A Method for Quantifying the Benefits of Immersion Using the CAVE

Abstract

Immersive virtual environments (VEs) have often been described as a technology looking for an application. Part of the reluctance from industry to make regular use of immersive VEs stems from the fact that there is no empirical evidence suggesting that immersive applications provide any benefit over non-immersive ones. We present a method in this paper that can be used to statistically support claims for the benefits of immersion. In attempting to quantify the benefits of immersion, we focus on two major components: physical immersion and head-based rendering. Prior experiments intended to demonstrate the benefits of immersive VEs have tended to confound variables or examine aspects other than immersion. Our method using a CAVE permits the independent evaluation of the effects of physical immersion and head-based rendering. The number and location of active display surfaces controls the level of physical immersion, while enabling/disabling the head tracker and permitting freedom of movement control head-based rendering. This approach has been used successfully in our pilot study regarding the benefits of immersion in abstract information visualization. This method can also be used in a wide variety of other experiments, and we present a brief description of some for possible future study.

Keywords

Virtual Environments, Virtual Reality, Immersion, CAVE

Introduction

A large number of immersive VE applications have been proposed or prototyped in the research lab, but only a few have become production applications [Brooks99]. One reason for this lack of commercial acceptance is that there is no empirical evidence that immersive applications provide any benefits over non-immersive applications. Quantifying the benefits of immersion is obviously an important challenge for the VE community. If attractive cost/benefit ratios can be demonstrated for particular tasks and domains, industries will be much more likely to invest in VE technology. Our work is aimed at producing empirical evidence of increased user performance, usability, understanding, etc. due to the use of immersive VEs.

To avoid potential confusion, we begin by introducing some terminology related to VEs and follow with descriptions of two of the major components of immersion.

When describing the characteristics of a VE, the terms *immersion* and *presence* are often confused, and used interchangeably. They are actually separate and distinct concepts. Immersion is related directly to the technology used in a VE system. The higher the fidelity of perceptual outputs from the system and user input to the system, the greater the level of immersion. Presence is the user's *response* to a given VE system, or more generally, "the sense of being there." Presence can vary a great deal from person to person; one person may feel a vastly different intensity of presence than someone else in the same VE with identical levels of immersion [Slater03].

Two components of VE technology (contributing to levels of immersion) central to our study are *physical immersion* and *head-based rendering*. Physical immersion is the degree to which the virtual world appears to visually surround the user in space. The *field of regard* (FOR) is equivalent to the level of physical immersion, and can be defined as the visual angle surrounding the user within which the virtual world is displayed to the user. Note that field of view (FOV) is not the same as FOR. The FOV is the area (also measured in visual angle) within which the user sees the virtual world at any instant of time. In a tracked head-mounted display (HMD), for example, the horizontal FOV may be 60 degrees (the user sees the virtual world in 60 degrees of his visual field at any instant), but the horizontal FOR is 360 degrees, since the user sees the virtual world no matter which direction he looks. With a non-head-tracked large projection screen, the horizontal FOV might be 90 degrees. The horizontal FOR in this case would also be 90 degrees, since the user can see the entirety of the display at a glance.

Head-based rendering means that the virtual world is drawn from the user's viewpoint. This is usually accomplished by the measurement of the user's 3D head position and orientation, from which the location of the user's eyes is inferred. Head-based rendering provides an intuitive method of viewing from various perspectives in the VE. In an HMD-based system, orientation tracking is critical, because it is necessary to produce a 360-degree FOR. In a projection-based VE system, however, orientation tracking is less important, as the projected view changes only slightly as the user's head turns. Position tracking is important with both types of displays, as it can provide a strong depth cue (motion parallax) and a natural means of obtaining different views of the environment.

Prior efforts aiming to measure the benefits of immersion, in our view, were unable to distinguish between the components of immersion (see section 2). This, and other factors that were not held constant among the systems used for comparison, has led us to believe that a new approach is needed. This new method should provide statistical evidence for the claims of benefits from immersion, and should allow independent evaluation of many components of immersion. We present a method, based on CAVE-like displays [Cruz-Neira93] that addresses these challenges and is focused on the evaluation of the effects of physical immersion and head-based rendering.

In the next section we begin with a survey of related work, then proceed with a description of the method. Next, we explain proof of concept evaluation, in which the method was applied to an information visualization application. , We conclude with a list of further experiments where this method could be used, a summary of our findings, and areas for future work.

2 Related Work

In general, prior research attempting to measure the benefits of immersion has made use of two different strategies: practical experiments and controlled experiments.

2.1 Practical experiments

Practical experiments compare systems that are similar to those that would be used in the real world. In the study of immersion, therefore, a typical experiment would compare a desktop (non-immersive) application with a reasonably equivalent immersive version of the application. While the results of such studies are directly applicable to real-world choices between systems, they are not generalizable because the systems being compared typically have many differences (confounds). Below are three examples.

In a study of information visualization applications, a desktop tool set named XGobi was compared with a similar application running in a CAVE-like system [Arns98]. Similar datasets were displayed on both devices, and users performed a series of tasks in order to see if there were any tangible benefits of viewing datasets in an immersive environment. The authors hypothesized that viewing high dimensional statistical data would be more efficient in an immersive VE. The authors found that some tasks requiring spatial understanding were performed more quickly in the VE than on the desktop. Confounds in this study included FOV, FOR, resolution, display device, input device, and interaction techniques.

A similar, more recent, comparison looked at the effects of immersion on oil well path editing tasks [Gruchalla04]. This study also compared user task performance on a desktop and in a CAVE, but tried to eliminate many of the differences between the systems. The two systems had the same stereoscopic rendering and the same resolution. However, they differed in many respects as well – input device, display device, the user's virtual frame of reference, head-based rendering, physical immersion, and interaction techniques were all confounded in this study. Gruchalla found significantly better task performance in the CAVE, which is a useful result when choosing between these two specific system configurations. Because of the confounded experimental design, however, we cannot know what caused these differences, and we cannot generalize these results to say that *immersion* resulted in better task performance.

Another comparative study involved the use of a desktop and an HMD to determine if an immersive VE provided any benefit when spatial attributes of the data were important in task completion [Datey01]. The author found a trend for increased task performance in the immersive condition, but this trend was not statistically significant. This study attempted to control many variables, but still confounded immersion with display device, input device, and interaction techniques. Moreover, physical immersion and head-based rendering were confounded in this study.

2.2. Controlled experiments

Controlled experiments have well-defined independent variables that are explicitly managed by the experimenter. All other factors in a controlled experiment are held constant (ideally). This has the advantage of allowing precise statistical analysis of the results, so that if a difference (in performance, usability, etc.) is found, the cause of that difference is clear. The tradeoff, of course, is that in achieving this level of control, the systems being compared are often different than those that would be used in a real application, so it may be difficult to apply the results of the experiment.

The best example of a controlled experiment attempting to quantify the benefits of immersion is the work done by Pausch and his colleagues [Pausch97]. In this work, a comparison was made between an HMD with head tracking and a stationary HMD with hand-based input for navigation and viewpoint control. They hypothesized that users would be able to find a target faster in the head-tracked condition, but did not find this to be the case. However, the head-tracked HMD users were able to determine if a target was *not* located in an environment significantly more quickly. This suggests perhaps, that the head-tracked subjects built a cognitive map of the space more quickly, and avoided redundant searching. Unlike the studies described in section 2.1, in this study there was only one difference between the two conditions – the method of setting the viewpoint orientation (via the head or hand). The display device (HMD) and all its associated properties (e.g. FOV, resolution), the input device (a tracker), the rendering style, and all other characteristics of the two systems were held constant. Thus, any differences between the conditions must have been due to immersion.

Unfortunately, this setup is still not sufficient to study the *components* of immersion. The important distinction here is that in an HMD-based VE system, head tracking *produces* physical immersion, so there is no way to separate these two components. This is the primary impetus behind our use of the CAVE for quantifying the benefits of immersion.

In summary, existing methods for investigating the effects of immersion either confound multiple variables or measure the effects of something other than immersion. In our approach, there is no confounding of variables since the display devices, input devices, interaction method, user's posture, and physical level of encumbrance are equivalent under all testing conditions.

3 Description of Method

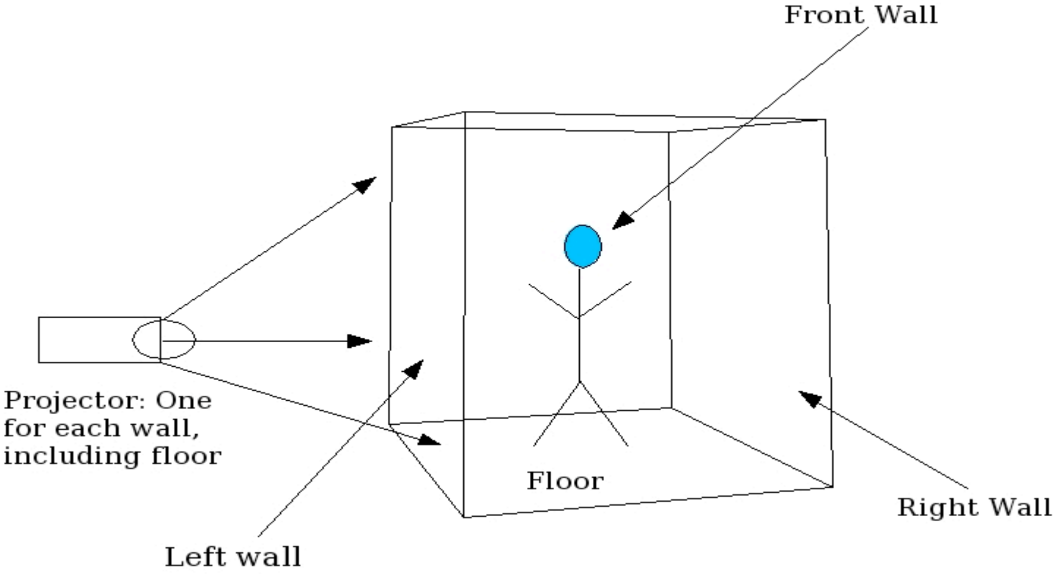

As noted in the introduction, two primary components of immersion are physical immersion and head-based rendering. In order to evaluate the benefits of immersion, we needed a method that allowed us to separate the effects of these two components, and does not confound variables or investigate aspects other than immersion. We have already seen that HMD-based VE systems cannot achieve this separation of variables, since head tracking is required for physical immersion in such systems. Using the CAVE [Cruz-Neira93] or another surround-screen VE display, however, the components of immersion can be varied independently. A typical CAVE has four projection surfaces (three walls and a floor), and uses stereo projection technology and head tracking to display an immersive virtual world.

In our method, the number and location of active display surfaces determines the level of physical immersion (i.e. the FOR). When one screen is in use, a low degree of physical immersion is produced, while the use of all four screens creates a high degree of physical immersion. While FOR is varied explicitly, FOV is held constant in our method. When the user wears the active stereo glasses in the CAVE, the horizontal FOV is restricted to approximately 100 degrees (much less than humans' natural horizontal FOV, which is greater than 180 degrees). In the one-screen (low FOR) condition, a user standing at the center of the CAVE can see the entire screen at a glance, so FOV = FOR. In the four-screen (high FOR) condition, the user sees 100 degrees at a glance, but can view 270 degrees total horizontal FOR (see figure 1). We did not measure the effect of FOV in our proof-of-concept study, but could in future experiments as mentioned in section 5.

To control head-based rendering in our method, the head tracking system is enabled or disabled. In the non-head-tracked condition, the user is assumed to be at the center of the CAVE, looking straight ahead. The user is allowed to rotate his head to view the scene (e.g. in the four-wall condition), but neither this rotation nor any involuntary head movements change the way the scene is rendered (for greater accuracy, we should use an "OmniStereo" projection [Simon04] in this condition; with the naïve approach, objects that appear on the side walls will be distorted due to the assumption that the user is looking straight at the front wall). In the head-tracked condition, the position and orientation of the user's head are used to draw the scene. While the effects of head rotation may be difficult to perceive, the effects of head translation are quite obvious in this condition. The user can change his viewpoint simply by walking, crouching, leaning, jumping, etc.

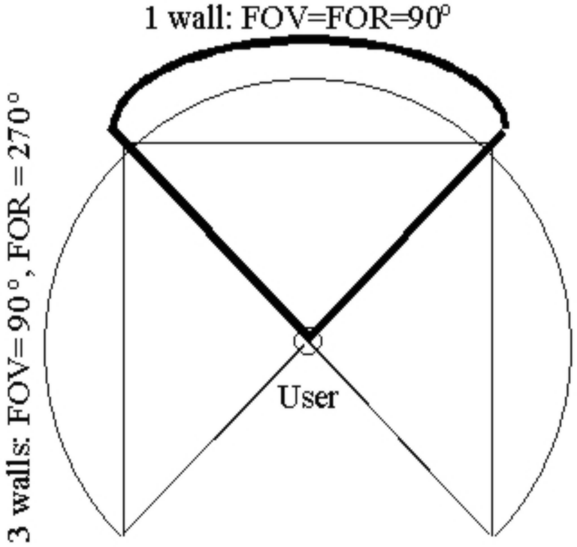

Figure 1. Top-down view of the CAVE indicating FOV and FOR for the two levels of physical immersion

The combination of our two independent variables – two levels of physical immersion and two levels of head-based rendering – lead to four conditions that can be evaluated in an experiment (table 1).

	Head-Based Rendering	**No Head-Based Rendering**
Low Physical Immersion	One screen CAVE	One screen CAVE, no Head Tracking
High Physical Immersion	Four screen CAVE	Four screen CAVE, no Head Tracking

TABLE 1 –Experimental Conditions

Although an extremely simple idea, this experimental design can be used to quantify the effects of immersion for a wide variety of application domains and user tasks, using a wide variety of metrics. It would allow us to measure the effects of immersion on learning in an educational application, on spatial understanding in a geographic information system, or on decision-making in an architectural design context, just to name a few possibilities. See section 5 for more examples of the potential use of the method.

4 Proof-of-Concept Evaluation

As an initial demonstration of the use of our method for quantifying the benefits of immersion, we evaluated CaveDataView, an immersive information visualization application that displays 3D scatterplots in the CAVE. Our motivation was some previous work that had suggested that immersion might be beneficial for some information access and understanding tasks in data visualization [Datey02].

CaveDataView represents data points by yellow cubes laid out on a 3D grid (figure 3). The user can navigate by moving his head and body (physical movement) and by using a wand input device (virtual movement). The user can also select data points to see their ID number (figure 4).

FIGURE 3 – CaveDataView in the CAVE

FIGURE 4 – Ray-casting and data point labeling in CaveDataView

4.1 Procedure

Each user performed a series of timed tasks in the environment, using each of the four conditions shown in table 1. We also collected difficulty and usefulness ratings on a scale of 1-7 (7 being most difficult and useful, respectively) after each task was performed. We administered a post-experiment questionnaire containing more questions about the various immersive conditions and their effects on the performance of the tasks. Users were then asked to rate which conditions were most favored.

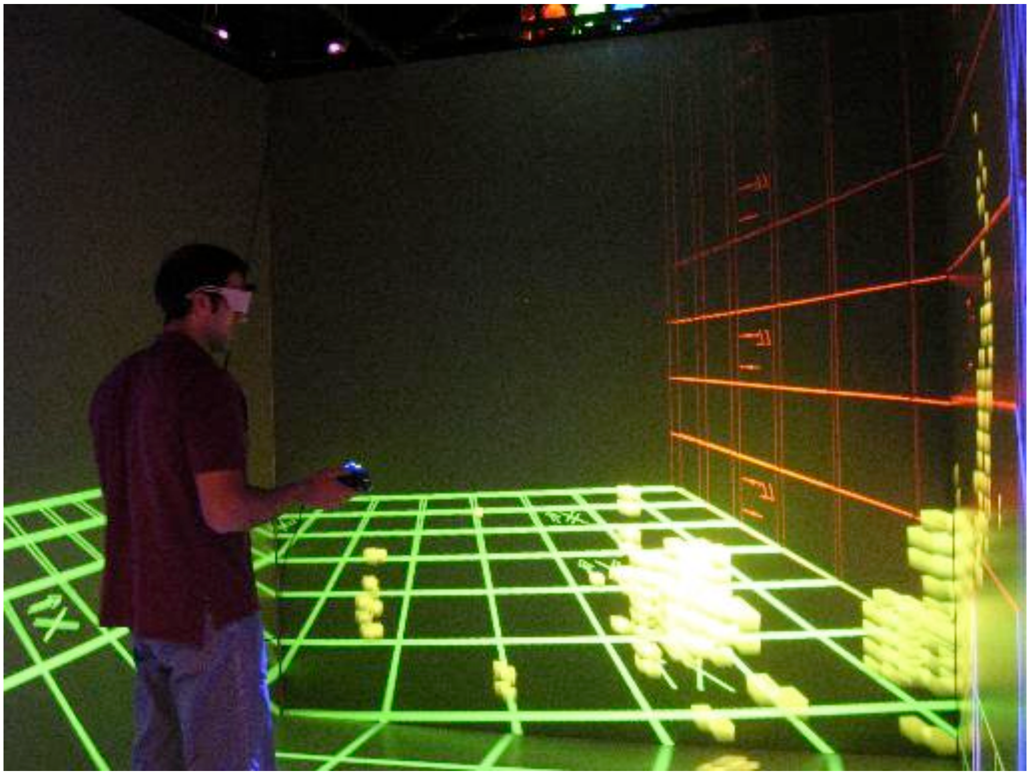

4.2 Results

The full results of this study have been reported elsewhere [Raja04]. The subset of results presented below are only intended to show that information about the effects of physical immersion, head-based rendering, and the combination of the two can be obtained using our method.

To ease the discussion of the results, we will identify each condition from table 1 with a number:

Condition 1 = four screens with head-based rendering
Condition 2 = four screens with no head-based rendering
Condition 3 = one screen with head-based rendering
Condition 4 = one screen with no head-based rendering

To examine the effects of physical immersion, we compared condition 1 with condition 3, and condition 2 with condition 4. Similarly, we analyzed the effect of head-based rendering by comparing conditions 1 and 2, and conditions 3 and 4. Finally, we examined the combined effect of both components by comparing the most immersive condition (1) with the least (4).

Timing metric: For two user tasks, users performed slightly slower in conditions with a high degree of physical immersion (conditions 1 and 2), while for two other tasks, users were much faster in the high physical immersion conditions.. Out of the four tasks that were completed under both conditions, three of them were faster in the physically immersive condition. We also found that subjects performed tasks more quickly in conditions 1 and 3, suggesting that head-based rendering can lead to greater efficiency.

Subjective ratings: User ratings indicated similar levels of disorientation for conditions 1 and 3, but higher disorientation in condition 2, indicating that head-based rendering helped users to maintain spatial orientation. Head-based rendering also appeared to make a difference in the users' preference: the average perceived usefulness was higher in condition 1 than in condition 2.

These results might be attributable to the intuitive nature of head-based rendering which permits the user to view the datasets from various perspectives in a natural manner. The user is able to move through the dataset in finely grained increments by moving her head. This could be especially useful in dense areas where there is quite a bit of occlusion from the data points.

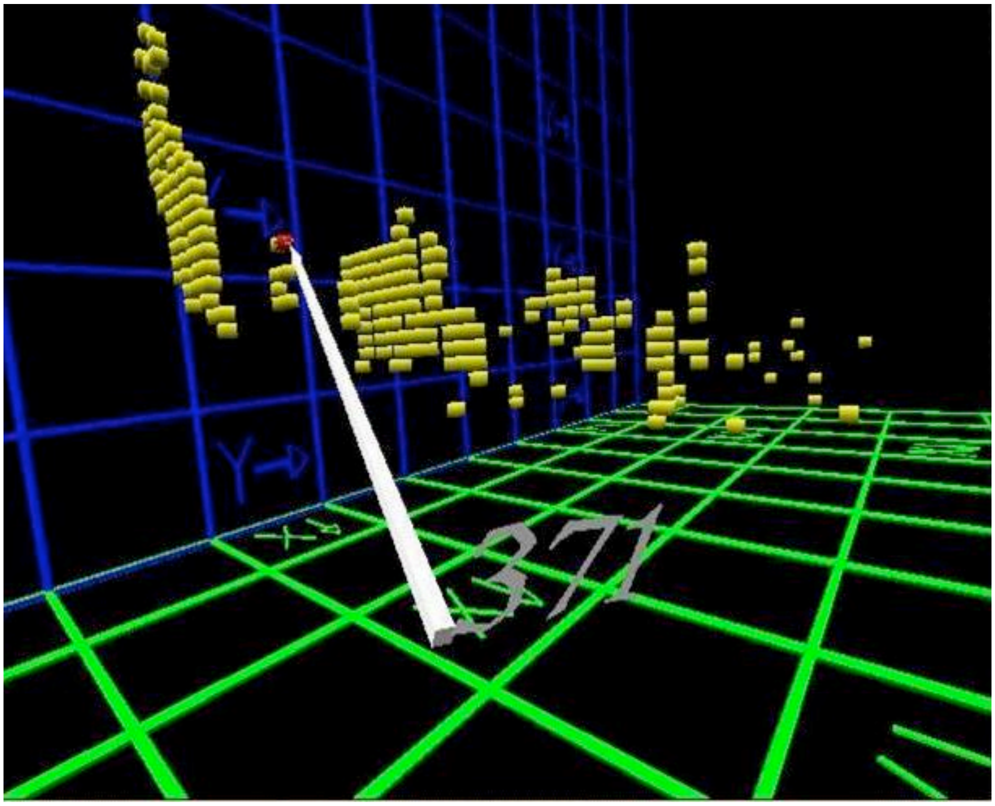

User comments and observations: All users stated that when four screens were used, it was much easier to view large datasets. One subject comment read, "Four walls very useful in ability to view as much of the data set as possible. Six walls would be even better!" This was reinforced by observing the subjects' behavior when certain datasets were displayed in the one-screen condition. Some subjects turned to look at the side walls expecting points to be displayed there, and when none appeared, seemed a bit frustrated that they had to turn back to the front wall to manipulate the dataset further. From this behavior, it is easy to see the effect that physical immersion has on some users. If this were simply a practical comparison between a desktop and immersive VE, the cause for these comments would not be so easily attributed to that particular aspect of immersion.

5 Future Experiments

Our method for quantifying the benefits of immersion is applicable in a wide variety of domains – anything where immersion is thought to be useful. For example, in the treatment of certain phobias using VEs, It would be interesting to see how much of an impact immersion has. In one system designed to assuage a patient's fear of flying [Hodges96], the subject is immersed inside a virtual aircraft much like a commercial jetliner. Would a high degree of immersion correlate with effective treatment of the phobia? With our method, an experiment could be run which systematically varies the levels of immersion. The outcome of an experiment like this could have profound implications on the commercial viability of such systems. If the study found immersion is not important in the treatment of phobias, then systems with a relatively low degree of immersion could be built for a fraction of the cost of more highly immersive ones.

Education is another area where the method would be useful. For example, the ScienceSpace project [Dede96] used immersive VEs to teach principles of science to high school students. The authors hypothesized that immersion would improve the retention rates of concepts that are difficult to grasp. Our method would provide a means of determining whether immersion is beneficial for these educational goals, and if so, which components of immersion are the most important.

Our method can also be extended to examine other aspects of immersion. For example, we could easily add the ability to test the potential benefits of increased FOV. We currently hold the horizontal FOV constant, but physical or virtual “blinders” could be used to artificially limit the FOV. Other aspects of immersion that could be considered as variables would include the use of stereoscopic graphics, the display resolution, or the level of detail of the rendered environment, among many others. When all of these variables are truly independent, we can form an accurate model of their individual and combined effects on usability, task performance, presence, or any other useful metric.

One drawback to our current approach is the limited FOR (physical immersion) that can be achieved in a four-sided CAVE. With a fully-surrounding display such as Iowa State’s C6, we could examine any possible combination of FOV and FOR. For example, we could simulate a common HMD by artificially limiting the FOV (e.g. to 60 degrees horizontal) but by using a 360-degree FOR.

6 Conclusion

Demonstrating the benefits of immersion is a critical challenge for the VE community. The acceptance and real-world usage of immersive VEs depends on proof of their usefulness. The simple method outlined in this paper provides a framework for empirical studies that can provide such proof.

In the information visualization application we used for proof-of-concept, we separated and analyzed two important components of immersion quite readily. While there were not enough subjects in our informal study to conduct a thorough statistical analysis, we were able to independently evaluate these components of immersion. And this is the key point and most salient feature of this method: it can be applied more generally to examine and independently evaluate immersion's components in any application domain, task, or scenario where immersion might provide a benefit.

Acknowledgements

The authors wish to thank Prof. Chris North, John Lucas, and the subjects who participated in our exploratory pilot study.

References

Ahlberg, C. and Wistrand, E. “An information visualization & exploration environment.” In Proceedings of IEEE Information Visualization, 1995, pp. 66-73.

Anupam, V., Dar, S., Leibfried, T., and Petajan, E.. “DataSpace: 3-D visualizations of large databases.” In Proceedings of IEEE Information Visualization, 1995, pp. 82-89.

Arns, L., Cook, D., Cruz-Neira, C. “The Benefits of Statistical Visualization in an Immersive Environment.” In Proceedings of IEEE Virtual Reality, 1999, pp. 88-95.

Cruz-Neira, C., Sandin, D., DeFanti, T. “Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE.” In Proceedings of ACM SIGGRAPH, 1993

Datey, A. “Experiments in the Use of Immersion for Information Visualization.” Unpublished Masters thesis, Virginia Tech Computer Science Department. Available at http://scholar.lib.vt.edu/theses/available/etd-05092002-151043/

Dede, C., Salzman, M. C., Loftin, R. B. “ScienceSpace: Virtual Realities for Learning Complex and Abstract Scientific Concepts.” In Proceedings of IEEE VRAIS, 1996, pp. 246-252.

Gruchalla, K. “Immersive Well-Path Editing: Investigating the Added Value of Immersion.” In Proceedings of IEEE Virtual Reality, 2004, pp. 157-164.

Hodges, L. F., Rothbaum, B., Watson, B., Kessler, G. Drew., Opdyke, D. “A Virtual Airplane for Fear of Flying Therapy.” In Proceedings of IEEE VRAIS, 1996

Hollands, J.G., Pierce, B., Magee, L., “Displaying Quantitative Information in Two and Three Dimensions.” In Proceedings of Human Factors and Ergonomics Society 39th Annual Meeting, pp. 1425-1429.

Kelso, J., Arsenault, L., Satterfield, S., and Kriz, R. “DIVERSE: A Framework for Building Extensible and Reconfigurable Device Independent Virtual Environments.” In Proceedings of IEEE Virtual Reality, 2002

Lin, J., Duh, H., Abi-Rached, H., Parker, Furness, T.“Effects of Field of View on Presence, Enjoyment, Memory, and Simulator Sickness in a Virtual Environment.” In Proceedings of IEEE Virtual Reality, 2002, pp. 164-168.

Mine, M., “Virtual Environment Interaction Techniques.” TR95-018, UNC Chapel Hill, Department of Computer Science, 1995

Pausch, R., Proffitt, D., and Williams, G. (1997), “Quantifying immersion in virtual reality.” SIGGRAPH'97.bertson, G., Card, S., and Mackinlay, J. (1993). “Information visualization using 3D interactive animation.” Communications of the ACM, 36(4), April 1993.

Raja, D., Bowman, D., Lucas, J., and North, C. (2004), “…”

Salzman, M. C., Dede, C., & Loftin, B. (1996). “ScienceSpace: Virtual realities for learning complex and abstract scientific concepts.” Proceedings of IEEE Virtual Reality Annual International Symposium, (pp. 246-253). New York: IEEE Press.

Slater, M., Usoh, M., Linakis, V., & Kooper, R. (1996). “Immersion, presence and performance in virtual environments: An experiment with tri-dimensional chess.” Proceedings of 1996 ACM VRST, 163-172.

Slater, M., (2003). "A Note on Presence Terminology." www.presence-connect.com, January 2003

http://www.spotfire.com/products/decision.asp

Wickens, C., LaClair M., and Sarno K.(1995) “Graph Task Dependencies in 3D Data: Influence of 3D and Color.” Proc Human Fact

---

Simon, A., Smith, R., and Pawlicki, R. “OmniStereo for Panoramic Virtual Environment Display Systems” In Proceedings of IEEE Virtual Reality, 2004, pp. 67-74.