



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/76426/>

Version: Published Version

Article:

Ruddle, RA, Payne, SJ and Jones, DM (1999) Navigating large-scale virtual environments: What differences occur between helmet-mounted and desk-top displays? *Presence: Teleoperators and Virtual Environments*, 8 (2). 157 - 168. ISSN: 1054-7460

<https://doi.org/10.1162/105474699566143>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Roy A. Ruddle
Ruddle@CARDIFF.AC.UK

Stephen J. Payne
Dylan M. Jones
School of Psychology
Cardiff University
PO Box 901
Cardiff CF1 3YG UK

Navigating Large-Scale Virtual Environments: What Differences Occur Between Helmet-Mounted and Desk-Top Displays?

Abstract

Participants used a helmet-mounted display (HMD) and a desk-top (monitor) display to learn the layouts of two large-scale virtual environments (VEs) through repeated, direct navigational experience. Both VEs were “virtual buildings” containing more than seventy rooms. Participants using the HMD navigated the buildings significantly more quickly and developed a significantly more accurate sense of relative straight-line distance. There was no significant difference between the two types of display in terms of the distance that participants traveled or the mean accuracy of their direction estimates. Behavioral analyses showed that participants took advantage of the natural, head-tracked interface provided by the HMD in ways that included “looking around” more often while traveling through the VEs, and spending less time stationary in the VEs while choosing a direction in which to travel.

1 Introduction

Experimental studies have shown that people can ultimately develop accurate route- and survey-type (map perspective) spatial knowledge when they navigate large-scale virtual environments (VEs) (Ruddle, Payne, & Jones, 1997). However, this process typically takes a considerable amount of time.

Anecdotal and experimental evidence suggests that immersive displays such as helmet-mounted displays (HMDs) and BOOMs are more effective than desk-top displays when VEs or teleoperation systems are used to perform some close-range tasks, such as protein molecule modeling, target detection, and target tracking (Hartshorn & Hubbard, 1993; Pausch, Shackelford, & Proffitt, 1993; Tachi & Yasuda, 1994). However, other studies found no significant differences between the two types of display when they were used to target radiotherapy beams (Chung, 1992), or to view a three-dimensional sculpture that was subsequently reconstructed in the real world (Mizell, Jones, Jackson, & Pickett, in press).

The physical movements that people make when they use immersive displays provide kinesthetic and vestibular feedback to changes in their orientation that is not present when people use desk-top displays, and this additional feedback may help people to develop spatial knowledge. This article presents the results of an experiment that investigated the differences that occurred when participants used one type of immersive display (an HMD) and a desk-top display to

repeatedly navigate large-scale, virtual buildings. However, we first summarize the key differences between a range of types of VE display, making reference to studies that have investigated aspects of these differences in both virtual and real-world environments.

2 Differences Between VE Displays

VEs may be viewed using devices that include desk-top displays, HMDs, BOOMs, and spatially immersive displays (e.g., CAVEs and screen-projected theaters). These devices are typically used with user interfaces that have different characteristics, support different fields of view (FOVs), and may lead people to perceive VEs in different ways. All of the displays support a variety of resolutions (a factor that is principally affected by cost) and may be viewed in either a nonstereo or a stereo mode.

2.1 User Interface

People typically use abstract interfaces (e.g., mouse, keyboard, joystick, or a Spaceball) to control their translationary movements and changes of direction with desk-top displays. With immersive displays, people typically use abstract interfaces to perform translationary movements, but physically turn around to change direction.

When using desk-top displays, people receive feedback on their movements from visual changes in the displayed scene and the motor actions of their fingers on the interface devices. Visual continuity during changes of view direction is achieved by constraining the rate at which the view direction is allowed to change; even with a graphics supercomputer, the equivalent of a glance over the shoulder takes one to two seconds. The process of glancing becomes more like an implicit instruction to “rotate until you are facing the intended direction and then rotate back.” This changes the work required to integrate the information that is gained during the rotation with the user’s existing spatial knowledge.

By contrast, the visual feedback that people receive when using immersive displays is supplemented by ves-

tibular and kinesthetic feedback from their changes of direction. The effect of this additional feedback on the user’s ability to navigate is not known, but data from some real-world studies suggest that vestibular and kinesthetic feedback help users to develop spatial knowledge, and that physical changes of direction are more important than physical translationary movements for the development of that knowledge (Presson & Montello, 1994; Rieser, 1989; Rieser, Lockman, & Pick, 1980).

In the study by Reiser et al. (1980), participants who were congenitally blind (and therefore received no visual feedback), were shown to have developed reasonably accurate route- and survey-type knowledge of a building with which they were familiar, although their knowledge was significantly less accurate than that of normally sighted participants. However, the importance of vestibular and kinesthetic feedback to the congenitally blind participants, when compared with other methods of learning spatial knowledge such as counting steps or estimating travel time, is not known. In the other two studies (Presson & Montello, 1994; Rieser, 1989), participants estimated directions significantly more accurately when they made physical changes of direction than when they imagined they had changed direction, but no significant differences occurred when participants made physical and imagined translationary movements.

Despite the differences identified above, experimental studies have not shown significant effects when immersive and desk-top displays were used for navigation tasks. A study by Slater, Alberto, and Usoh (1994) found no significant difference between the time that participants took to find one object in a VE model of a real building when using either a HMD or a desk-top display, and also found no significant difference between the time that participants subsequently took to find the same object in the real building. However, participants had no knowledge of the building prior to using the VEs and this effectively means that the participants’ goal was to perform an exhaustive search of the VE until the object was found. Differences may have been found if participants had repeatedly navigated the VEs until they had learned the precise location of the object. The primary purpose of another study (Waller, Hunt, & Knapp, 1998) was to

investigate the transfer of spatial knowledge learned in immersive and desk-top VEs to the real world. In a two-part procedure that was repeated six times, participants navigated a virtual maze and then navigated a real-world version of the maze while blindfolded. No significant differences were found between the immersive and desk-top VEs, but perhaps different effects would have been found if participants had spent more time learning the virtual maze before being first exposed to the real maze, or if the experimental assessment had taken place in the virtual maze.

Interfaces that allow people to make physical translational movements when they travel through VEs are difficult to construct, but one, the omni-directional treadmill, that is currently under development is designed to allow people to run or walk in any direction across a 2-D surface (Darken, Cockayne, & Carmein, 1997). An alternative to a treadmill is an interface that allows people to walk on the spot (walk in place) while traveling through VEs. In one study that investigated this type of interface, participants' sense of presence was greater when they walked on the spot (i.e., made steps without forward motion) than when they did not, but there was insufficient evidence to suggest that the walking interface also helped participants to navigate more efficiently (Slater, Usoh, & Steed, 1995). Although people are likely to become tired if they travel large distances using walking interfaces, this expenditure of physical effort is a desirable feature in some military training applications (Darken et al., 1997).

2.2 FOV

Navigating VEs with a restricted FOV increases the angle to which (and the number of times) users must rotate their head in order to notice what they are walking past. The geometric field of view (GFOV) of VEs may be made greater (or smaller) than the physical field of view (PFOV) by altering the viewing parameters of the VE, but this produces a distortion. For example, setting the GFOV to be greater than the PFOV produces the effect of looking through a wide-angle camera lens, and makes objects seem farther away and smaller than they actually are. A small GFOV (e.g., ten degrees) ad-

versely affects the user's sense of presence in VEs when compared with larger GFOVs (e.g., fifty or ninety degrees) (Hendrix & Barfield, 1996).

When users navigate VEs with displays that do not provide peripheral vision, they sometimes accidentally travel past their targets (Ruddle, Payne, & Jones, 1998). This situation is less likely to occur when people use displays such as CAVEs, because these displays allow people to use their normal (real-world) FOV. Peripheral vision may be simulated with immersive or desk-top displays by providing peripheral viewports at the sides of the displays. (See Slater & Usoh, 1993.)

2.3 Perceptual Differences

Other, more general, perceptual differences seem to occur between different types of display. In one case, design engineers reported having a more correct understanding of the actual size of a CAD model when they used a BOOM than when they used a desk-top display (Mizell et al., in press). A user's sense of presence may also vary between "being inside" immersive VEs and "looking into" desk-top VEs, but the effect of presence on the user's ability to navigate in VEs remains to be investigated.

3 Experiment

The following study investigated differences that occurred when participants navigated large-scale, virtual buildings while using an HMD and a desk-top display. Participants were run individually and were told that the study was being performed to investigate the development of people's spatial knowledge when they navigated VEs. A repeated-measures design was used, with each participant first navigating around one virtual building four times while using one type of display (e.g., the HMD), and then navigating around the other virtual building four times using the other display (e.g., the desk-top). Data from other studies such as Ruddle et al. (1997) have shown that this type of design is a practical way of controlling for the differences that occur in individuals' ability to navigate VEs.

3.1 Method

3.1.1 Participants. A total of twelve participants (ten men and two women) took part in the experiment. They were all psychology undergraduates who volunteered for the experiment and received course credit for their participation. They were randomly assigned to one of four groups, and these were used to counterbalance the order in which the displays were used (half of the participants used the desk-top display first and the other half used the HMD first), the order in which the buildings were navigated, and the different building/ display combinations.

3.1.2 VE. The experiment was performed on a Silicon Graphics Crimson RealityEngine, running a C++ Performer application that we designed and programmed. The desk-top VE was displayed using a 21-inch (53 cm) monitor (1280 × 1024 resolution). It was positioned on a desk in front of a participant and viewed from a “normal” distance (approximately 60 cm). Participants were not restrained in any way. The HMD was a Virtual Research VR4 (247 × 230 resolution), and head-tracking was performed using a Polhemus FAS-TRAK sensor and the MR Toolkit (Green, 1995). Both types of VE used the same application update rate (20 Hz) and a nonstereo graphics input.

The two virtual buildings were designed so that they each had a similar number of rooms, and the same number and complexity of navigation decision points (corridor intersections). (See Figure 1.) Each building contained a lobby, and five named rooms (e.g., the seminar room) that were filled with 3-D models of characteristic furniture to enable their easy identification. The remainder of each building was split into 75 (Building 1) or 71 (Building 2) empty rooms of roughly equal size. Each VE was texture-mapped. (A typical interior view is shown in Figure 2.) Apart from features caused by the building’s shape and the furniture in the named rooms, there were no landmarks (e.g., plants and pictures).

The user interface was similar for both displays and was mastered without difficulty by all the participants. A mouse and keyboard was used with the desk-top display, and a handheld button box, together with physical head

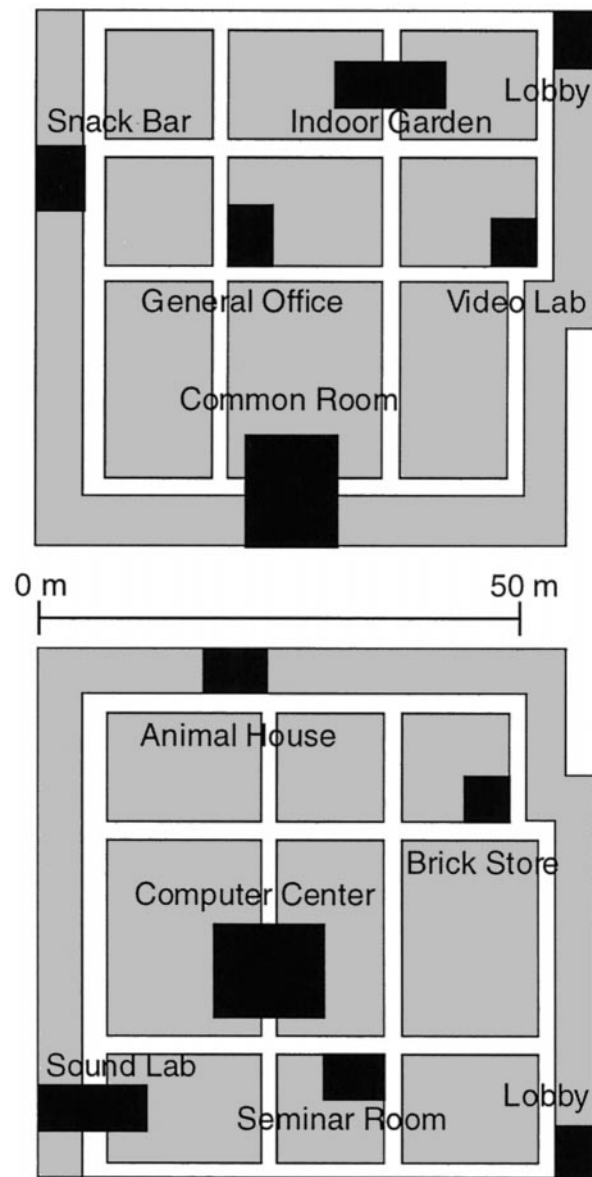


Figure 1. The layout of Building 1 (above) and Building 2 (below). The target locations are black, other rooms are gray, and the corridors are white.

and body movements, was used with the HMD. The controls that are described in the following sections refer to the desk-top display. In general, the HMD controls substituted buttons for keys on the keyboard.

Participants’ viewing and movement directions were decoupled, and this allowed them to look around easily while traveling in a straight line down corridors in the




Figure 2. A view inside Building 1. The view is from beside the video lab, looking towards the snack bar.

VEs. When using the desk-top display, participants controlled their view direction in a combination of two ways: (1) by moving the mouse from side to side, the view direction could be changed by ± 45 degrees, and (2) by holding down the left or right mouse button, a full 360-degree rotation could be performed. When using the HMD, the participants' view direction was determined by their physical head and body movements and was recorded by the Polhemus sensor. The sensor was used to provide two rotational degrees of freedom (DOFs) (heading and pitch), but only one rotational DOF (heading) was used with the desk-top display. This simplified the rotational movements made when using the desk-top display, and we do not believe the difference between each display's number of rotational DOFs caused any of the differences recorded in the experimental data.

Participants controlled their speed and altered their direction of movement by using five keys on the keyboard. Four of the keys allowed participants to slow down, stop, speed up, and move at the maximum allowed speed (4.8 km/h). The fifth key changed the participants' direction of movement to the current view direction. At all times, a green triangle was projected at foot level to indicate the participants' current direction of movement. Participants were prevented from walking

through walls by a collision-detection algorithm, and doors opened automatically when approached.

Both displays used the same geometric, horizontal FOV (ninety degrees), which was chosen so that participants could "stand" at a corridor junction and simultaneously see down both corridors. The physical, horizontal FOVs were approximately forty degrees (desk-top; the angle subtended by a 21-in. monitor when seen from a normal viewing distance) and fifty degrees (HMD). It should be noted that, when using the HMD, this gave rise to a conflict between the apparent direction of objects in the VE and the amount a participant needed to physically turn their head/body in order to face the object (consider an object displayed at the right-hand edge of the HMD; it would appear to be 25 degrees to the right, but to face it a participant would have to turn 45 degrees to the right). However, none of the participants commented about this conflict. The height above the building "floor" at which viewing took place (effectively a participant's virtual "eye" height) was set to be the same as each participant's actual eye height. When using the HMD, participants stood inside a circular barrier (900 mm high and 1 m in diameter) that was similar to the barriers found in some immersive VE games arcades and exhibition stands. It prevented participants from moving outside the working range of the Polhemus sensor and helped to keep the HMD cable away from participants' feet.

3.1.3 Procedures. The familiarization and test procedures were the same for each participant. First, a participant was familiarized with the desk-top VE controls using a simple, practice, virtual building that contained a figure-eight arrangement of corridors and two rooms. Then, they were familiarized with the immersive VE controls using the same practice building. Participants were allowed an unlimited amount of time for the familiarization process, and it was usually completed in 45 to 60 minutes.

Each participant underwent two tests (one in each virtual building and one using each display) that each had the same four stages. The order in which participants used the two displays and the allocation of the displays to the buildings was counterbalanced. Each test

was completed during a single day. In each stage, a participant started in the lobby, visited the five named rooms in a specific order, and then returned to the lobby. In each particular building test stage, all the participants visited the rooms in the same order, but the order was varied across stages. The minimum possible distance traveled in each test stage was approximately 280 m.

During Stage 1, the experimenter read out directions describing the shortest route from the participant's current location to the next target location (i.e., one of the named rooms or the lobby). An example route description was "turn right out of the door, second left, and go through the door at the end." The participant traveled to the location and, if they made an error, they were corrected immediately. When they reached the location, they pressed the "y" key. This triggered the display of a message on the screen, specifying the name of the next location to be visited. The message was removed after a few seconds, but could be redisplayed at any time if the participant pressed the "h" key. Then, the experimenter read out the description of the shortest route to the next location. This sequence of events continued until the participant had successfully visited each of the five target rooms and returned to the lobby.

In Stage 2, the participant followed verbal descriptions of the shortest route to the five named rooms, but was told to find their own way back to the lobby, for which the following "2.5-minute rule" applied. If, after 2.5 minutes, the participant had not reached the lobby, the experimenter gave verbal instructions describing the shortest route to the lobby, which the participant then followed. However, if, after 2.5 minutes, the participant was traveling directly towards the lobby but had not yet arrived, they were allowed to continue unaided but were immediately given verbal instructions if they deviated from the shortest route.

During Stage 3, the participant had to find their own way to the five rooms and the lobby, but the 2.5-minute rule was applied for each of these locations to place a limit on the amount of time participants spent inside the VE. Stage 4 was identical to Stage 3, except that, when the participant pressed the "y" key to indicate their arrival in each of the five rooms, the VE software moved them to the center of the room, and then the participant

made estimates of direction and of straight-line distance (respectively, the VE-orientation and VE-straight data) to the other four rooms. To make the direction estimates, participants rotated their direction of view until they thought they were facing directly toward the appropriate room and indicated this by pressing the "y" key, which caused the view direction to be recorded. For each distance estimate, a Motif window was displayed into which either the participant typed their estimate (desk-top) or the experimenter typed the participant's spoken estimate (HMD). The Motif window was visible in the HMD.

A five-minute break was taken between each test stage. To help alleviate eyestrain during the HMD tests, participants also took a five-minute break approximately every fifteen minutes (if the stage lasted longer than fifteen minutes, which Stage 4 always did). After completing the HMD test, participants completed a short questionnaire (see Annex 1) that was based on the questionnaire used in the NASA Hubble Space Telescope study (Loftin & Kenney, 1994) and which was primarily designed to investigate any side-effects felt by the participants. Three participants withdrew while using the HMD, citing nausea, and were replaced.

3.2 Results

3.2.1 Navigation and Spatial Knowledge. Each participant's route-finding ability was measured both in terms of the distance that they traveled between the locations and the time that they took. The distance-traveled metric was calculated by computing the distance a participant traveled in each test stage as a percentage in excess of the minimum possible distance (PE-distance). The time-taken metric (PE-time) used a similar calculation, with the minimum possible time calculated by dividing the minimum possible distance by the maximum speed of movement (4.8 km/h). Participants' appreciation of relative distance in the buildings was calculated by correlating their VE-straight distance estimates with the corresponding actual distances. The distribution of this correlation was then normalized using Fisher's *r*-to-*z* transformation. The accuracy of participants' direction estimates was determined by calculating the




Figure 3. Mean percentage extra time (PE-time) taken by participants when using the HMD and the desk-top display. Error Bars indicate MSE.




Figure 4. Mean percentage extra distance (PE-distance) traveled by participants when using the HMD and the desk-top display. Error bars indicate MSE.

mean angular error of their VE-orientation estimates. The order in which participants used the two types of display was treated as a between-groups variable, but this was not significant for any of the analyses.

Participants' PE-time data for the unguided stages (Stages 3 and 4) were analyzed using a repeated-measures analysis of variance (ANOVA). Figure 3 shows that, as expected, participants navigated significantly more quickly in Stage 4 than in Stage 3 ($F(1, 10) = 8.07, p < 0.05$). There was no main effect of display type ($F(1, 10) = 2.43, p > 0.05$), but planned contrasts showed that participants navigated significantly more quickly with the HMD than with the desk-top display in Stage 3 ($F(1, 10) = 10.01, p < 0.05$). The difference for Stage 4 was not significant, and there were no significant interactions.

Participants' PE-distance data for the unguided stages were also analyzed using a repeated-measures ANOVA. Figure 4 shows that, contrary to expectations, there were no main effects for stage number ($F(1, 10) = 2.70, p > 0.05$), or for the display type ($F(1, 10) = 0.12, p > 0.05$), and there were no significant interactions. Planned contrasts showed that there were no significant differences between the two types of display for Stage 3 or Stage 4.

The accuracy of participants' survey knowledge was assessed using their VE-straight and VE-orientation estimate data. A repeated-measures ANOVA of the data, illustrated in Figure 5, showed that participants made




Figure 5. Participants' mean straight-line distance (VE-straight) correlations when using the HMD and the desk-top display. The data is transformed from participants' mean Fisher's z values. Error bars indicate MSE.

significantly more accurate estimates of relative straight-line distance when they used the HMD than when they used the desk-top display ($F(1, 10) = 5.77, p < 0.05$). There was no significant difference in the absolute percentage error of participants' straight-line distance estimates (HMD, $M = 54\%$; desk-top, $M = 77\%$), and participants showed no consistent tendency to under- or overestimate the distances with either type of display.

The trend found in the data for the PE-time and relative straight-line distance estimates (which suggests that participants learned some components of spatial knowledge more quickly with the HMD) was also found in the




Figure 6. Participants' mean VE-orientation direction-estimate errors when using the HMD and the desk-top display. Error bars indicate MSE.

direction estimate data. (See Figure 6). However, a repeated-measures ANOVA showed that the difference between the two types of display was not significant ($F(1, 10) = 4.43, p > 0.05$).

3.2.2 Behavioral Differences. Participants aided their navigation by looking around at decision points and glancing into rooms as they passed. As has already been noted, one of the primary differences between HMDs and desk-top displays lies in the mechanism that users use to control their direction of view (head and body movements when using an HMD but, typically, a mouse when using a desk-top display). Participants' mean head movement rates (the rates at which participants looked around while navigating the VEs) were measured by calculating the total movement of their head during each test stage (in degrees) and dividing by the length of time they took to complete the stage (excluding the time taken to perform the distance and direction estimates and any rests taken to relieve eye-strain). These data were then filtered to remove head movements smaller than the filter amount (e.g., thirty degrees) from the calculation. Figure 7 shows that participants looked around more when using the HMD than when using the desk-top display. To see into a room as they approached, participants had to look approximately thirty degrees to one side. A repeated-measures ANOVA, computed using a thirty-degree filter, showed that participants looked around significantly more when they used the HMD than when they used




Figure 7. Participants' mean head movement rates when using the HMD and the desk-top display. The rates are calculated using five-degree filter increments.

the desk-top display ($F(1, 11) = 61.49, p < 0.0001$).

When using the desk-top display, participants often stopped before looking to each side, either to look into a room or to look down a corridor. Another repeated-measures ANOVA, illustrated in Figure 8, showed that participants spent less time stationary when using the HMD than when using the desk-top display ($F(1, 11) = 14.44, p < 0.005$), and planned contrasts showed that this difference was significant for both Stage 3 ($F(1, 11) = 26.99, p < 0.0005$), and Stage 4 ($F(1, 11) = 18.67, p < 0.005$).

3.2.3 Questionnaire. Data from the questionnaire are summarized in Table 1. All twelve of the participants who completed the study reported at least two "slight" side effects (Annex 1, Question 2). Three participants withdrew from the study while using the HMD. One of these participants withdrew during the HMD practice and reported "moderate" nausea and dizziness. The other two participants withdrew during the first (guided) HMD test stage and reported "severe" nausea and either "moderate" or "severe" dizziness. (See Table 2.) None of the participants who either completed the study or withdrew reported experiencing doublevision. The evaluative questions (question 4 through question 9) indicated that all participants who completed the study enjoyed taking part. However, partici-




Figure 8. Mean amount of time that participants spent stationary when using the HMD and the desk-top display. Error bars indicate MSE.

Table I. Physical, HMD Side Effects of the Twelve Participants Who Completed Study

Side effect	None	Slight	Moderate	Severe
Nausea	8	2	2	0
Dizziness	4	7	1	0
Disorientation	0	9	2	1
Eyestrain	2	6	4	0
Sweating	6	4	2	0
Vertigo	11	1	0	0
Blurred vision	5	5	2	0
Headache	8	3	1	0
General discomfort	6	2	4	0
Claustrophobia	11	1	0	0
Other	12	0	0	0

participants were not asked to complete the questionnaire after using the desk-top display and, therefore, it was not possible to compare these data between the two display types.

3.3 Discussion

The overall goal of the study was to perform a baseline investigation that compared navigation and spatial knowledge acquisition in large-scale virtual buildings using an HMD and a desk-top display. Each participant spent an average of thirty minutes traveling approximately 1.5 km during the unguided stages of each test.

Table II. Physical, HMD Side Effects of the Three Participants Who Withdrew From the Study

Side effect	None	Slight	Moderate	Severe
Nausea	0	0	1	2
Dizziness	0	0	2	1
Disorientation	0	2	1	0
Eyestrain	0	3	0	0
Sweating	1	2	0	0
Vertigo	0	3	0	0
Blurred vision	3	0	0	0
Headache	0	1	2	0
General discomfort	1	2	0	0
Claustrophobia	2	1	0	0
Other	3	0	0	0

On average, participants navigated the virtual buildings twelve percent quicker when using the HMD. However, participants did not travel a correspondingly shorter distance when using the HMD. (In fact, they traveled an average of three percent farther, but the difference was not statistically significant.) Instead, the majority of participants' increased speed was derived from changes in behavior between the two display types. When using the desk-top display, participants often stopped before altering their direction of view to look into rooms to see if they contained any furniture or to look down corridors at junctions. By contrast, participants spent approximately eight percent less time stationary when using the HMD and "looked around" more while they were moving.

One explanation for this behavioral difference may be that the HMD provided an interface in which changes in view direction were natural (i.e., head and body movements) and required less effort (e.g., quick glances, rather than holding down a mouse button). This explanation is supported by data that showed that the rate at which participants looked around was significantly greater when they used the HMD than when they used the desk-top display. (See Figure 7.) The pattern of participants' head movements also differed slightly, with 66% of participants' view direction changes being greater than ninety degrees when using the desk-top display, but a more gradual range of movements being used with the

HMD (57% greater than ninety degrees). This suggests that people who use HMDs in VEs are less likely to travel past something (for example, an object which is set to one side) without noticing it. By contrast, informal observation suggests that participants sometimes developed “tunnel vision” when they used the desk-top display, traveling past their desired location without modifying their view direction to see it.

Participants developed a significantly more accurate sense of relative straight-line distance when they used the HMD than when they used the desk-top display. Knowledge of the relative distances between locations is an important component of spatial knowledge, but the reason for the difference is unclear. However, it may be related to general perceptual differences that are caused by being “inside” immersive VEs and the more realistic perception of size that was found by Mizell et al. (in press).

Witmer and Kline (1998) found that people tend to underestimate distances in VEs. In the present study, participants showed no consistent tendency to either under- or overestimate distances with either type of display, but factors that are likely to have affected these data are the discrepancy between the GFOV of the VEs and PFOVs of the displays, and the ability of individual participants to estimate distances in the real world. Contrary to expectations, the accuracy of participants’ direction estimates was not significantly enhanced by the kinesthetic feedback that was provided when the HMD was used.

The present study is only one step toward our overall goal of understanding the effects of using immersive displays to navigate VEs. Further studies are planned in which participants will navigate VEs for longer periods of time. This will allow us to investigate whether or not the time advantage of using an HMD is present only when people initially navigate a VE, as suggested by the PE-time data. (There was a significant difference between the displays for Stage 3 but not for Stage 4.) Other planned changes from the method used for the present study include using more-complex virtual buildings, with an increased time penalty for poor navigational learning, the use of a between-participants design to overcome the learning effect that was probably present when participants navigated the second building

(see Ruddle et al., 1998), and the use of virtual buildings that have corridors that are curved or intersect at angles other than ninety degrees. Other studies are required to investigate the effects of simulated-walking interfaces, treadmills, and roving interfaces (i.e., those that use extended-range sensors to allow people to walk around inside VEs).

Finally, landmarks have been shown to aid navigation in VEs (Darken & Sibert, 1996; Ruddle et al., 1997). Participants in the present study looked around more when they used the HMD. Therefore, it is possible that landmarks will be noticed more when an HMD is used and, as a result, will provide a greater aid to navigation.

References

- Chung, J. C. (1992). A comparison of head-tracked and non-head-tracked steering modes in the targeting of radiotherapy treatment beams. *Proceedings of the 1992 ACM Symposium on Interactive 3D Graphics* 193–196.
- Darken, R. P., Cockayne, W. R., & Carmein, D. (1997). The omni-directional treadmill: A locomotion device for virtual worlds. *Proceedings of ACM User Interface Software & Technology Conference (UIST '97)* (pp. 213–221). New York: ACM.
- Darken, R. P., & Sibert, J. L. (1996). Navigating large virtual spaces. *International Journal of Human-Computer Interaction*, 8, 49–71.
- Green, M. (1995). *The MR Toolkit Version 1.4* [Computer program]. Department of Computing Science, University of Alberta, Canada.
- Hartshorn, M., & Hubbard, E. (1993). Interactive protein modelling. In K. Warwick, J. Gray, & D. Roberts. (Eds.), *Virtual reality in engineering* (pp. 25–32). London: IEE.
- Hendrix, C., & Barfield, W. (1996). Presence within virtual environments as a function of visual display parameters. *Presence: Teleoperators and Virtual Environments*, 5, 274–289.
- Loftin, R. B., & Kenney, P. J. (1994). Virtual environments in training: NASA’s Hubble space telescope mission. *Proceedings of 16th Interservice/Industry Training, Systems & Education Conference (IITSEC '94)*. <http://www.vetl.uh.edu/Hubble/hubble.html>.
- Mizell, D., Jones, S., Jackson, P., & Pickett, D. (in press). Is VR better than a workstation? A report on human performance experiments in progress. In M. Göbel (Ed.), *Virtual Environments*. Springer-Verlag.

- Pausch, R., Shackelford, M. A., & Proffitt, D. (1993). A user study comparing head-mounted and stationary displays. *Proceedings of IEEE Symposium on Research Frontiers in Virtual Reality*, 41–45.
- Presson, C. C., & Montello, D. R. (1994). Updating after rotational and translational body movements: Coordinate structure of perspective space. *Perception*, 23, 1447–1455.
- Rieser, J. J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 1157–1165.
- Rieser, J. J., Lockman, J. L., & Pick, Jr., H. L. (1980). The role of visual experience in knowledge of spatial layout. *Perception & Psychophysics*, 28, 185–190.
- Ruddle, R. A., Payne, S. J., & Jones, D. M. (1997). Navigating buildings in “desk-top” virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied*, 3, 143–159.
- Ruddle, R. A., Payne, S. J., & Jones, D. M. (1998). Navigating large-scale “desk-top” virtual buildings: Effects of orientation aids and familiarity. *Presence: Teleoperators and Virtual Environments*, 7, 179–192.
- Slater, M., Alberto, C., & Usoh, M. (1994). In the building or through the window? An experimental comparison of immersive and non-immersive walkthroughs. *Proceedings of Virtual Reality Environments in Architecture and Design*.
- Slater, M., & Usoh, M. (1993). Simulating peripheral vision in immersive virtual environments. *Computers & Graphics*, 17, 643–653.
- Slater, M., Usoh, M., & Steed, A. (1995). Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Transactions on Computer-Human Interaction*, 2, 201–219.
- Tachi, S., & Yasuda, K. (1994). Evaluation experiments of a teleexistence manipulation system. *Presence: Teleoperators and Virtual Environments*, 3, 35–44.
- Waller, D., Hunt, E., & Knapp, D. (1998). The transfer of spatial knowledge in virtual environment training. *Presence: Teleoperators and Virtual Environments*, 7, 129–143.
- Witmer, B. G., & Kline, P. B. (1998). Judging perceived and traversed distance in virtual environments. *Presence: Teleoperators and Virtual Environments*, 7, 144–167.

Annex 1: Questionnaire

Congratulations! You have survived the immersive/ helmet phase of this experiment. Thank you very much for your time and effort. The following questions refer specifically to your experience of virtual reality using the helmet and handheld control box. Please answer each question and, if you wish to add any comments, please feel free to do so.

1. Have you ever used a virtual reality system of any type before participating in this experiment?
yes () no ()
2. If you experienced any of the following side effects while inside the virtual environment, please indicate the level of discomfort (mark as many as apply).

SIDE EFFECT	NONE	SLIGHT	MODERATE	SEVERE
A. nausea	()	()	()	()
B. dizziness	()	()	()	()
C. disorientation	()	()	()	()
D. eyestrain	()	()	()	()
E. sweating	()	()	()	()
F. vertigo	()	()	()	()
G. blurred vision	()	()	()	()
H. headache	()	()	()	()
I. general discomfort	()	()	()	()
G. claustrophobia	()	()	()	()
K. Other?	()	()	()	()

Annex 1: Questionnaire (Continued)

3. Did you experience double-vision? yes () no ()

4. How enjoyable/interesting was your interaction with the virtual environment?
(Please indicate by circling the corresponding number rating).

1	2	3	4	5
Boring				Captivating

Why/ Comments?

5. Was it easy for you to become "absorbed" in the virtual environment?

1	2	3	4	5
Difficult				Very easy

Why/ Comments?

6. How difficult was it for you to maneuver through the environment with the handheld controls?

1	2	3	4	5
Very difficult				Very easy

Why/ Comments?

7. How often did you have to spend time thinking about the hand controls with respect to the actual task?

1	2	3	4	5
Never				Very frequently

Why/ Comments?

8. How much did using the hand controls distract you from concentrating on the virtual environment task?

1	2	3	4	5
Not at all distracting				Extremely distracting

Why/ Comments?

9. How distracting was it for you that there was another person in the room whilst you were immersed in the virtual environment?

1	2	3	4	5
Not at all distracting				Extremely distracting

Why/ Comments?