

Empirical Comparisons of Virtual Environment Displays

Doug A. Bowman¹, Ameya Datey¹, Umer Farooq¹, Young Sam Ryu², and Omar Vasnaik¹

¹Department of Computer Science

²The Grado Department of Industrial and Systems Engineering
Virginia Polytechnic Institute and State University

Contact Author's Address:

Doug Bowman
Dept. of Computer Science (0106)
660 McBryde Hall
Virginia Tech
Blacksburg, Virginia 24061 USA
Voice: +1 (540) 231-2058
Fax: +1 (540) 231-6075
{bowman, adatey, ufarooq, yryu, ovasnaik}@vt.edu

Abstract

There are many different visual display devices used in virtual environment (VE) systems. These displays vary along many dimensions, such as resolution, field of view, level of immersion, quality of stereo, and so on. In general, no guidelines exist to choose an appropriate display for a particular VE application. Our goal in this work is to develop such guidelines on the basis of empirical results. We present two initial experiments comparing head-mounted displays with a workbench display and a four-sided spatially immersive display. The results indicate that the physical characteristics of the displays, users' prior experiences, and even the order in which the displays are presented can have significant effects on performance.

Keywords and Phrases

3D displays, virtual environments, 3D interaction, empirical evaluation, head-mounted displays, spatially immersive displays, tabletop displays.

Empirical Comparisons of Virtual Environment Displays

Abstract

There are many different visual display devices used in virtual environment (VE) systems. These displays vary along many dimensions, such as resolution, field of view, level of immersion, quality of stereo, and so on. In general, no guidelines exist to choose an appropriate display for a particular VE application. Our goal in this work is to develop such guidelines on the basis of empirical results. We present two initial experiments comparing head-mounted displays with a workbench display and a four-sided spatially immersive display. The results indicate that the physical characteristics of the displays, users' prior experiences, and even the order in which the displays are presented can have significant effects on performance.

1 Introduction

Virtual environments (VEs) have often been described as a “technology in search of a problem.” Over the last decade, researchers have shown the effectiveness of VEs for tasks such as psychiatric therapy [7] and training [14], and many more application areas seem on the verge of acceptance. However, a new problem now presents itself: designers of VE applications must choose from a wide array of possible technology. Should a haptic display be used? If so, which one? What input devices will be most effective? Is head tracking required? Very little guidance exists to aid designers in these choices.

In particular, all VEs must use some sort of visual display device. Indeed, many VEs are characterized by the display device they use (e.g. a “CAVE⁺ application” or a “head-mounted display system”). We have some intuitive understanding of the relative strengths and weaknesses of some of the most common VE display devices (see section 3), but there are few empirical results to back up these feelings. As Fred Brooks noted in his review of the current state of the art in VEs, one of the most crucial challenges in the field is “choosing which display best fits each application” [2]. Our goal is to develop guidelines that create a mapping between an application and a display, or, more specifically, between an application's requirements and a display. For example, a useful result might be “If an application requires the highest level of perceived immersion, use a spatially immersive display.”

⁺ Throughout this paper, we use the term ‘CAVE’ to refer to a spatially immersive display using multiple flat surrounding screens. ‘CAVE’ is a trademark of Fakespace Systems, Inc.

We present two initial studies aimed at producing this type of guideline. First, we compare head-mounted displays (HMDs) to a workbench display to examine the differences between an egocentric and exocentric point of view for a search task. Second, we compare HMDs to a spatially immersive display (SID) to study the effects of manual rotation vs. natural rotation of the user’s viewpoint. These experiments provide some insight into the subtle differences between the display types, and suggest many interesting lines of future work.

2 Related work

There has been a great deal of effort in the VE community aimed at developing new displays [e.g. 13] and improving existing display types [e.g. 9]. However, there is little work that objectively compares different VE displays. A SIGGRAPH ’96 panel [11] asked prominent VE researchers to make a case for HMDs or SIDs as the “future of virtual reality.” Panelists in general proposed that displays should be chosen based on the tasks and requirements of a particular application, but no empirical results were given.

Our first experiment looks at a naïve search task in a VE. Darken [5] has studied this task extensively as it relates to wayfinding in VEs. He concludes that some search tasks are inherently egocentric (first person), while others require information from an exocentric (third person) reference frame. We test a similar hypothesis in our comparison. Our second experiment looks at natural (physical) viewpoint rotation vs. manual rotation. Research suggests that physical motion allows better estimates of the magnitude of rotation [1]. Chance et al. [3] asked subjects to traverse a virtual corridor similar to the ones used in our experiment. They found that natural turns produced significantly higher levels of spatial orientation than manual turns, a fact which has important implications for the choice of VE displays given the results of our experiment.

3 Display devices

Three-dimensional (3D) visual display technology for use in virtual environments has become quite varied. For many years, the head-mounted display (HMD) was considered the most common VE display. A range of new display devices has appeared recently that promises to be useful for VE systems. These include workbench displays, desktop stereoscopic displays, retinal displays, and spatially immersive displays (SIDs). In this section we explore the characteristics of the three display types compared in our experiments.

3.1 Head-mounted displays

The HMD (figure 1) is considered the canonical VE display. This device usually consists of two LCD or CRT screens that are mounted on a helmet-like device so that they are fixed relative to the wearer's eye position. This device portrays the virtual world by obtaining the user's head position and orientation from a tracking system. Imagery is magnified with a set of optics.

There are several attributes of HMDs that may have an effect on a user's performance [6]. First, HMDs may present biocular (same image to both eyes) or stereoscopic (different images to each eye) imagery. Second, HMDs come in a wide range of resolutions. Resolution usually trades off with field of view (FOV), which is measured in degrees of horizontal visual angle. A lower FOV results in "tunnel vision" and may decrease immersion, but higher FOVs involve spreading out the available pixels, which can decrease resolution and introduce distortion. Finally, there are ergonomic issues related to HMDs such as the display's size and weight and the ability to adjust various visual parameters.



Figure 1. User wearing a head-mounted display (HMD)

3.2 Tabletop displays

Tabletop VE displays such as the Responsive Workbench [10] are achieved by projecting stereoscopic imagery onto a planar surface in the physical world (figure 2). This is usually done using a graphics projector and a screen of frosted glass. The user sees the stereo images by wearing a pair of active stereo glasses, which sync with the time-multiplexed graphics to ensure that each eye sees only the image intended for it. The user's head is also tracked so that the proper perspective can be displayed.

Workbench displays have several perceived advantages over HMDs. The headgear is much lighter and less cumbersome. Also, the 3D imagery may appear to "sit" upon the physical surface of the workbench much like an architect's scale model. Unlike HMDs, several users can view the imagery at once, although the stereo and perspective are generally only correct for a single head-tracked user. Finally, since the

imagery is on a physical surface with a small area, the user can generally reach out to touch 3D objects directly, interacting with a stylus, pinch gloves, or other devices. However, tabletop displays are not generally intended to immerse the user within the virtual world.



Figure 2. Interacting with 3D objects on a tabletop display

3.3 Spatially immersive displays

Rather than using a single set of displays that follow the user's head movements, spatially immersive displays (SIDs) use multiple displays that surround most of the user's field of view (figure 3). The most common example of a SID is the CAVE [4]. CAVEs generally use four large back-projected screens set at right angles (three walls and a floor), onto which stereo graphics are projected. The user again wears stereo glasses and a head tracker. In this way, a 3D environment can be displayed to the user that is more immersive than a tabletop display.

The increased immersion is the most attractive feature of CAVEs and other SIDs. The user is also less encumbered than with HMD systems. In addition, the FOV in a CAVE (assuming the user looks at the front wall) can be more than 180 degrees, which matches the eye's perceptual limit (this FOV is usually not achieved, however, due to the limitation of the FOV through the stereo glasses).

On the other hand, CAVEs may be less immersive than HMDs because of the two missing surfaces of the cube (back wall and ceiling) that break the illusion of being within a 3D space. In a standard CAVE, the

user must have some method for rotating the environment in order to see objects behind them. Also, the projected graphics tend to be less bright, and thus require a darkened room. Finally, CAVEs are currently many orders of magnitude more expensive than HMD-based systems, making them impractical for many potential users.



Figure 3. User within a spatially immersive display (SID)

4 Egocentric vs. exocentric search

Our first experiment compared an HMD and a tabletop display. HMDs provide an egocentric point of view into the virtual world, while tabletop displays, depending on the orientation of the screen, may provide a more exocentric, or “god’s-eye” viewpoint. This implies that HMDs are appropriate for egocentric tasks, while tabletop displays map well to exocentric tasks. The problem lies in the definition of exocentric and egocentric tasks. We define an exocentric task as one in which the user’s reference frame is not the same as the environment’s reference frame – the user reaches into the world. An egocentric task is one in which the user and the environment share a reference frame – the user is immersed within the world.

A search task could be done in either an exocentric or egocentric fashion (in fact, it is difficult to imagine tasks that cannot be done in both ways). We chose to examine a difficult naïve search (the user has no prior knowledge of the location of the target, and the target is well-hidden) in exploring the performance of these two displays. Our hypothesis was that this type of search task is most efficiently performed from

an egocentric point of view, and is therefore most appropriate for an HMD. The user's objective in the experiment is to find the hidden object inside a warehouse in the shortest possible time.

4.1 Method

4.1.1 Subjects

Fifteen people, both students and non-students participated in the experiment. Participants had varying degrees of computer skills, and three participants had used VE display devices before.

4.1.2 Apparatus and implementation

The HMD in the experiment was the Virtual Research V8. It supports a resolution of 640x480, with a sixty-degree diagonal field of view. The HMD presented biocular images to the user. We used an Intersense IS-900 VET tracking system to track the head and one hand of the user. The hand tracker was a stylus with a button that could be depressed by the user to navigate through the environment in the direction he was pointing. The HMD application was developed using the SVE toolkit [8] and ran on a Windows NT PC.

For the tabletop display, we used a Fakespace Immersive Workbench, with the screen lying completely flat (parallel to the floor). Users wore Crystal Eyes stereo glasses, and held a tracked joystick in one hand. Tracking was done by a Polhemus Fastrak system. The joystick was used to rotate the environment about its vertical axis. To rotate the world, users depressed a joystick button and turned the joystick in the desired direction. To view the world more closely, users simply leaned forward so that the virtual model was nearer to their viewpoint. The workbench application was developed using the DIVERSE toolkit (see www.diverse.vt.edu).

4.1.3 Environment

The environment used in this experiment was a warehouse (figure 4). Walls bound the warehouse on all sides, and there are a number of boxes of various sizes scattered randomly. Care was taken to ensure that the target object (a spring) was well hidden and not obvious to the subject at first glance. Two configurations were designed for the experiment with the position of boxes at different locations. For each configuration, the object was hidden in five distinct places, leading to ten trials with each display. The same 3D models were used in both the HMD and workbench applications.

4.1.4 Experimental design

We used a within-subjects design, in which all subjects used both display types. Time for completion of the search task (measured in seconds) was the only dependent variable. The independent variable was

display device (HMD and workbench). Each subject completed ten trials (in randomized order) with each display. The order in which the displays were used was counterbalanced between subjects.




Figure 4. Warehouse environment

4.1.5 Procedure

Subjects were first given a demographic questionnaire. This questionnaire asked for demographic information such as age, gender, and occupation (or major field of study), and for information on the subject's use of computers and prior experience with VEs.

With each display, subjects were instructed in the use of the headgear and input devices, and were given ample time to practice the navigation/rotation techniques and the search task in a sample environment. When the subject was ready, they began the experimental trials. Timing began upon the subject's first movement after the environment was loaded. Subjects were free to use any combination of physical and virtual navigation/rotation techniques to find the target object. When they spotted the object, they reported this verbally, and timing was stopped.

After each set of ten trials, the subjects were given a user comfort questionnaire. This questionnaire elicited subjective ratings of comfort on a ten-point scale for arm strain, hand strain, dizziness, and nausea.

4.2 Results

The average time for task completion with the HMD was 32.6 seconds, and the average time on the workbench was 33.6 seconds. We performed a single-factor analysis of variance (ANOVA) on the data and

found that this difference was not statistically significant ($F(1,28) = 0.022$, $p > 0.5$). This result could indicate that there is no difference in performance between the two displays, that there was too much variance between the average times of the subjects, or that something more subtle was taking place. We found one possible explanation in the order in which subjects used the displays.

4.2.1 Ordering effect

The subjects who started with the HMD had an average of 23.4 and 27.1 seconds respectively on the HMD and the workbench. However, subjects starting on the workbench (we considered only seven subjects in order to obtain an equal sample size) had an average of 35.2 seconds on the HMD and 35.6 seconds on the workbench. In other words, subjects using the HMD first performed better on both displays (figure 5). We performed a two-factor ANOVA using display type and display order as the independent variables, and found that the difference between the two orderings was nearly statistically significant ($F(3,24) = 3.92$, $p < 0.06$).




Figure 5. Effect of display ordering on performance

This result suggests that subjects who started with the HMD were able to perform the task efficiently, and also learned something about the task or the environment that transferred to the workbench. Conversely, subjects starting with the workbench were inefficient with this device, and they developed some bias that caused their HMD performance to be worse than it should have been. We suggest two

possible explanations. First, as we hypothesized, the search task might be most natural from a first-person point of view. Subjects using the HMD first saw the world from a first-person viewpoint, and were therefore able to build up a mental model of the warehouse and the possible hiding places within it based on landmark and route knowledge. As these subjects progressed to the workbench, they were able to use this spatial knowledge and combine it with the map-like character of the workbench display to maintain a high level of performance. This is contrasted with subjects who started with the workbench. It provides a survey view of the environment, but that view is not helpful when the target object is well hidden. Subjects are forced to place their viewpoint within the virtual model to find the object. Moreover, this strategy does not transfer well to the HMD. A progression from the HMD to the workbench mirrors the typical pattern of spatial knowledge acquisition – from landmark and route knowledge to survey knowledge [15].

There is also a possible physical explanation for this result. The best strategy when searching in the HMD is to utilize head tracking heavily to turn in all directions, bend down to look under boxes, and lean to look around objects. On the workbench, the best strategy also uses a great deal of physical movement (placing one's viewpoint within the model and looking at possible hiding places). Subjects who used the HMD first discovered this strategy, and carried it with them to the workbench. Those who started on the workbench did not discover the optimal strategy, since it does not correspond to the way we are used to viewing models on a table. Thus, they remained relatively still, and also tried to use this strategy in the HMD.

Whatever the explanation, the ordering effect is clearly crucial. It also implies that our original hypothesis (that this type of search task is egocentric and therefore more appropriate for an HMD) is correct for users who have not been biased by performing the task on another display. A between-subjects design in which each subject uses only one of the display types should show this conclusively.

In summary, there is a clear difference in the way this difficult search task was performed using the HMD and workbench displays, although this difference was not statistically supported until we looked at the effect of display order. Based on the results of the subjects who used the HMD first, we can say that such a task can be done effectively using either display type (the lower line in figure 5). For naïve users, in general, the HMD provides a more natural point of view for object search. Effective search from an exocentric point of view may require a strategy that is somewhat anti-intuitive.

4.2.2 Other results

There was a positive correlation between subjects' reported level of computer usage "for fun" and performance on the workbench ($r = 0.65$). This may be due to the high levels of spatial ability (and therefore ability to understand the third-person representation of the warehouse as it rotates) developed by frequent computer gamers. It may also reflect on these users' prior experience with using a joystick of the type used with the workbench.

We also found high correlation levels between the subjects' reported levels of hand strain ($r = 0.66$) and arm strain ($r = 0.68$) in the HMD environment and their performance in this environment. Subjects using the HMD and stylus have nowhere to rest their arms, and many of them kept the stylus in front of their faces so that the virtual representation of the stylus (see figure 4) could be seen. The strain produced by this obviously had a negative effect on performance.

5 Natural vs. manual rotation

Most SIDs produce excellent perceived immersion because of their high resolution, excellent stereo, and wide field of view. However, the most common SIDs do not have complete physical immersion, since two sides of the six-sided cube are usually missing. Because of this, users cannot physically turn 360 degrees to view the world all around them. Rather, they must manually rotate the world to see what is behind or above them. In HMDs, on the other hand, perceived immersion may be less, due to a low field of view, lower resolution, and so on, but physical immersion is complete – the user sees the virtual world no matter what direction she looks.

We wanted to test the implications of this tradeoff in the two displays, so we implemented a series of corridors in which users could choose between physical and manual turns. Since manual turns are less natural and reduce spatial orientation, and since they might be less efficient, we expected that CAVE users would only use manual rotation when it was required. We further conjectured that HMD users would not use the manual rotation option frequently. By collecting empirical data about the way the user turns, we could obtain information about the appropriate kind of interaction techniques for navigation in various VE applications based on the type of display.

5.1 Method

5.1.1 Subjects

The experiment was conducted with 18 subjects. Out of these, experiments with two subjects were aborted due to discomfort on the part of the subject. Consequently, 16 subjects (eight male, eight female) completed the experiment.

5.1.2 Apparatus and implementation

The Virtual Research V8 HMD with biocular graphics and the Intersense IS-900 tracking system were also used in this experiment. The hand tracker was the Intersense wand, which has an integrated joystick. The joystick was used to navigate and to perform manual rotations. As in the previous experiment, the HMD application ran on a PC and was developed using SVE.

The SID in this experiment was a Fakespace CAVE. It is a theater 10x10x9 feet, made up of three rear-projection screens for the front, right and left walls and a down-projection screen for the floor. A Silicon Graphics Power Onyx with three Infinite Reality Engines is used to create the imagery that is projected onto the walls and floor. The tracking system and the joystick is identical to the equipment that is used with the HMD setup. The application for this setup was developed using DIVERSE.

5.1.3 Environment

Eight corridors were created using a 3D modeling tool, and were used in both the HMD and CAVE applications. The walls had a texture of a brick and were made higher than the user, so that the user could not see over the walls and get any additional spatial information. Corridors had no choice points, so there was no decision making about the path to be followed. All turns were at right angles. Figure 6 shows a typical view within a virtual corridor.




Figure 6. User's view of a virtual corridor

The subject could press the joystick forward or backward to move forward or backward respectively in the direction he is facing (gaze-directed steering). The subject could turn his head or entire body naturally

while making a turn. We call this *natural rotation*. The subject could also choose to use the joystick to turn the world; this is *manual rotation*. Finally, the subject could combine both the techniques while making a turn. This is *combination rotation*. Figure 7 shows a map-like view of a virtual corridor. In a 4-wall CAVE, the user would be forced to rotate the environment manually at the position marked by a star, assuming the user starts by facing the front wall of the CAVE. This forced manual rotation is referred to as *mandatory manual rotation*.




Figure 7. Top-down view of a virtual corridor with mandatory manual rotation

5.1.4 Experimental design

The experiment used a within-subjects design. The single independent factor in these experiments was the display device (HMD and CAVE). Therefore each subject participated in two sessions, one for each display device. Both sessions were conducted at the same time, one after the other. The order of sessions was alternated; half of the subjects started with the HMD first, and the rest started with the CAVE. The sequence of corridors was counterbalanced by the use of a Latin Square design.

The number of turns using each of the techniques was observed and recorded. Also, the time needed to complete each trial was recorded. These two factors were the dependent variables.

5.1.5 Procedure

The subjects were first given a demographic questionnaire (see section 4.1.5). Subjects were then placed in a practice corridor to get them acquainted with the setup. When the subjects felt comfortable, the experimental trials began. The subjects were told to try and navigate out of the corridor as quickly and efficiently as they could. Upon completion of one corridor, the next corridor was loaded. Subjects completed eight trials using each device, and were free to use natural, manual, or combination rotation at each turn. Evaluators closely watched both the subject and the graphics display to determine the technique used for each turn.

5.2 Results

5.2.1 Rotation types

Figure 8 shows the percentage of the different types of rotations with the two display devices. It is quite clear that the percentage of natural rotations is smaller in the CAVE than the HMD, whereas the percentage of manual rotations is greater in the CAVE than in the HMD. Based on a single-factor ANOVA, we found that display type significantly affected the percentage of natural turns and the percentage of combination turns ($F(1,30) = 7.40, p < 0.02$ and $F(1,30) = 10.09, p < 0.005$ respectively). Combining these results with the results from Chance et al. [3], we can conclude that HMD users are more likely to maintain spatial orientation than CAVE users in an environment where turning is frequent.




Figure 8. Percentage of different types of rotations in HMD and CAVE.

There was a noticeable difference in rotation preferences based on gender. Females seemed to prefer using natural rotation, and tried to use natural rotation as much as possible. Considering only the female subjects, only the natural turns were significantly affected ($F(1,14) = 7.04, p < 0.02$), whereas in the case of male subjects, significant results were observed on combination turns ($F(1,14) = 8.30, p < 0.02$).

The results show that natural rotation is the preferred technique in the HMD. Participants also used manual turns at places where it was not mandatory, especially in the CAVE. The CAVE's missing back wall often came as a surprise to subjects who turned naturally in the CAVE, realized the wall was missing, and then used manual rotation. After such an episode some subjects then continued using manual rotation even when it was not required.

Figure 9 compares the two display devices in terms of *elective* (non-mandatory) manual rotations. These ratios were calculated by dividing the total number of non-mandatory manual rotations by the total

number of possible non-mandatory manual rotations for each corridor (in the HMD, all manual rotations are non-mandatory, and every turn is a possible non-mandatory manual turn; in the CAVE, only those turns where the user had a choice between natural, manual, and combination were considered). The figure shows that in a majority of the corridors, subjects chose to perform manual rotation more often in the CAVE even when it was not required. Overall, however, the difference was small: HMD users elected to use manual rotation 39 percent of the time, while CAVE users elected to do so 41 percent of the time.




Figure 9. Percentage of elective manual rotations in the two display devices for each of the corridors

Although not significant, we also saw a trend towards an ordering effect of the two displays. Subjects who used the HMD first also used more natural turning in the CAVE. Subjects who started with the CAVE preferred more manual rotation even in the HMD.

5.2.3 Time for traversal

Figure 10 shows the evaluation of the time for traversal through the corridors. Time has been normalized to control for the effect of corridor length. It can be seen that the average time in the HMD was less than the average time in the CAVE in most of the corridors. Overall, it took an average of 15.32 seconds to traverse a corridor in the HMD and an average of 16.37 seconds in the CAVE. Although this was a strong trend, there was not a significant difference ($F(1,30) = 3.46, p < 0.08$). In the case of natural turns, the proprioceptive sense of the subject allows fast natural turns without the loss of spatial orientation.

Even though manual turns can be made faster than natural turns, increased speed of manual turns can disorient the user.




Figure 10. Average time (normalized) for traversal through each of the corridors

6 Conclusions and future work

We have presented the results of two empirical evaluations comparing performance using common virtual environment displays. The first experiment showed that subjects performed differently on a difficult search task depending on which display they encountered first. We believe that this ordering effect supports our hypothesis that this task is more appropriate in a first-person display, such as an HMD, and we plan further testing to prove this empirically. The second experiment showed that HMD users are significantly more likely than CAVE users to use natural rotation in a VE. This will produce higher levels of spatial orientation, and may make navigation more efficient.

In the future, we plan further comparative explorations of VE display types, with the goal of producing a set of guidelines facilitating the mapping between applications and displays. The current studies in particular suggest three follow-on experiments. First, we will identify tasks that are more clearly exocentric or egocentric in nature in order to demonstrate their appropriateness for exocentric or egocentric display types, respectively. For example, solving a jigsaw puzzle is a task for which it is difficult to conceive an egocentric strategy – a tabletop display should prove most effective for this task. Second, the missing sides

in SIDs may also have an effect on the sense of presence. We can test this using presence questionnaires [16] or more objective measures, such as memory for object locations. Finally, some authors have suggested that non-isomorphic rotations could be used in the CAVE to allow a 360-degree view even with a missing back wall [12]. It would be instructive to test the effects of such a technique on spatial orientation. There is much work to be done in this area, but the results of such experiments should allow VE developers to choose effective visual display devices for their applications.

Acknowledgments

The authors would like to acknowledge the time and effort of the experimental subjects. They also appreciate the helpful comments given by members of the 3DI research group and the Virtual Environments class. Finally, they thank John Kelso, Lance Arsenault, and Randy Levensalor for their support of the development process.

References

1. Bakker, N., Werkhoven, P., and Passenier, P. Aiding Orientation Performance in Virtual Environments with Proprioceptive Feedback. In *Proceedings of IEEE Virtual Reality Annual International Symposium*, Atlanta, Georgia: IEEE Computer Society Press, 1998, pp. 28-35.
2. Brooks, F., *What's Real About Virtual Reality?* IEEE Computer Graphics & Applications, 1999. **19**(6): p. 16-27.
3. Chance, S., et al., *Locomotion Mode Affects the Updating of Objects Encountered During Travel*. Presence: Teleoperators and Virtual Environments, 1998. **7**(2): p. 168-178.
4. Cruz-Neira, C., Sandin, D., and DeFanti, T. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. In *Proceedings of ACM SIGGRAPH*, 1993.
5. Darken, R. and Cevik, H. Map Usage in Virtual Environments: Orientation Issues. In *Proceedings of IEEE Virtual Reality*, 1999, pp. 133-140.
6. Davis, E., *Visual Requirements in HMDs: What Can We See and What Do We Need to See?*, in *Head-Mounted Displays: Designing for the User*, J. Melzer and K. Moffitt, Editors. 1996, McGraw-Hill: New York. p. 207-249.
7. Hodges, L., et al., *Virtual Environments for Treating the Fear of Heights*. IEEE Computer, 1995. **28**(7): p. 27-34.

8. Kessler, G., Bowman, D., and Hodges, L., *The Simple Virtual Environment Library: An Extensible Framework for Building VE Applications*. Presence: Teleoperators and Virtual Environments, 2000. **9**(2): p. 187-208.
9. Kijima, R. and Ojika, T. Transition Between Virtual Environment and Workstation Environment with Projective Head Mounted Display. In *Proceedings of IEEE Virtual Reality Annual International Symposium*, Albuquerque, New Mexico: IEEE Computer Society Press, 1997, pp. 130-137.
10. Kruger, W., et al., *The Responsive Workbench: A Virtual Work Environment*. IEEE Computer, 1995. **28**(7): p. 42-48.
11. Lantz, E. The Future of Virtual Reality: Head Mounted Displays versus Spatially Immersive Displays. In *Proceedings of SIGGRAPH*: ACM Press, 1996, pp. 485-486.
12. LaViola, J., et al. Hands-Free Multi-Scale Navigation in Virtual Environments. In *Proceedings of ACM Symposium on Interactive 3D Graphics*, Research Triangle Park, North Carolina: ACM Press, 2001, pp. 9-15.
13. Meyer, M. and Barr, A. ALCOVE: Design and Implementation of an Object-Centric Virtual Environment. In *Proceedings of IEEE Virtual Reality*, Houston, Texas: IEEE Computer Society Press, 1999, pp. 46-52.
14. Tate, D., Sibert, L., and King, T. Virtual Environments for Shipboard Firefighting Training. In *Proceedings of the Virtual Reality Annual International Symposium*, 1997, pp. 61-68.
15. Thorndyke, P. and Hayes-Roth, B., *Differences in spatial knowledge obtained from maps and navigation*. Cognitive Psychology, 1982. **14**: p. 560-589.
16. Witmer, B. and Singer, M., *Measuring Presence in Virtual Environments: A Presence Questionnaire*. Presence: Teleoperators and Virtual Environments, 1998. **7**(3): p. 225-240.