Exploring the Benefits of Immersion in Abstract Information Visualization

Dheva Raja, Doug A. Bowman, John Lucas, Chris North

Virginia Tech
Department of Computer Science
Blacksburg, VA 24061
{draja, bowman, jolucas, chnorth1}@vt.edu

ABSTRACT

The benefits of immersion with regard to information visualization applications have rarely been explored. In this paper, we describe a user study designed to explore a variety of information visualization tasks in immersive and non-immersive 3D scatterplots. In the non-immersive version the information was displayed using only one wall of the CAVE, while the immersive version used all four walls. We also examined the effects of head tracking, giving a total of four conditions: four walls in the CAVE with head tracking, four walls without head tracking, one wall with head tracking and finally, one wall without head tracking. By separating the variables in this way, we can independently evaluate the effects of immersion and head tracking. In general, we found the fully immersive condition, (four walls with head tracking) to be most useful in viewing the datasets and performing the tasks.

Keywords

Immersion, Head tracking, CAVE, Virtual Environment, Virtual Reality, Information Visualization, 3D Scatterplots

1. INTRODUCTION

Numerous visualization applications have been developed for use in immersive virtual environments (VEs). These primarily are comprised of, but not limited to, scientific visualizations, architectural walkthroughs, and simulations of various devices in commercial and industrial enterprises. As one might imagine, not all applications have an industrial or scientific purpose for their origin. Some quite stunning artistic demonstrations, as well as quite a few games have been developed or ported for use in VEs.

Physical immersion and *head tracking* are the primary characteristics that make these applications so compelling. Physical immersion is the degree to which the virtual world surrounds the user in space. A VE where the user feels a strong sense of "presence" generally has a high degree of physical immersion with a wide field of view.

Head tracking is the measurement of the user's head position and orientation, which is then used to render the world from the user's point of view in space. It provides an intuitive method of viewing from various perspectives in the VE. Taken together, we call these two characteristics *immersion*.

Immersion has been shown to be beneficial in applications where spatial knowledge of an environment is useful. In an information visualization application, physical immersion and head tracking could allow for more efficient identification of trends in data, greater spatial understanding of the entire data set, and easier identification of single data points.. The CAVE [Cruz-Neira93] is unique in that it allows data to be seen in a wide field-of-view without moving one's head and without distorting the spatial relation of the data. By contrast, in a head mounted display (HMD), the user is still immersed but does not benefit from the extent of the peripheral vision that one does in the CAVE.

Our goal is to understand the benefits of immersion for information visualization. If immersion is proven to be beneficial, immersive VEs could become more popular for analyzing 3D data sets. Our results in this paper are limited to 3D scatterplots, and thus are not necessarily applicable to all types of 3D data visualizations, but further studies could show immersion's usefulness. Other 3D data visualization techniques that could benefit from immersion are: multi-dimensional function visualizations, 3D histograms, line graphs and surface maps.

We have developed a CAVE-based information visualization application (CaveDataView), and conducted an exploratory pilot study in which we evaluate the benefits of immersion for 3D scatter plots. The goal of this research was to show that there are tangible benefits to viewing a generic information visualization application in a physically immersive, head-tracked environment such as the CAVE.

To be specific, our hypotheses were:

1. **A high degree of physical immersion will allow higher levels of task performance and greater user satisfaction when visualizing datasets represented by 3D scatterplots.**

2. **Head tracking will allow higher levels of task performance and greater user satisfaction when visualizing datasets represented by 3D scatterplots.**

We begin with a survey of relevant related work, followed by a description of our immersive information visualization application. Next we describe the design of our evaluation and its results. Finally, we conclude and present our ideas for future work.

2. RELATED WORK

2.1 3D Information Visualization Techniques

3D information visualizations take a complex and abstract dataset and organize it into a 3D visual representation, which can be navigated and accessed by the user. Abstract properties of the data are mapped into perceptual qualities, such as position, orientation, size, shape, color, or motion, and relationships between pieces of data are represented spatially. The resulting visualization can reveal patterns in the data that may not be obvious from the original dataset. A number of 3D information visualization tools have been developed. Dataspace [Anupam95] is a system for interactive 3D visualization of large databases. IVEE [Ahlberg95] is a 3D environment that uses a number of techniques such as maps, star fields and query mechanisms for visualizing a database. Work from Xerox PARC [Robertson93] provides additional examples of the use of interactive 3D graphics for information visualization. Spotfire [Spotfire] is a commercial desktop application that allows users to load in a dataset of their choice, and visualize it using 2D and 3D graphs. 3D scatterplots in Spotfire are very similar to the plots we display in CaveDataView.

2.2 Information Visualization Applications in VEs

There are not many current applications for information visualization in immersive VEs, and very few specifically intended for use in the CAVE. One notable exception is the C2 statistical program, which was loosely based on a desktop tool called XGobi. [Arns98]

VR Vibe [Hollands95] is an HMD-based application, which creates a visualization of bibliographies for information retrieval. Users specify keywords in 3D space, and representations of the documents are then displayed in the space according to how relevant each document is to each of the keywords. The position of a document depends on the relative importance of each of the keywords to it, which is computed using document-matching algorithms.

The LEADS system developed at University of Nottingham [Ingram95] applies concepts based on urban planning for database visualization and abstract information domains. The system uses a city metaphor based on districts; nodes and edges connected by paths and landmarks facilitate formation of cognitive maps to avoid getting 'lost' in the information space. The LEADS system shows how using an easily recognized metaphor simplifies information visualization using immersive VEs.

The precursor to this research was the Wizard application [Datey01], an HMD-based immersive VE for exploring 3D scatterplots. Wizard provided both a small hand-held overview of the dataset and a larger version of the data through which the user could navigate.

2.3 Comparative Studies

In the most closely related work to our study, a desktop visualization toolset named XGobi was compared with a C2 virtual environment system [Arns98]. Similar datasets were displayed on both devices, and a series of user tasks were performed in order to see if there were any tangible benefits to 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 a good deal of spatial understanding were performed more quickly in the VE than on the desktop.

Wickens [Wickens95] compared conventional 2D graphs with 3D graphs for presenting 3D data. Users were asked questions about the data that ranged from focused attention on a single data point, to questions that

integrated the entire data space. Here the authors found that the 3D display resulted in the longest times for focused attention tasks, but this diminished markedly as the questions became more integrative in nature.

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 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

3. SYSTEM DESIGN AND DESCRIPTION

3.1 CaveDataView

CaveDataView is the VE application we developed to test our hypotheses. It can display scatterplots of 3D datasets in a CAVE (figure 1). We developed the application in C++ using the DIVERSE application programming interface (API) [Kelso02]. DIVERSE allowed us to create an application that would work, with little to no modification, on a CAVE, and HMD or a desktop.

DPF, the Diverse interface to OpenGL Performer, was used for the generation of the scene graphics. DPF and Performer were powerful enough to allow us to render scenes with large numbers of data points. We tested our application under DPF 2.3.1 which is available for both the Irix and Linux platforms. To interface with the tracking hardware, we used the Diverse Toolkit (DTK). DTK encapsulates the tracking system and places the needed

information into shared memory making it easily available to the application. Diverse uses a Dynamic Shared Object structure that allows application modules to be quickly added or removed. This structure was invaluable in that it allowed us to develop on desktop computers and move relatively seamlessly to a CAVE display system.


<FIGURE 1 - A 3D scatterplot in CaveDataView>

CaveDataView reads in tab-delimited files of three column data. Each line of the file is a distinct data point. Once all points have been read, the data space is scaled to fit in a reasonably sized VE. Data points are represented by yellow cubes laid out on a 3D grid. The grid itself is a 3D object that stretches from the origin along the three coordinate axes to a size three times that of the dataset. This is important so that the user does not get disoriented while navigating through and around the data. Lighting and shading are added to the scene to make the edges and corners of the cubes easier to see. Shading also helps the user orient themselves to the cubes as each side had different shading applied to it due to the direction of the light.

<FIGURE 2 – The Virgina Tech CAVE>

Our CAVE (figure 2) is a four walled display, 3 sides and a floor, that uses stereo projection technology and head tracking to display an immersive VE. The CAVE uses an Intersense IS900 VET tracking system that comes with a 6-DOF Head Tracker and a 6-DOF Wand control device. The head tracker tracks the user's movement through the environment and causes the scene to render the correct perspective for the user's position and orientation. The wand device has a small joystick and four buttons, and is used for navigation and manipulation of the environment. Wand navigation is used for moving quickly through the virtual world; for fine movements the user can walk or just move his head when head tracking is enabled. Figure 3 shows a user standing in the CAVE while CaveDataView is running.


<FIGURE 3 – CaveDataView in the CAVE>

We realized that in order to test our hypotheses we would need to give the user some ability, beyond basic navigation, to interact with the environment. More specifically, the user would have to be able to point at individual objects of interest to them or that were important to the tasks they were asked to perform. To do this, we implemented simple ray-casting [Mine95]. A virtual ray is projected from the tip of the wand in the direction the wand is pointed. When the ray intersects a data point, the ray shrinks to the point of intersection on the side of the data point. This shrinking provides a valuable visual cue to highlight the data point of interest. To further highlight the data point of interest, the color of the cube is changed to red when intersected by the ray.

Aside from pointing, we required a method for the user to report verbally which data point they had selected. To do this each data point was assigned an ID number that had no relation to its location or any of the attributes of the data. We first attempted to put this ID number on the side of each cube, but quickly exceeded the rendering power of the SGI Onyx that renders the scene. The decision was made to only label the currently selected cube with its ID number. As this ID number was difficult to read from a distance, we also placed it at the start of the selection ray above the wand (see figure 4). This labeling method worked very well

as when close to the data points the labels on the sides of the cube were very readable, and the wand label was readable when selecting cubes from a distance.

<FIGURE 4 – Ray-casting and data point labeling in CaveDataView>
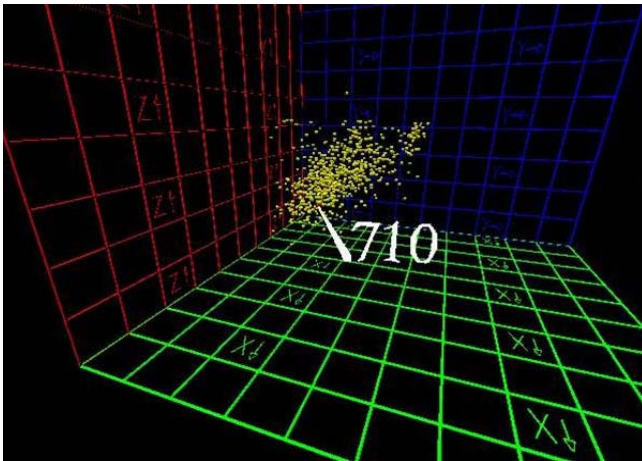
4. EVALUATION

4.1 Experimental Design

Quantifying the benefits of immersion is obviously an important challenge for the VE community. If attractive cost/benefit ratios can be proven for particular tasks and domains, industries will be much more likely to invest in VE technology. However, it is also very difficult to design controlled experiments to measure these benefits. Previous work has taken two approaches.

First, some researchers have attempted "practical" comparisons of immersive and desktop systems (e.g. [Datey01]), in which the same application is used with an immersive display such as a head-mounted display (HMD) and a non-immersive display, such as a desktop monitor. The problem with these evaluations is lack of experimental control. There are many differences between the two conditions besides immersion, including input device, interaction techniques, resolution, brightness, field of view, the user's posture (standing or sitting), etc. Thus, if a difference between the conditions is found, it is not at all clear that the difference is due to immersion.

The second approach is to control these additional factors by using the same display in both conditions. As we have already seen, Pausch et al. [Pausch97] used an HMD along with two ways for the user to rotate the viewpoint: via head turning and via hand turning. This study showed that the head-tracked condition was superior in some ways for a search task. We would claim, however, that there was still a confound between head tracking and physical immersion in this study – in the head-tracked condition, the user was physically immersed (the virtual world appeared to surround him in space), while in the hand turning condition, no physical immersion existed. In other words, in an HMD-based system, head tracking *creates* physical immersion, and there is no way to separate the two.

In an immersive projection display, such as the CAVE, however, we can separate these two variables and examine their effects separately. This has led to our experimental design, shown in table 1. In the physically immersive condition, we use all four walls of the CAVE, while in the physically non-immersive condition, only the front wall is used. Separately, we can control the use of head tracking (on or off), leading to four possible combinations. Two of these conditions are novel as compared to Pausch's experiment. In the physically immersive, non-head-tracked condition, the virtual world surrounds the user, but the user cannot rotate her head to view different parts of the world, nor can she translate her head to get a different perspective on the world. In the physically non-immersive, head-tracked condition, the virtual world only appears in front of the user, but she is free to turn or translate her head to see the world from a different point of view.

Tasks involving different datasets were performed in these four conditions, and the times to complete the tasks were measured. Subjects provided difficulty ratings and disorientation ratings on a seven-point scale after each task was performed in a particular condition. Additionally, at the conclusion of the subjects' series of tasks, a questionnaire was administered. The results were then gathered and evaluated.

	Head Tracking	No Head Tracking
Non-Immersive	One wall CAVE	One wall CAVE, no Head Tracking
Immersive	Four wall CAVE	Four wall CAVE, no Head Tracking

<TABLE 1 –Experimental Conditions>

Because of the exploratory nature of this pilot study, no statistical analysis was performed on the data.

4.2 Tasks

We developed 6 tasks with which to test our hypotheses. The tasks themselves represented typical tasks that would be performed when analyzing a data set. As noted above, all tasks were timed, although for some tasks completion time would not necessarily be the most important factor. Tasks were chosen so that they would have only a single correct answer.

1.   One Axis Distance

One Axis Distance asked the subject to find the point with the highest Y value. This task tested their ability to judge distances along one axis in the scatterplot. We felt this was an important basic task in that it is a key component in gaining understanding of a single data point.

2.   Two Axis Distance

Two Axis Distance required the subject to locate the point with the both the lowest X and lowest Y value. We were concerned about the subjects' ability to orient themselves so that they could compare distances along two axes at once. This proved to be one of the more difficult tasks.

3. Trend Determination

Trend Determination required the subject to get a general sense of the layout of the data in order to spot trends. Subjects were asked to report the trend in this format: as A increases/decreases, B increases/decreases. Subjects could also report that the data had no trend. All datasets used in the study exhibited some sort of trend.

4. Clusters

The Cluster finding task asked subjects to locate clusters of data points greater than 20 points. We felt this was another important task for data visualization, but as the definition of "cluster" can be subjective, we had to loosely define what we thought a cluster was, as well as use data that constrained what the user would identify as a cluster.

5. Single Point Search

The Single Point Search task had subjects locate a differently colored point in a densely packed group of points. The point was not visible from the subject's starting position and required them to navigate to find it. This task would be important in a real-world application where a data point is highlighted in another view and must then be located in the VE view (e.g. a brushing-and-linking task).

6. Outliers

The Outlier task was designed specifically to answer a question about VE scatter plots. We asked the subject to find the two data points that were furthest away from the main group of data points. In the VE it was easy to miss data points that were very far above the user or just out of the field of view. We expected the fully immersive conditions to perform better for this task.

4.3 Subjects

Four subjects were recruited to participate in the study. We wanted to use experienced VE users in order to minimize problems associated with VE navigation or sickness. The subjects performed all 6 tasks once, and saw all four conditions at least once. Table 2 shows the design we used. The entries in the table show the order in which the tasks were performed. We ordered the tasks so that each task was performed for each condition. Since this was an exploratory study, we felt four subjects were sufficient. We made sure to get as much input as possible from each subject so that we could refine our evaluation method.

		Condition			
		1	**2**	**3**	**4**
Subject	**1**	1, 2	3, 4	5	6
	2	3, 4	5	6	1, 2
	3	5	6	1, 2	3, 4
	4	6	1, 2	3, 4	5

<TABLE 2 – Subject Task Order>

4.4 Results

This section is organized by exploring the results pertaining to the first hypothesis regarding physical immersion, followed by the second hypothesis regarding head tracking.

Listed below are the four conditions. Four walls indicate a high degree of physical immersion; one wall indicates a low degree of physical immersion.

Condition 1 = four walls with head tracking

Condition 2 = four walls with no head tracking

Condition 3 = one wall with head tracking

Condition 4 = one wall with no head tracking

4.4.1 Physical immersion

In examining physical immersion, tasks completed under condition one were compared with those completed under condition three. Conditions two and four were not compared due to missing data. The average time to complete tasks one and two was 34.1 seconds under condition one. The time to complete those tasks under condition three was 30.5 seconds. For tasks four and five, the time for condition one was 16.4 seconds. For condition three the time for those tasks was 25.7 seconds. For tasks one and two, a high degree of physical immersion was a few seconds slower, and for tasks four and five it was much faster. Out of the four tasks that were completed under both conditions, three of them were faster in the physically immersive condition. From this data we can see a trend for greater task efficiency in a physically immersive environment when head tracking is also used.

The average perceived disorientation and usefulness levels of condition one were 1.75, (1-7, 7 being most disorienting) and 5.5 (1-7, 7 being most useful), respectively. For condition three, the average disorientation and usefulness ratings were 1.5 and 5.3. The average difficulty across condition one for completing the tasks was 3.3 (1-7,7 being most difficult.) For condition three the average difficulty was 3.2. These three metrics, therefore, did not indicate any benefit for physical immersion in completing the tasks.

All users stated that when four walls of the CAVE 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-wall 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. Certain tasks seemed to benefit particularly from a high degree of physical immersion. This might explain why tasks four and five were completed significantly faster under condition one than condition three. Task #4 involved finding clusters, and in task #5, the subject had to find a colored cube in a densely packed dataset. A wider field of view permitted a greater area of the dataset to be visualized at once, resulting in markedly lower completion times.

4.4.2 Head tracking

For head tracking, conditions one and two were compared. Conditions three and four were not compared due to missing data. For all tasks, condition one (with head tracking) was much faster than condition two (no head tracking). The average disorientation and usefulness ratings of condition one were, 1.75 and 5.5, respectively. For condition two those numbers were 2.75 and 4.25. The average. difficulty for condition one was 3.3, for condition two it was also 3.3. We see from this data a trend for greater efficiency, lower disorientation, and greater utility in a head-tracked immersive environment.

Overall, head tracking was perceived to be beneficial in viewing the datasets. One subject commented, "Head tracking was very useful, it was much easier to view the data when moving around." This was noted in several subjects' behavior during task completion. In particular, the clusters task (#4) seemed to be quite well suited for head tracking. There was, however, one subject out of the four who did not find head tracking useful.

5. CONCLUSIONS AND FUTURE WORK

Due to the small number of subjects and design of the experiment, we cannot make any definitive statistical conclusions based on the data recorded. However, the data do show trends in certain instances that are promising for future study:

- A high degree of physical immersion resulted in generally lower times than a low degree of physical immersion.

- Head tracking showed a strong trend in favor of its use. This is apparent in not only task completion times, but disorientation and usefulness ratings as well.

- The combination of a high degree of physical immersion and head tracking seemed to yield the best results, as completion time for most tasks across condition one were lower than for any other condition.

Based on the more subjective results and observations, we can say with certainty that both physical immersion and head tracking overall were perceived to be useful and beneficial, especially when viewing large datasets.

We believe this research to be an important first step toward proving tangible benefits of immersion in information visualization applications, and look forward to continuing this work in greater depth. Our novel strategy for separating the effects of physical immersion and head tracking (in itself a contribution) will be used in the future for more formal experiments.

For a more formal study, we will use a full-factorial within-subjects design, with replication, in order to achieve the maximum statistical power. Additionally, a greater number of tasks need to be designed that take full advantage of head tracking then those currently used. The majority of the current set of tasks may be accomplished by viewing the datasets from afar, where head tracking is less significant.

Several changes in the interaction methods might also prove beneficial in the future:

- Display the numbers of multiple cubes close to the user. This could aid in the selection of a particular data point without having to point directly at it.

- Label the axes with the names of the attributes being visualized. This could potentially reduce confusion and aid in orientation and understanding of the dataset.

- Introduce a different method of navigation and manipulation of the dataset to rotate the data about a point, in addition to the current method of navigating through it (first person), and add these two methods as controlled variables to find out the impact of the navigation methods when looking for benefits of immersion.

6. ACKNOWLEDGEMENTS

7. REFERENCES

[Ahlberg95] Ahlberg, C. and Wistrand, E. (1995). IVEE: An information visualization & exploration environment. Proc. IEEE Information Visualization, IEEE Computer Press, 66-73.

[Anupam95] Anupam, V., Dar, S., Leibfried, T., and Petajan, E.. (1995). Research report: DataSpace: 3-D visualizations of large databases. Proc. IEEE Information Visualization, IEEE Computer Society Press, 82-89.

[Arns98] Arns, L., Cook, D., Cruz-Neira, C.(1998). The Benefits of Statistical Visualization in an Immersive Environment, Proceedings of IEEE VR'99, pp. 88-95

[Cruz-Neira93] Cruz-Neira, C., Sandin, D., and DeFanti, T. (1993). Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. Proc. ACM SIGGRAPH.

[Datey01] 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/

[Hollands95] Hollands, J.G., Pierce, B., Magee, L., Displaying Quantitative Information in Two and Three Dimensions. Proc Human Factors and Ergonomics Society 39th Annual Meeting, 1425-1429

[Kelso02] Kelso, J., Arsenault, L., Satterfield, S., and Kriz, R. (2002).DIVERSE: A Framework for Building Extensible and Reconfigurable Device Independent Virtual Environments. IEEE Virtual Reality, Orlando, FL,

[Lin02]Lin, J., Duh, H., Abi-Rached, H., Parker, Furness, T.,(2002) Effects of Field of View on Presence, Enjoyment, Memory, and Simulator Sickness in a Virtual Environment, Proceedings of the IEEE Virtual Reality Conference 2002, p.164

[Mine95] Mine, M. (1995). Virtual Environment Interaction Techniques. TR95-018, UNC Chapel Hill CS Dept.

[Pausch97] Pausch, R., Proffitt, D., and Williams, G. (1997), Quantifying immersion in virtual reality, SIGGRAPH'97.

[Robertson 93] Robertson, G., Card, S., and Mackinlay, J. (1993). Information visualization using 3D interactive animation. Communications of the ACM, 36(4), April.

[Slater96] 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.

[Spotfire] http://www.spotfire.com/products/decision.asp

[Wickens95] Wickens, C., LaClair M., and Sarno K.(1995) Graph Task Dependencies in 3D Data: Influence of 3D and Color. Proc Human Factors and Ergonomics Society 39th Annual Meeting, 1420-1424