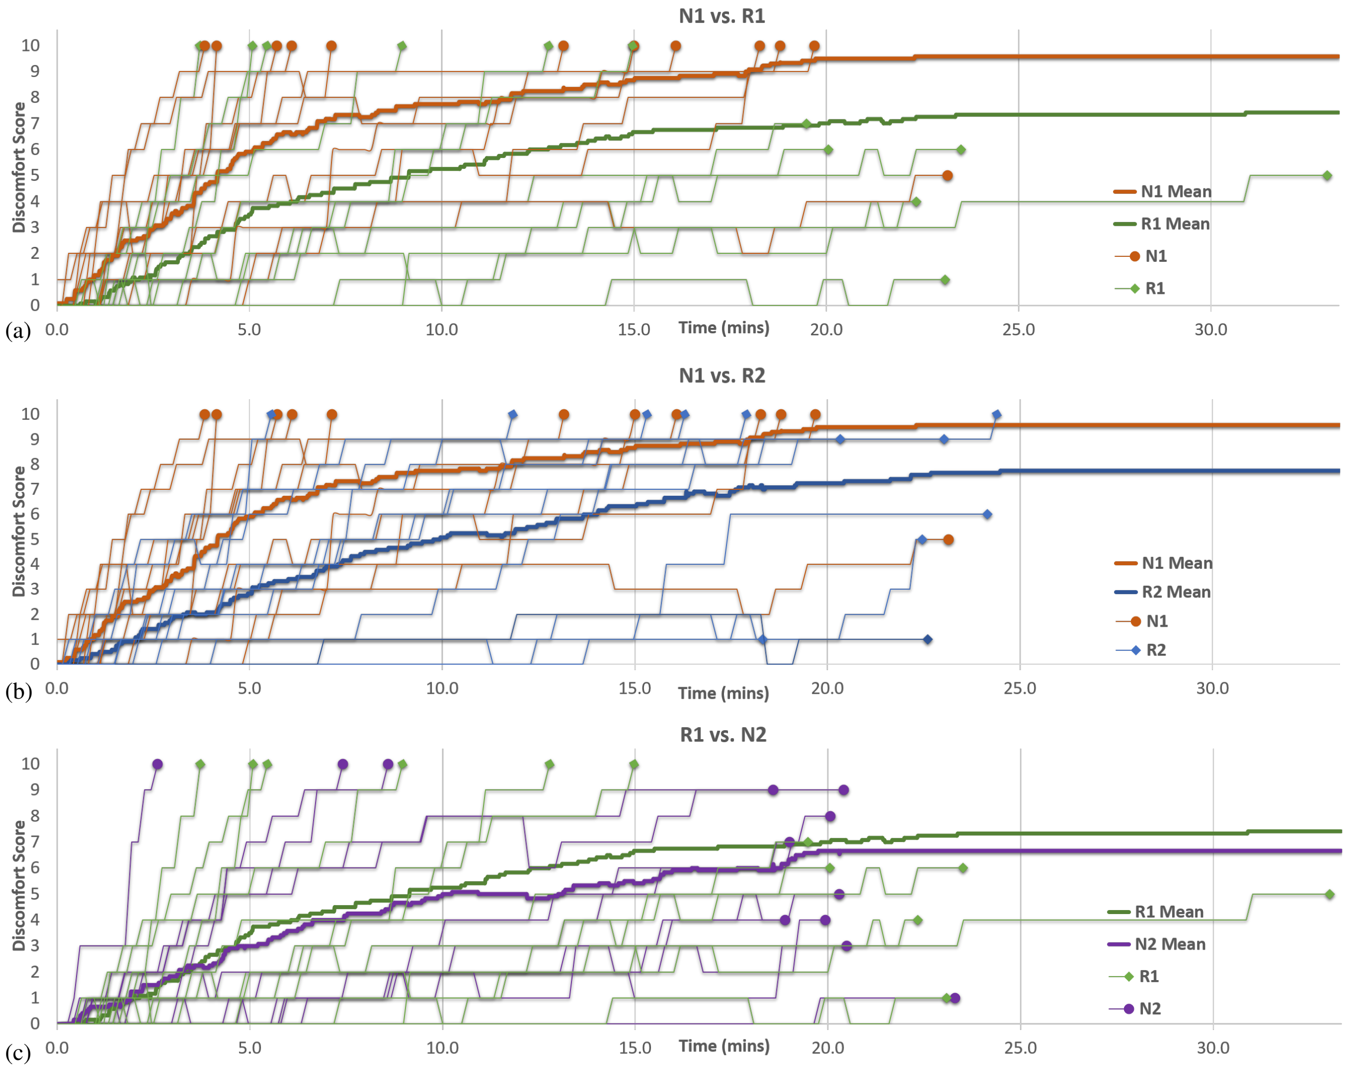

Figure 8: Discomfort scores over time for individual participants and mean discomfort scores over time by session. (a) $N1$ vs. $R1$. (b) $N1$ vs. $R2$. (c) $R1$ vs. $N2$. (Same color coding as Figure 4.) Though discomfort scores change instantaneously, graph uses sloped lines for legibility.

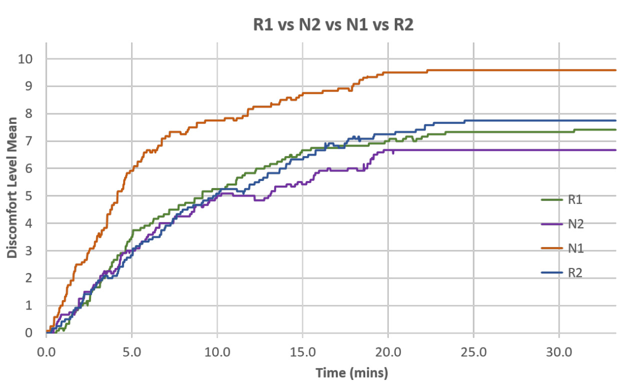

Figure 9: Mean discomfort scores over time by session. (Same color coding as Figure 4.)

$R1$ participants to produce the thick green line. Figure 9 compares these mean discomfort scores for $R1$, $N2$, $N1$, and $R2$.

Comparing ADS for $N1$ and $R1$ (Table 1) supports H1 (two-sample t-Test, $df = 22$, $t = 1.717$, $p = 0.00057$*). Comparing SSQ scores did not yield a significant difference (two-sample t-Test, $df = 17$, $t = 0.336$, $p = 0.370$). (We compare only post-exposure SSQ scores throughout.) Since these are participants in their first session, there are no order effects to compromise our results.

While ADS and SSQ scores may be good measures of subjective discomfort at the time at which a participant completes the experiment, they may not be the best measure of a participant's relative performance if they terminate early. For example, two participants could finish the session with the same ADS and SSQ scores, yet one participant could have spent much more time in the VE before finishing. While we could compare ADS and SSQ scores for just those participants who completed the session, we did not do this because the pool of participants who completed sessions was relatively small (Table 1, "Completed"). To address this, we propose a metric that takes time spent in the VE into account, which we call the *Relative Discomfort Score* (*RDS*).

To calculate RDS for each participant, we first find the time t_{max} at which the last participant finished. In our experiment, the longest time a participant spent in the VE was 1945s ($\approx$ 32.5 mins). However, this participant was an outlier. Using the times only of completed sessions as input to Grubb's test, this value was the only outlier ($n = 22$, mean = 1288s, SD = 180s, Z-crit = 2.75, Z= 3.65). Therefore, we set t_{max} to the second longest completion time of

1470s (24.5 mins, $Z = 0.615$).

If a participant terminated before t=1470s, their last discomfort score (10) is repeated each second from the time the participant terminated until t_{max}. Similarly, if a participant finished early with a score less than 10, this finishing score is repeated. In both these cases, we refer to the stop time as t_{stop}, the discomfort score at t_{stop} as DS_{stop}, and the discomfort score at each second i prior to t_{stop} as DS_i. Next, we integrate the resulting discomfort score values from the time of the participant's first movement, to t_{max}, and then divide that value by t_{max} to yield:

$$RDS = \frac{\sum\limits_{0 \le i < t_{stop}} DS_i \; + \; (t_{max} - t_{stop} + 1)DS_{stop}}{t_{max}}.$$

Note that if a participant terminated the experiment with DS_{stop}= 10, at the beginning of the experiment, their RDS is approximately 10, yet if the participant completed the experiment without ever incrementing their discomfort score, RDS = DS_{stop} = 0.

Comparing RDS for $N1$ and $R1$, using the Mann–Whitney U-Test, Z = 1.6454, p = 0.04947$^+$. Assuming the two groups are normal and comparing them using a two-sample t-Test, df = 19, t = 1.729, p = 0.0258$^+$. While these values are not significant after Bonferroni correction, given our small sample size, we believe that these results indicate an encouraging trend, in support of the restrictors being effective.

H2: *N*1 vs. *R*2

Mean discomfort scores over time for $N1$ participants and $R2$ participants can be compared in Figure 8(b). Comparing ADS for $N1$ and $R2$ does not support H2, after Bonferroni correction (two-sample t-Test: df = 11, t = 2.1436, p = 0.02763 $^+$). While comparing SSQ scores does not indicate a significant diference (paired t-Test: df = 11, t = 1.25, p = 0.119), comparing RDS supports H2 (Wilcoxon Signed-Rank Test, $Z = -2.7456$, $p = 0.00298$*).

We hypothesized that there would be a few factors contributing to these results. First, we expected that participants would generally perform better in their second session due to order effects [9]. However, we also expected that some participants would perform worse because an unpleasant $N1$ session could make them intolerant to $R2$. This idea was supported by a few of our participants. A participant who terminated earlier in $R2$ than in $N1$, wrote in the additional information section, "After the previous experiment, I felt really bad for a very long time, more than 8 hours. I had small headaches and nausea for the whole time. They wouldn't disappear before I got asleep. For today experiment, I didn't last as long as the last experiment because I didn't want to feel as bad as last time. I stopped right after I started having a little headache." In addition, as mentioned earlier, one participant whose results we did not analyze did not show up for their second session, because of an unpleasant $N1$ experience. If it wasn't for this sort of aversion, $R2$ scores could on average be even lower. Yet generally, participants performed better as shared by a participant after $R2$, "The second experiment, compared to the first was much easier, I felt used to it. Thus, I didn't get nauseous as the first time. I was able to go further and enjoy more the experiment."

Taking this into account, it would be interesting to examine how $R1$ compared with $N2$. From Figure 8(b), we see that $N2$ performance was similar to that of $R1$. Comparing $R1$ and $N2$, we see no statistical significance in any of our measures. This is likely due to the combination of two effects. First, $N2$ follows $R1$, noting that order effects likely play a role here. Second, an N condition is likely more uncomfortable than an R condition. If order effects were the sole contributor to better performance, we would expect a larger drop in scores from $R1$ to $N2$. Thus, if the FOV restrictors did not make a difference, we would also expect $N1$–$R2$, and $R1$–$N2$ session pairs to have similar differences in scores. This is not the case: we found the mean score differences for participants between their $N1$ and $R2$ sessions were ADS (mean = 1.02, SD = 1.65) and RDS (mean = 2.28, SD = 2.24), and the mean score differences for participants between their $R1$ and $N2$ sessions were ADS (mean = -0.192, SD = 2.21) and RDS (mean = 0.551, SD = 2.49). We do not find significant differences in ADS (two-sample t-Test, df = 20, t = 1.52, p = 0.0718), but when comparing RDS, given that our samples were normally distributed, our data suggests that this could be the case (two-sample t-Test, df = 22, t = 1.79, p = 0.0437$^+$). It could also be the case that $R1$ helps participants acclimate to future N sessions, beyond order effects. Because of this, the $R1$ session might reduce the intolerance that we saw in $N1 - R2$, which could be another reason why $N2$ scores on average are lower, rather than higher.

One participant finished $N1$, while nine finished $N2$. Comparing ADS using a two-sample t-Test: df = 20, t = 2.6264, p = 0.0081*. Comparing SSQ scores reveals no significant difference (two-sample t-Test: df = 19, t = .814, p = 0.2130. Comparing RDS, using two-sample t-Test: df= 18, t = 2.5797, p = 0.0094*; Wilcoxon Signed-Rank Test, Z-Score is 2.0496, p = 0.02018*).

As can be seen, participants who use R first are affected less by the N condition. Noting that $N1$ was more uncomfortable than $R1$ and that $N2$ participants did not report a significantly different level of discomfort than $R2$ participants, this suggests that participants who start with the R condition generally have a better experience in their combined VR sessions, and further suggests that FOV restrictors might help VR users have a more comfortable launch,"warming up" to their initial VR experiences.

H3: FOV: 80° vs. 90°

While both 80° and 90° minimum FOVs ($minFOV$s) reduce discomfort, we found no significant differences when comparing R-80° ADS (mean = 3.74, SD = 2.21) and R-90° ADS (mean = 3.16, SD = 1.43) using a two-sample, two-tail t-Test: df = 19, t = 0.761, p = 0.456271878. Comparing R-80° RDS (mean = 5.19, SD = 2.87), and R-90° RDS (mean = 5.037, SD = 3.01), we found no signficance using a two-sample two-tail t-Test: df = 22, t = 0.128, p = 0.899.

A participant who noticed the 80° FOV restrictors, when asked in the questionnaire if the FOV restrictors made him more or less comfortable, stated, "Less comfortable. I had to rotate a lot and the rotation was making me nauseous." As his vision was constrained, he had to rotate more using the gamepad to find his path, and according to most participants, it was rotation using the gamepad that contributed the most to nausea. Thus, to understand if $R2$-80° indeed worked effectively, it would be valuable to look at how performance with only natural tracked rotation without a gamepad, would compare. Finding a technique to reduce nausea felt through such rotation could potentially help resolve this issue.

Presence: *N* vs. *R*

Considering the original 30 participants, the difference in Witmer and Singer presence scores [32] was not signficant between N-90° (mean = 154, SD = 13.1) and R-90° (mean = 153.7, SD = 15.0) conditions, using a paired t-Test: df = 14, t = 0.0935, p = 0.463, even when taking order effects into account. We see that smaller FOV restrictor size did not result in a signficant decrease in presence scores either for N-80° (mean = 144.5 SD =18.9) or R-80° (mean = 147.7, SD = 19.8), using a paired two-tail t-Test: df = 14, t = -0.7910, p = 0.4421.

Taking into account all 30 N and R sessions, the differences in presence scores were not significant between N (mean = 149.267, SD = 16.72) and R (mean = 150.733, SD = 17.53), using a paired two-tail, t-Test: df = 29, t = -0.598, p = 0.554.

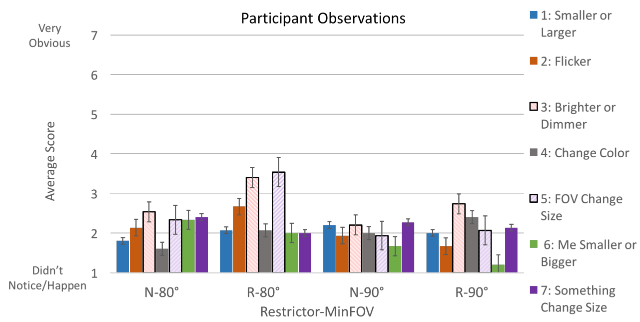

Figure 10: Average Scores for participant observation questions, where 1 was "Did not notice or did not happen" and 7 was "Very obvious." Error bars show Standard Error.

H4: Participant Observation Questions

Figure 10 shows how our two target questions compare with the distractor questions. Note that FOV observations in R-90° are generally between the means of the distractor questions 1, 2, 4, 6 and 7 (described in Section 5.4). However, this is not necessarily the case with brightness. As shown in Figure 11, 3 of 15 R-90° participants noticed the environment dim with a score greater than 5, while only 2 of 15 participants noticed the FOV shrinking with a score greater than 3. These individuals had a score of 7. Still, the majority did not notice these changes. For example, 11 of 15 R-90 participants marked a 1 ("Did not notice or did not happen") for FOV change and 9 of 15 marked a 1 for brightness change, clearly indicating that it was imperceptible to them.

We will now look for differences between the same participants in their N-90° and R-90° conditions. Using the Wilcoxon Signed-Rank Test for comparison, N-90° participants (mean = 2.20, SD = 1.97) and R-90° participants (mean = 2.73, SD = 2.52) did not notice the environment darken, W=4, W-$crit$ (n = 5, $p < 0.05$) = 0, so the result is not significant at $p < 0.05$. These participants also did not notice FOV reduce within the N-90° condition (mean = 1.93, SD = 1.58) and R-90° condition (mean = 2.07, SD = 2.12). Using the Wilcoxon Signed-Rank Test for comparison $W = 8.5$, W-$crit$ ($n = 6, p < 0.05$) = 2. Thus, this result is also not significant at $p < 0.05$.

However, the distributions of the answers to the FOV-related questions relative to the distractor questions seem visually anomalous for R-80°, as shown in Figure 10. In addition, looking at Figure 11, we see a shift of values from the left to the right, between N (top) and R (bottom) graphs. We analyzed whether R-80° participants felt that the restrictors dimmed the environment (mean = 3.40, SD = 2.10) more than N-80° participants (mean = 2.53, SD = 1.81), using a Wilcoxon Signed-Rank Test (Z=-1.89, p=0.0294) and whether R-80° participants sensed that something was happening to their FOV (mean = 3.50, SD = 2.29) more than N-80° participants (mean = 2.33, SD = 1.63), also using a Wilcoxon Signed-Rank Test (Z=-1.78, p = 0.03754^{+}). However, these results were not significant under Bonferroni correction. Though participants noticed the changes using the 80° FOV restrictor, they were still, on average, uncertain about the change, given that a score of 1 was "Did not notice or did not happen" and 7 was "Very obvious."

Even though some participants noticed aspects of the restrictors, were the restrictors a desirable part of the VE? To explore this question we will take a deeper look at the self-response section of the questionnaire.

Six of 30 participants marked a 6 or higher after their R session, regarding noticing FOV change. Three of these participants also marked a 6 or higher on noticing the brightness change (as did an additional two participants not among the six). One other participant mentioned noticing an FOV change, when asked if he noticed anything changing, but marked a 3, when asked the same question

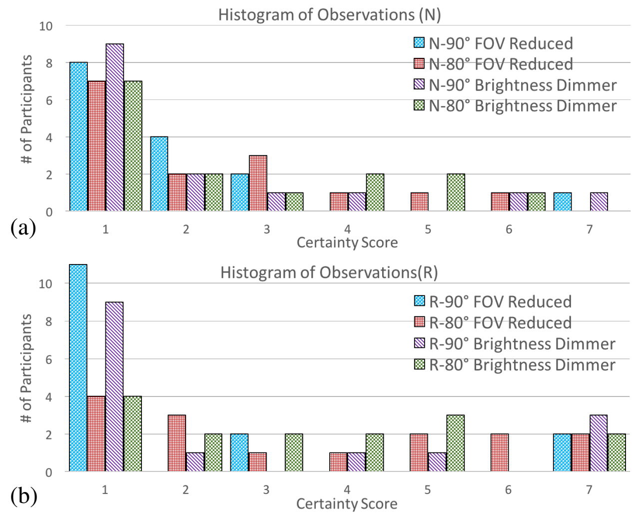

Figure 11: Histogram of responses to participant observation questions relevant to FOV restrictors. (a) N sessions. (b) R sessions.

in the form of a Likert-scale question. No other comments were made about brightness or FOV. Thus, there were nine participants in total who mentioned changes in either brightness or FOV in R. Five of these nine participants either left a comment on this issue or answered the question, "Do you think this made the virtual environment more or less comfortable. If so, how?", and "Rate on a scale of 1–7 whether you would want this to be included in future virtual experiences."

One participant who noticed the light dim after her R1-90° session mentioned that, "Dimming of the lights was more comfortable." She selected a 5 for having it in future experiences. The second participant, who mentioned after his R2-80° session that he noticed "Change in size of vision (i.e. could see less in the periphery)" selected a 7, definitely wanting this feature to be available in future virtual experiences. According to him, the experience was, "More comfortable. The head movements made me less dizzy and seemed more natural." "Yes. I feel better today than I felt yesterday." The third participant after his R2-90° session said, "my view at the corner of my eyes was non-existent. so I could only see things that were directly in my line of sight. rest were blurred or removed. brightness was reduced." When asked, "Do you think this made the virtual environment more or less comfortable. If so, how?" he responded with a 7 ("Definitely want in future virtual experiences"), explaining, "Much comfortable. I wasn't able to complete the study yesterday, I was much more comfortable today. I think removing extra information helps...More enjoyable, as I was able to adjust easily there was less acclimitizaton." The fourth person who selected a 7 in brightness and FOV reduction after his R2-80° session mentioned "The second experiment, compared to the first was much easier, I felt used to it. Thus, I didn't get nauseous as the first time. I was able to go further and enjoy more the experiment." Finally, as mentioned earlier, one participant in his R2-80° session, mentioned that additional head rotation necessitated by the reduced FOV was undesirable, giving it a 1 for future virtual experiences.

8 Conclusions and Future Work

We performed a two-session within-subject user study that explored the effects of dynamically, yet subtly, changing a seated participant's FOV in response to visually perceived motion as they virtually traverse a VE. Even though we had a relatively small number of participants, our data indicates that FOV restrictors helped participants stay in the VE longer and feel more comfortable than they did in the control condition. Our data also suggests that FOV restrictors helped participants experience less discomfort on their first experience, which in turn helped them transition into their next (*N*) session.

FOV restrictors were unnoticed by the majority of participants, with 15 out of 30 participants selecting 1 (definitively stating "Did not notice or did not happen") in response to whether they noticed the FOV decrease during the session in which they experienced the FOV restrictors. Furthermore, 11 of 15 *R*-90 participants marked a 1 for FOV change and 9 of 15 marked a 1 for brightness change, clearly indicating that it was imperceptible to the majority experiencing that condition. Yet those who did notice the restrictors generally preferred to have them. FOV restrictors thus seem helpful, as long as they do not restrict the FOV to an undesirable level. Future work should explore using different restrictor types and shapes. Further experimentation with scaling parameters could make the FOV restrictors more subtle and effective. While we reduced FOV as a simple function of speed and angular velocity, it would be interesting to see how this would compare to reducing FOV based on parameters such as biosignals (e.g., Kim et al. [10]) or optical flow (similar to how Argelaguet [1] automatically adjusts navigation speed based on optical flow).

Future studies should also examine how FOV restrictors could help acclimate users to VR experiences. A larger number of participants would have helped increase the certainty of our results. In addition, there are two conditions that could provide valuable data points to help validate the effectiveness of FOV restrictors in future studies: A group of participants who use *N* in both their sessions, and another group who use *R* twice, to help us compare performance between successive sessions with and without FOV restrictors, given the known tendency for VR sickness to decrease in later sessions [9].

ACKNOWLEDGEMENTS

This material is based upon work supported in part by the National Science Foundation under Grants IIS-0905569 and IIS-1514429. We thank Luis Tolosa and Matthieu Gavaudan for helping coordinate and run our original study, and Carmine Elvezio, Mengu Sukan, Lena Mamykina, and Marissa Burgermaster for their feedback and suggestions.

REFERENCES

[1] F. Argelaguet. Adaptive navigation for virtual environments. In *3DUI, 2014 IEEE Symposium on*, pages 123–126. IEEE, 2014.

[2] K. W. Arthur. *Effects of field of view on performance with head-mounted displays*. PhD thesis, Dept. of CS, UNC Chapel Hill, 2000.

[3] M. Bolas, J. A. Jones, I. McDowall, and E. Suma. Dynamic field of view throttling as a means of improving user experience in head mounted virtual environments, Sept. 18 2014. US Patent App. 14/216,220.

[4] J. J. Cummings and J. N. Bailenson. How immersive is enough? A meta-analysis of the effect of immersive technology on user presence. *Media Psychology*, (ahead-of-print):1–38, 2015.

[5] P. DiZio and J. R. Lackner. Circumventing side effects of immersive virtual environments. In *Proc. HCI International*, volume 2, pages 893–896, 1997.

[6] R. Kennedy, D. Lanham, C. Massey, J. Drexler, and M. Lilienthal. Gender differences in simulator sickness incidence: Implications for military virtual reality systems. *Safe Journal*, 25(1):69–76, 1995.

[7] R. S. Kennedy, J. E. Fowlkes, K. S. Berbaum, and M. G. Lilienthal. Use of a motion sickness history questionnaire for prediction of simulator sickness. *Aviation, Space, and Env. Medicine*, 63(7):588–593, 1992.

[8] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *Int. Jnl. of Aviation Psyc.*, 3(3):203–220, 1993.

[9] R. S. Kennedy, K. M. Stanney, and W. P. Dunlap. Duration and exposure to virtual environments: Sickness curves during and across sessions. *Presence*, 9(5):463–472, 2000.

[10] Y. Y. Kim, E. N. Kim, M. J. Park, K. S. Park, H. D. Ko, and H. T. Kim. The application of biosignal feedback for reducing cybersickness from exposure to a virtual environment. *Presence*, 17(1):1–16, 2008.

[11] E. M. Kolasinski. Simulator sickness in virtual environments. Technical report, DTIC Document, 1995.

[12] J. J. LaViola Jr. A discussion of cybersickness in virtual environments. *ACM SIGCHI Bulletin*, 32(1):47–56, 2000.

[13] J. J.-W. Lin, H. B. Duh, D. E. Parker, H. Abi-Rached, and T. Furness. Effects of field of view on presence, enjoyment, memory, and simulator sickness in a virtual environment. In *Proc. IEEE VR*, pages 164–171. IEEE, 2002.

[14] M. E. McCauley and T. J. Sharkey. Cybersickness: Perception of self-motion in virtual environments. *Presence*, 1(3):311–318, 1992.

[15] Oculus VR. Health and safety warning. http://static.oculus.com/documents/health-and-safety-warnings.pdf.

[16] G. D. Park, R. W. Allen, D. Fiorentino, T. J. Rosenthal, and M. L. Cook. Simulator sickness scores according to symptom susceptibility, age, and gender for an older driver assessment study. In *Proc. HFES Ann. Meeting*, volume 50, pages 2702–2706. Sage Publications, 2006.

[17] E. Patrick, D. Cosgrove, A. Slavkovic, J. A. Rode, T. Verratti, and G. Chiselko. Using a large projection screen as an alternative to head-mounted displays for virtual environments. In *Proc. CHI*, pages 478–485. ACM, 2000.

[18] R. Pausch, T. Crea, and M. Conway. A literature survey for virtual environments: Military flight simulator visual systems and simulator sickness. *Presence*, 1(3):344–363, July 1992.

[19] J. Prothero and H. Hoffman. Widening the field-of-view increases the sense of presence in immersive virtual environments. *HIT Lab Tech. Report TR-95, U. Wash*, 1995.

[20] S. Razzaque, Z. Kohn, and M. C. Whitton. Redirected walking. In *Proc. EUROGRAPHICS*, pages 105–106, 2001.

[21] L. Rebenitsch and C. Owen. Individual variation in susceptibility to cybersickness. In *Proc. UIST*, pages 309–317. ACM, 2014.

[22] scsioverdrive. VR Training in Tuscany, for those with simulation sickness, url = https://forums.oculus.com/viewtopic.php?t=1545/, urldate = 2013-05-16,.

[23] A. F. Seay, D. M. Krum, L. Hodges, and W. Ribarsky. Simulator sickness and presence in a high FOV virtual environment. In *Proc. IEEE VR*, pages 299–300. IEEE, 2001.

[24] S. Sharples, S. Cobb, A. Moody, and J. R. Wilson. Virtual reality induced symptoms and effects (VRISE): Comparison of head mounted display (HMD), desktop and projection display systems. *Displays*, 29(2):58–69, 2008.

[25] D. J. Simons, S. L. Franconeri, and R. L. Reimer. Change blindness in the absence of a visual disruption. *Perception-London*, 29(10):1143–1154, 2000.

[26] D. J. Simons and D. T. Levin. Change blindness. *Trends in Cognitive Sciences*, 1(7):261 – 267, 1997.

[27] K. M. Stanney and R. S. Kennedy. The psychometrics of cybersickness. *CACM*, 40(8):66–68, 1997.

[28] K. M. Stanney, R. S. Kennedy, and J. M. Drexler. Cybersickness is not simulator sickness. In *Proc. HFES Ann. Meeting*, volume 41, pages 1138–1142. SAGE Publications, 1997.

[29] E. Suma, S. Clark, S. Finkelstein, D. Krum, Z. Wartell, and M. Bolas. Leveraging change blindness for redirection in virtual environments. In *Proc. IEEE VR*, pages 159–166. IEEE, 2011.

[30] M. J. Wells and M. Venturino. Performance and head movements using a helmet-mounted display with different sized fields-of-view. *Optical Engineering*, 29(8):870–877, 1990.

[31] D. Whittinghill, B. Ziegler, T. Case, and B. Moore. Nasum virtualis: A simple technique for reducing simulator sickness. http://gtp.autm.net/technology/view/70738, 2015.

[32] B. G. Witmer and M. J. Singer. Measuring presence in virtual environments: A presence questionnaire. *Presence*, 7(3):225–240, 1998.

[33] R. Yao, T. Heath, A. Davies, T. Forsyth, N. Mitchell, and P. Hoberman. Oculus VR best practices guide. *Oculus VR*, 2014.

[34] C. Youngblut. What a decade of experiments reveals about factors that influence the sense of presence. Technical report, DTIC Document, 2006.