The Effects of Immersion on 3D Information Visualization

Dheva Raja

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
In
Computer Science

Dr. Doug Bowman, Chair
Dr. Ron Kriz, Committee Member
Dr. Chris North, Committee Member

May 22, 2006
Blacksburg, VA

Keywords: 3D Information Visualization, CAVE, Virtual Environments, Scatter Plots, Surface Plots

The Effects of Immersion on 3D Information Visualization

Dheva Raja

Abstract

The effects of immersion with respect to information visualization have rarely been explored. In this thesis, we describe a methodology, two information visualization applications that were developed using the CAVE, and three user studies in order to explore, examine and attempt to quantify the effects of immersion. We focus on three major components of immersion: field of regard (FOR), head-based rendering (HBR), and stereoscopic viewing.

We hypothesize that a high degree of FOR will result in increased task performance and user satisfaction when visualizing data represented by scatter and surface plots. We also hypothesize that HBR and stereoscopic viewing will result in increased task performance, but the effects of these components would be greater in the scatter plots than surface plots.

We have conducted three user studies with the information visualization applications developed for this research. In the first study, an exploratory pilot study, we observed a trend in favor of using high FOR and HBR. In the second exploratory pilot study, we observed a slight trend in favor of high FOR. In the third study, thirty-two subjects performed tasks using both the scatter plots and surface plots with eight test conditions. We observed trends in favor of high levels of FOR, HBR and stereoscopic viewing in scatter plots, a slight trend in favor of HBR for surface plots, and a significant interaction effect between FOR and HBR in scatter plots for a particular task.

Acknowledgements

To my advisor Dr. Doug A. Bowman, whose tolerance, patience, and guidance made my circuitous (and sometimes uncertain) path in a fascinating subject come to fruition.

To my committee members Dr. Chris L. North, and Dr. Ron D. Kriz, a heartfelt thank you.

To Chad Wingrave, friend and labmate, who is always ready to lend a helping hand regardless of obligations, and who also happens to possess an incredible repository of knowledge of all things relating to Mac OS X. You will be missed, my undying El Rodeo lunch companion…regardless of gastro-intestinal ailments related to aforementioned restaurant.

To Umer Farooq, for being there in the early days and always providing sound, logical advice even when I did not want to listen. Your chicken kabob barbecues were missed almost as much as your companionship.

To Jian Chen, fellow labmate and 3DIer, for adding some life to the lab and fun conversation.

To Patrick Shinpaugh, for your never ending help with the complicated CAVE hardware and software, and for the fun banter my fellow Philly friend. Sometime in your life, wherever you are, expect some Philly cheesesteaks to be Fed-Exed to your doorstep.

To Andrew Ray, for your help, your version of DGL, and for your and Patrick's work on DIVERSE and the Linux cluster. You guys rock.

To Lance Arsenault, John Kelso, and Ron Kriz, the creators of DIVERSE. A few words must be said about this software here…to those who don't know, this is the software that makes seamless development between the desktop machines, and the CAVE possible. Using DGL and DIVERSE, I created CaveEquationView on my Apple laptop, then ran it on the CAVE Linux cluster with no modification. Amazing! There are a few of us who realize the power and achievement of this software…

To all the subjects who participated including Finance professors Mike Cliff, Raman Kumar and Randall Billingsley …thanks for your help and enthusiasm!

To my incredible family, without your unconditional love and support things would be incalculably more arduous and laborious.

And finally to Cleanth, your concern about me and prodding about my work mean more than you think.

--Dheva

Table of Contents

LIST OF FIGURES

LIST OF TABLES

Chapter 1: Introduction

This thesis explores the effects of immersion using the CAVE [Cruz-Neira93]. We have implemented two information visualization applications where user tasks were performed and evaluated to assist in examining three components of immersion. We first describe our general motivation for pursuing this work. Next, in an attempt to minimize confusion with regard to domain specific jargon, some widely used concepts and terms are then introduced and defined. Finally, the research questions central to this study, our hypotheses, and a description of our approach complete this chapter.

1.1 Motivation

When describing an immersive virtual environment (VE), it is important to understand the basic characteristics that make a VE "immersive." This immersive quality implies that the VE has several key attributes. These qualities can be, but are not limited to:

1. A "virtual" environment that physically envelops the user.
2. Rendering the image from the user's perspective by means of a device that tracks the position of the user's head.
3. Stereoscopic vision, which makes objects appear to "jump out" of the screen toward the user while also providing strong depth cues.

Very little research has been conducted with information visualization applications in immersive virtual environments. Although prior research in this domain is sparse, some studies have shown that immersion can be beneficial in applications where spatial knowledge of an environment is useful [Datey02]. In an information visualization application, we believe the key attributes we outlined above could allow for more efficient identification of trends in data, greater spatial understanding of the entire data set, easier identification of single data points and more rapid determination of local maxima and minima.

Most applications like Mathematica or Spotfire that are capable of representing data in 3D are used solely on desktop machines. Could viewing a dataset in an immersive environment provide an advantage over similar data visualizations on the desktop? The answer to this question could have profound implications on the manner in which data is displayed in the future and is one of the motivating factors for this research.

While exploring the potential benefits of viewing data in immersive VEs is an important issue, the larger question of whether there is a tangible benefit to using immersive VE technology remains unanswered. A reason for the lack of widespread commercial VE applications could be that there is no empirical evidence that proves any benefit of such systems. Producing empirical evidence proving benefits of immersive systems is the other motivating factor for pursuing this research.

1.2 Definition of Terms

1.2.1 Information Visualization

Commonly referred to as info viz, information visualization comprises an area of research within the field of Human-Computer Interaction (HCI). Info viz is the science and art of mapping abstract data to a visual representation that may be easily perceived, accessed and manipulated by the user[Datey01]. The reader is most likely familiar with the charting and display features in some of the software popular on desktop computers such as Excel, Mathematica, Matlab, SAS, etc. The following are examples of simple 2D graphs that convey abstract data in a readily perceptible form.

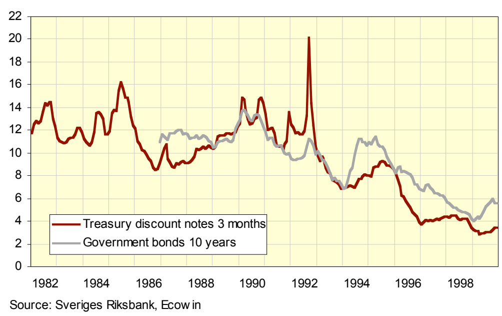

FIGURE 1.1 – BONDS VS. T-NOTES IN SWEDEN

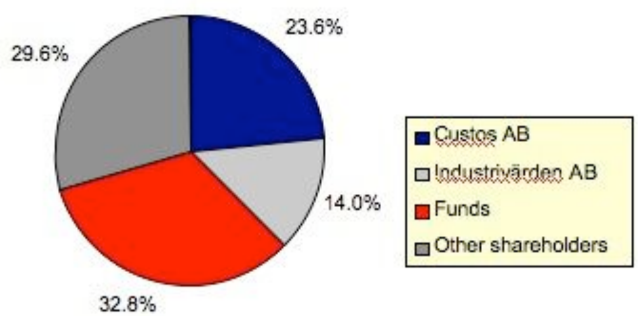

FIGURE 1.2 – DISTRIBUTION OF VOTING RIGHTS AT A LARGE MULTINATIONAL FIRM (SKANSKA)

These examples clearly show abstract data in a form that can more easily be digested and processed by the average person than if the data were only displayed in raw numerical format. While these graphs are indicative of common 2D information visualization techniques, this research concentrates on 3D visualization representations.

1.2.2 Virtual Environments

A virtual environment (VE) is a synthetic, usually 3D world, which is most often rendered in the first-person viewpoint [Bowman05]. This is distinguishable from common desktop environments in that the VE generally possesses a greater level of physical immersion. In our laboratory at Virgnia Tech, the devices used to display a VE are the head-mounted display (HMD) and the CAVE. Our research solely makes use of the CAVE, which will be examined in some detail shortly.

1.2.3 Head Mounted Display

A head mounted display (HMD) is a helmet-like contraption with separate liquid-crystal displays for each eye. The HMD is capable of stereoscopic viewing, and is often coupled with a device that tracks the user's head position and orientation [Bowman05].

FIGURE 1.3 - AN HMD

FIGURE 1.4 – A USER WEARING AN HMD

1.2.4 CAVE

The CAVE [Cruz-Neira93] at Virginia Tech has four screens (front, left and right sides, floor) 3 x 3 meters in size. Each screen has a surface that permits the clear viewing of rear-projected images. The user wears special glasses that synchronize the shutter speed of each eye with the refresh rate of images being drawn on the screens. This synchronization presents a separate image to each eye, which produces the stereoscopic viewing effect. Attached to the user's glasses is a device that tracks the head position and orientation. The computer then renders the image from the user's head position and orientation.

FIGURE 1.5 – A DEPICTION OF THE CAVE IN USE

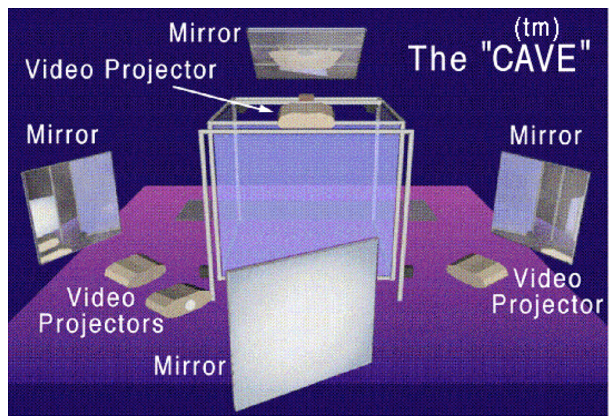

FIGURE 1.6 - COMPONENTS OF THE CAVE

FIGURE 1.7 - INSIDE THE CAVE AT VIRGINIA TECH

1.2.5 Field of View

Field of view (FOV) is the area (measured in visual angle) that the user sees in the virtual world at any instant in time [Bowman05].

1.2.6 Field of Regard

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 [Bowman05].

1.2.7 Clarification of Field of View and Field of Regard

Note that FOV is not the same as FOR. The FOV is the area 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.

1.2.8 Degrees of Freedom

Degrees of freedom (DOF) express the capability of a device to move and be detected in specified directions [Bowman05]. For example, many tracking devices in our visualization lab possess six DOF. A six DOF device typically sends position and orientation data in the form of *{x, y, z, h, p, r}*. The position is specified in the coordinate system as a location of *{x, y, z}*. The orientation is denoted as *{h, p, r}* which represent heading, pitch, and roll and are expressed in degrees. Heading is the angle of rotation about the *z*-axis (height), pitch is the angle of rotation about the *x*-axis, and roll is the angle of rotation about the *y*-axis.

1.2.9 Head Tracking

Head tracking (HT) refers to measuring the user's head location and orientation. A small device attached to the HMD or shutter glasses used in the CAVE feeds the positional data of the user's head into the computer rendering the image.

1.2.10 Head-Based Rendering

Head based rendering (HBR) 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 (as described in the above section on HT), from which the location of the user's eyes is inferred [Bowman05]. The computer rendering the image uses the positional data fed by the head tracker to calculate the image based on the user's viewpoint in the virtual world. Head-based rendering provides an intuitive method of viewing from various perspectives in the VE.

When describing characteristics of a VE, most would use the term head tracking interchangeably with head based rendering (HBR). We prefer to use the term HBR as it more accurately defines an important component of immersion in VEs.

1.3 Immersion and Presence

Immersion and presence are central to the experience of a VE. When describing the characteristics of a VE, these terms 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]. Furthermore, the level of immersion is independent of the user and is directly measurable, while the "presence" one feels is quite difficult to gauge. The level of immersion can readily be controlled in a VE, while presence is an individual's response to a given level of immersion and is entirely unpredictable.

1.4 Components of Immersion

As stated earlier, immersion is directly related to the technology used in a VE. There are a number of elements that comprise immersion in a VE that stimulate various senses in each person. Within each of these sensory channels, some examples of the constituent components of immersion are discussed along with an explanation of differing levels of immersion.

Visual

FOR, FOV, HBR, stereo, display resolution, realism of rendered graphics, and brightness of the display are all components of immersion that directly affect the visual perception of the user. Generally, fully visually immersive systems possess a high level of FOR, and/or FOV and HBR.

An example of a system providing a high level of immersion is the CAVE, which has a large FOR and wide FOV, HBR and stereoscopic viewing. The HMD has a lower level of immersion with a large FOR but smaller FOV, and a device such as a television provides a low level of immersion with narrow FOV, small FOR, and low display resolution.

We define physical immersion to be the extent to which the virtual world surrounds the user in space. The level of physical immersion changes by controlling the FOR, and varies in our experiments between 90° and 270°.

Auditory

Audio that is an element of an immersive environment could provide a realistic representation of actual sound conditions such as those in various simulators, or it may involve a simple beep when an object is selected. Though we did not use audio cues in our information visualization experiments due to the nature of the applications and consequential dubious beneficial influence, audio is thought to be an important immersive element. For example, in certain VEs, audio has a great impact in the treatment of certain phobias like the fear of flying, or in the treatment of war veterans for post-traumatic stress disorders [Brooks99].

Haptic

Haptic (sense of touch) feedback has been used in several applications where a simulated force is helpful in performing tasks when interacting with virtual objects. In the case of the nanoManipulator [Brooks99], the visual quality of the display type is far less important than the feedback provided by the haptic device, as the scientist may not be able to see the operation performed, but can feel quite accurately. In this case, the haptic feedback is very important in eliciting a sense of presence in the user.

In our research, we focus on three major components that fall in the *visual* category: physical immersion, HBR, and stereoscopic viewing. We select these three components because we feel that they are the most important elements with regard to information visualization applications using the CAVE. While resolution, audio cues, and haptic feedback are other important elements of immersion, we believe that they are less important in info viz apps because of the manner in which one interacts with the data. For example, having a high FOR would seem to be more beneficial in determining trends than an audio cue, since one would be able to view more of the dataset at any one glance.

HBR provides an intuitive method of viewing the virtual world from the first-person perspective as the scene is updated and redrawn depending upon the user's head position. This type of tracking provides a strong depth cue (motion parallax) and is quite helpful in navigating about the virtual world.

Stereoscopic viewing is achieved by providing separate images to each eye with the aid of special glasses. This can also provide a strong depth cue, though the perception may vary widely among users.

The above components are the ones we feel to be the most influential in comprising immersion and eliciting a sense of presence in the user when viewing info viz applications in the CAVE.

When a comparison is made between the FOR and another component of immersion we chose not to examine like the screen resolution, it is readily apparent that resolution is not as important as FOR. In our scatter plot application, cubes are drawn that do not require detailed resolution. The FOR, however, can be crucial in determining how one variable relates to another quickly by simply turning one's head. The HBR is instrumental in providing strong depth cues (along with stereo) that can aid in the selection of particular points in both scatter and surface plots. By comparison, an audio cue, would not tell the user much about the dataset in question. While these other components of immersion (resolution, audio, realism of rendered graphics) might be quite successful in creating a strong sense of presence in the user in other application domains, the three components we chose to examine we feel could have the greatest influence in information visualization applications.

1.5 Motivation for Using the CAVE

The CAVE permits 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 since most HMDs offer only a limited FOV (e.g. 60° diagonal).

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. The important distinction here is that in an HMD-based VE system, head tracking *creates* the high level of FOR, so there is no way to separate these two components. This is the primary impetus behind our use of the CAVE for quantifying the effects of immersion.

Prior efforts aiming to measure the effects of immersion, in our view, were unable to distinguish between the components of immersion. 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. We later present a method, based on CAVE-like displays that addresses these challenges and is focused on the evaluation of the effects of FOR, head-based rendering, and stereoscopic viewing.

1.6 Research Questions

There are several questions that this body of research attempts to address. After each question there is a brief section expanding upon the thoughts posed, and avenues for potential exploration.

1. Are there tangible benefits to an immersive environment while visualizing 3D data?

Would a user be able to determine a relationship in the data more quickly in an immersive environment than one that is non-immersive? Would a user be able to locate a particular point more quickly in an immersive VE? The ‘tangible’ in the research question refers to a measurable metric like time or user satisfaction with the particular condition.

First, in order to support a claim of an effect of immersion on a 3D information visualization application, it is important to know which component of immersion is producing the effect. If a high FOR and HBR are being used, and this system provides a clear advantage over the system not employing these components, is it the FOR, the HBR, or the combination of both components that produced the better results? When the component of immersion that created the benefit is known, efficient, cost effective systems can be designed to take advantage of this knowledge. Hence, we must devise a manner in which to separate the components of immersion and then measure each component's effect on the outcome. In our approach outlined shortly, we present a method where it is possible to statistically support claims that components of immersion affect the outcome in a measurable quantifiable manner.

2. *In densely packed datasets where many points are occluded, does a higher level of immersion provide any benefit?*

In a scatter plot for example, viewing a densely packed 3D dataset on a desktop where occlusion is an issue could result in difficulty locating a point or identifying trends. If the scaling were readjusted such that occlusion were not as much of an issue, a narrow FOR (like that of the traditional desktop) would make determination of trends a more laborious affair than those systems with a wide FOR.

3. *In large data sets, would an increase in the field-of-regard (FOR) in immersive environments help the user to more easily identify trends?*

As mentioned above, we think that a greater FOR would increase spatial awareness of the user potentially resulting in reduced task times, especially where occlusion is an issue. The user could simply turn her head to gain another perspective of the dataset as opposed to the more time-consuming method of using navigation to change viewpoint.

Several studies have shown that the greater spatial awareness and understanding by the user in an highly immersive environment should translate into a measurable and quantifiable difference when compared to the same scenario on a lower immersive system or desktop machine [Slater96] [Datey01] [Gruchalla04]. Specifically we are referring to faster response times when the user is queried regarding data trends.

Although a comparative study between the CAVE, HMD, and desktop was initially considered, we felt that the confounding of variables like FOR, resolution, interaction techniques, etc., would compromise the determination of the cause of potential differences between the hardware used. Furthermore, a number of comparative studies have been conducted with the aforementioned hardware with varying results. Please see Chapter 2 (Related Work) for summaries of comparative and practical studies.

4. *Do the effects of immersion in 3D information visualization produce similar outcomes across different representations of the data, i.e., scatter plots, surface plots, etc.?*

There are a quite a few representations of 3D information visualization applications. These include line, pie, and bar charts, histograms, scatter and surface plots, and many other designs. Our study examines the effects of immersion on 3D scatter plots and 3D surface plots in the CAVE. Other types of info viz representations might be examined in future work in order to determine the influence of various components of immersion across multiple formats of 3D info viz.

1.7 Statement of Impact

Very little research has been conducted to prove if there are tangible benefits of immersion. 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 effects 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. A substantial portion of this research is aimed at producing empirical evidence of increased user performance, usability, understanding, etc. due to the use of immersive VEs. In the info viz domain, many highly important analyses are conducted countless times daily, from medical research to financial data. If our research shows that certain elements of immersion produce beneficial results, then cost effective systems could help people that use info viz perform their tasks and analysis more efficiently.

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.

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 described in part of this research provides a framework for empirical studies that can provide such proof since we are able to independently evaluate the components of immersion. The last part of the prior sentence 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.

1.8 Hypotheses

1. *A high degree of FOR will allow higher levels of task performance and greater user satisfaction/ease of task completion when visualizing datasets represented by 3D scatter plots and 3D surface plots.*

We hypothesize that a larger FOR will result in reduced response times, and that the user will accomplish the task at hand more easily than in the low FOR conditions. Due to the increased spatial awareness that the high FOR condition affords, we think that trend identification in both 3D scatter plots and 3D surface plots will result in lower response times and greater user satisfaction.

2. *Head based rendering will allow higher levels of task performance and greater user satisfaction/ease of task comlpletion when visualizing datasets represented by 3D scatter plots and 3D surface plots. This effect will be less pronounced on surface plots, than in scatter plots.*

We think that HBR will also assist in the spatial awareness of the user and result in lower response times. HBR might not be as immediately detectable as a shift in the FOR, but we believe that HBR could be very helpful, especially in densely packed datasets where many points are occluded.

3. *Stereoscopic viewing will allow higher levels of task performance and greater user satisfaction/ease of task completion when visualizing datasets represented by 3D scatter plots and 3D surface plots. However, stereoscopic viewing will have a greater effect on 3D scatter plots than it will have on 3D surface plots.*

Stereoscopic rendering, in our view, may be even less apparent to the user initially than HBR. However in scatter plots, this characteristic (like HBR) may provide the user with strong depth cues. This could be quite useful in determining maximum and minimum points in a region.

We think that stereoscopic viewing would more helpful in scatter plots, than in surface plots due to the sense that one is able to obtain stronger depth perception with 3D scatter plots. When the user is very close to a portion of the surface plot, it is difficult to gauge depth, as the plot envelops the user and looks somewhat flat. Nonetheless, we feel that stereoscopic viewing would aid in the location of maximum and minimum points on a surface plot.

1.9 Our Approach

We have developed two 3D information visualization applications for use in the CAVE in an attempt to answer the research questions and test our hypotheses. The first is CaveDataView, which displays scatter plots in the CAVE that are read from a data file consisting of coordinates that denote points in space. The points are then displayed as cubes in 3D within the CAVE. The second application is CaveEquationView, which displays surface plots from a simple mathematical equation that the user types in, or from a file of points that have been computed from mathematical formulae. The plot is then displayed in 3D within the CAVE in a solid color, or with height dependant shading.

We conducted three separate experiments with these two 3D information visualization applications. The first experiment was a pilot study using CaveDataView that examined two components of immersion, FOR and head based rendering. The second pilot study used CaveEquationView to display stock options pricing and related component data. Here we examined FOR, head based rendering, and stereoscopic viewing. The third experiment was a formal one that combined two scatter plots from CaveDataView, and two surface plots from CaveEquationView. In this last experiment we examined FOR, head based rendering, and stereoscopic viewing.

In order to separate the components of immersion, which is crucial to understanding where the cause of potential benefits might originate, we explain our novel method in some detail below that explains how this separation may be achieved.

As noted in the introduction, the three primary components of immersion we examine in this research are FOR, head-based rendering and stereoscopic viewing. In order to evaluate the benefits of immersion, we needed a method that allowed us to separate the effects of these 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 FOR in such systems. Using the CAVE 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, the 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 1.8). We did not measure the effect of FOV in our research, but could in future experiments.

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.

In order to control the third component, stereoscopic viewing, a stereo or mono display driver is specified before the application is launched. The user still wears the special glasses during the mono condition to hold the FOV constant. If the glasses were removed, the brightness of the display would increase markedly, the FOV be confounded, and this would also provide a cue to the subject that stereo is about to be turned off. This confounding of the FOV and additional cue could be detrimental to the outcome of the experiment by obscuring the source of any effect.

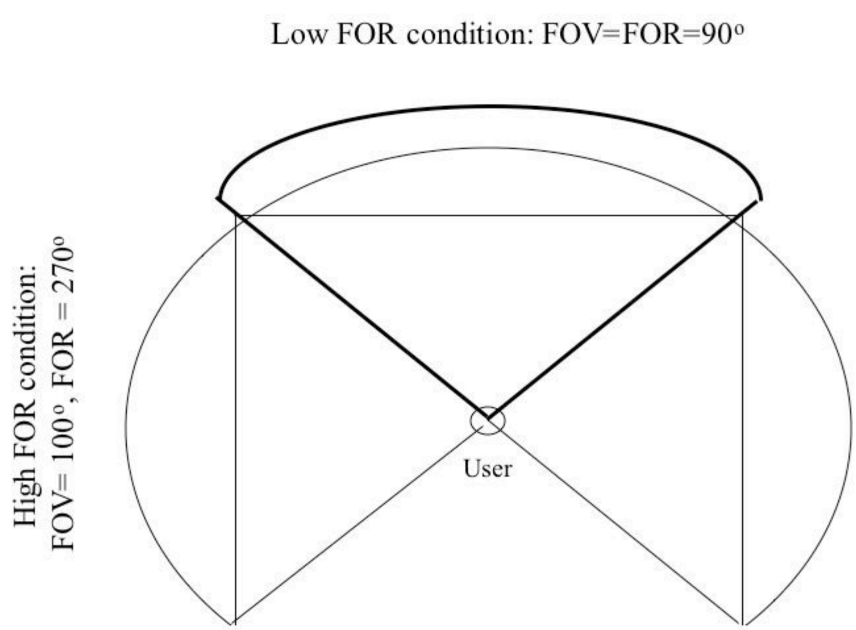

FIGURE 1.8 - TOP-DOWN VIEW OF THE CAVE INDICATING FOV AND FOR FOR THE TWO LEVELS OF PHYSICAL IMMERSION

High FOR				Low FOR		
	Stereo	**No Stereo**			**Stereo**	**No Stereo**
Head Based Rendering	4 screen CAVE, HBR, Stereo	4 screen CAVE, HBR, No Stereo		**Head Based Rendering**	1 screen CAVE, HBR, Stereo	1 screen CAVE, HBR, No Stereo
No Head Based Rendering	4 screen CAVE, No HBR, Stereo	4 screen CAVE, No HBR, No Stereo		**No Head Based Rendering**	1 screen CAVE, No HBR, Stereo	1 screen CAVE, No HBR, No Stereo

TABLE 1.1 - 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.

1.10 Overview

This thesis is divided into seven chapters. The first chapter introduces key terminology, concepts, and the motivation for this research. The requisite methodology and framework are described in detail in order to provide some answers to the research questions posed, and to test our hypotheses. The next section covers the related work that has influenced this work. Chapter 3 discusses the 3D information visualization applications developed for our experiments in some detail. In Chapter 4, we describe the first exploratory pilot study involving scatter plots and the results of the experiment. Chapter 5 describes the second pilot study with surface plots of options data and results. The formal experiment combining both scatter plots and surface plots is examined in Chapter 6. Finally we conclude with our contributions and suggest areas for future work in Chapter 7.

Chapter 2: Related Work

Some prior information visualization techniques and applications in VEs that have helped to influence and motivate our work are summarized below. We begin with an explanation of basic human visual perception relating to depth cues and the importance of this in relation to immersive qualities in task performance. These theories provide a bit of the foundation behind the idea that immersive technologies could prove beneficial in a number of different domains, including info viz.

In general, prior research attempting to measure the benefits of immersion has made use of two different strategies: practical experiments and controlled experiments. Relevant studies that can be placed in either experimental category are compared and contrasted, and the primary factors that differentiate our work from the reviewed studies are noted.

2.1 Visual Depth Cues and Immersive Environments

In examining the basic components of immersion, is there a reason to believe that FOR, HBR, and/or stereoscopic viewing might have a beneficial impact on task performance and spatial understanding? In attempting to answer this question, some background information about the manner in which humans perceive visual depth to gain an understanding of the 3D structure of the scene is crucial. Visual depth cues can be broken into four categories [Bowman05]:

Monocular, Static Cues

These cues refer to depth information that may be garnered by a single eye when viewing a static image. These cues include the following: relative size, height relative to the horizon, occlusion, shadows and lighting, and texture gradients.

Oculomotor Cues

These are depth cues that occur from the muscular tension in the human visual system, named *accommodation* and *convergence*.

Motion Parallax

Depth information is related to the viewer by objects closer to the viewer moving more rapidly than objects further away. This can happen by the objects moving relative to the viewer, the viewer moving relative to the objects, or some combination of the two. In a sparse 3D environment (usually the case in info viz applications) motion parallax can provide extreme strong depth cues.

Binocular Disparity and Stereopsis

This refers to the slightly different image that each eye sees of the environment. Binocular disparity is more pronounced as the object's proximity to the viewer nears. The fusion of the two images that each eye sees also provides a strong depth cue. This effect is termed stereopsis.

These depth cues described above are instrumental in assisting the viewer in developing spatial understanding of the environment. For example, in a picture of a city skyline, the human visual system can determine from the static depth cues in the photo (occlusion, relative position to horizon, etc.), which buildings are furthest from the camera. However, different types of displays produce varying levels in the perception in the "strength" of the depth cues. A monitor found on a typical desktop PC with HBR, produces weaker depth cues than surround screen displays like the CAVE with HBR, or in HMDs with HBR, which are generally considered more physically immersive environments [Bowman05]. Since the depth cues are of primary importance in helping the user to build a spatial understanding of the environment, a high physically immersive VE, like the CAVE, produces a better approximation of the cues one would experience in the real world than less physically immersive systems. Hence, the more realistic cues experienced in immersive systems has led us to believe that immersion might provide a benefit in viewing information visualization applications in the CAVE.

2.2 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 scatter plots in Spotfire are very similar to the plots we display in CaveDataView.

2.3 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 scatter plots. Wizard provided both a small hand-held overview of the dataset and a larger version of the data through which the user could navigate. The Wizard application was part of a study that can be categorized as a practical experiment, and is summarized in the next section (Figure 2.2).

2.4 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 must not be generalized since the systems being compared typically have many differences (confounds). Below are several 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, FOR, 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 similar and recent study used two general-purpose visualization packages available for desktop systems, and coupled their use to an existing, relatively mature VE system [Boyd03]. In the VIVRE system, IRIS Explorer and AVS/Express were used to display data visualizations on both desktop systems and the dVISE VE. Visualizations in this study included:

- Combustion and gas flow in industrial processes
- Potential leakage paths surrounding underground storage caverns
- Dynamic operation of a flexible control valve
- Complex natural microstructures and their properties

While the authors found that there was no performance difference between the desktop system and VE for small datasets, there was some consensus that for significantly large datasets the VE was superior for data manipulation. Additionally, the users reported the increased spatial awareness of the VE was beneficial when interacting with cutting planes in the 3D scene, especially when there were a large number of data points. In this study, FOR and resolution differed in the display devices, so it is difficult to explain the source of the benefit the VE provided the users.

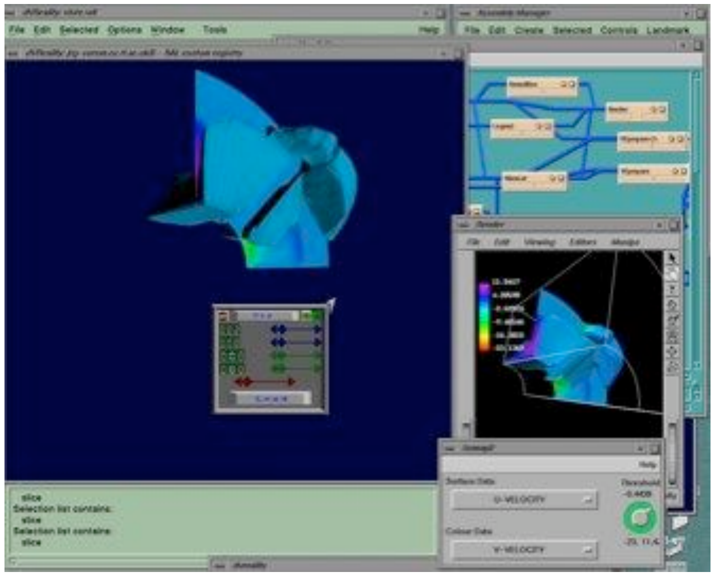

FIGURE 2.1 - VIVRE SYSTEM DISPLAYING FLUID FLOW IN THE VE *(LARGE WINDOW TO LEFT)* AND DESKTOP *(SMALL WINDOW, LOWER RIGHT)*

Finally, 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, FOR and head-based rendering were confounded in this study.

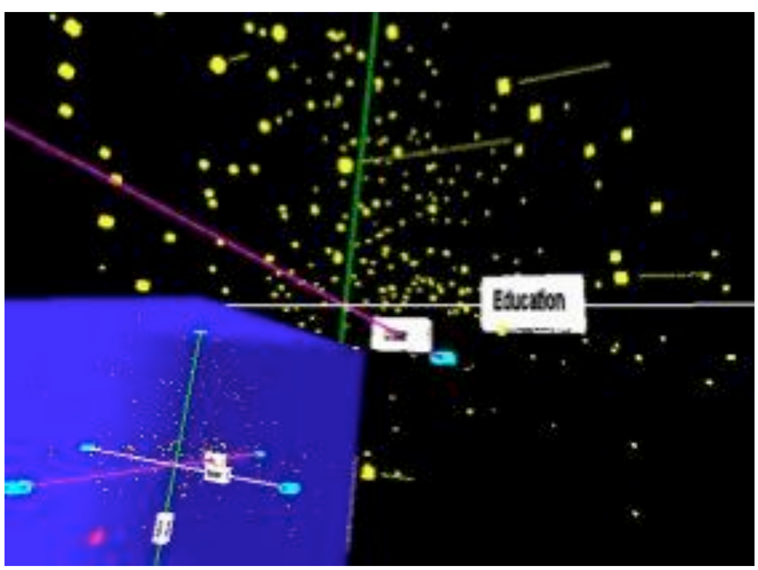

FIGURE 2.2 - WIZARD WITH OVERVIEW AND DETAIL DISPLAY OF DATASET

2.5 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* FOR, 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.

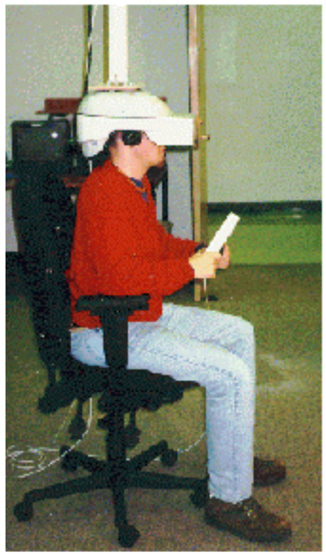

FIGURE 2.3 – NO, IT'S NOT A VE TORTURE DEVICE, THIS IS THE HMD USED IN PAUSCH'S STUDY

2.6 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 levels of encumbrance are equivalent under all testing conditions.

Chapter 3: System Design and Description

3.1 CaveDataView

CaveDataView is the first VE application we developed to test our hypotheses. It can display scatter plots 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 2.4.1 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 (DSO) 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.

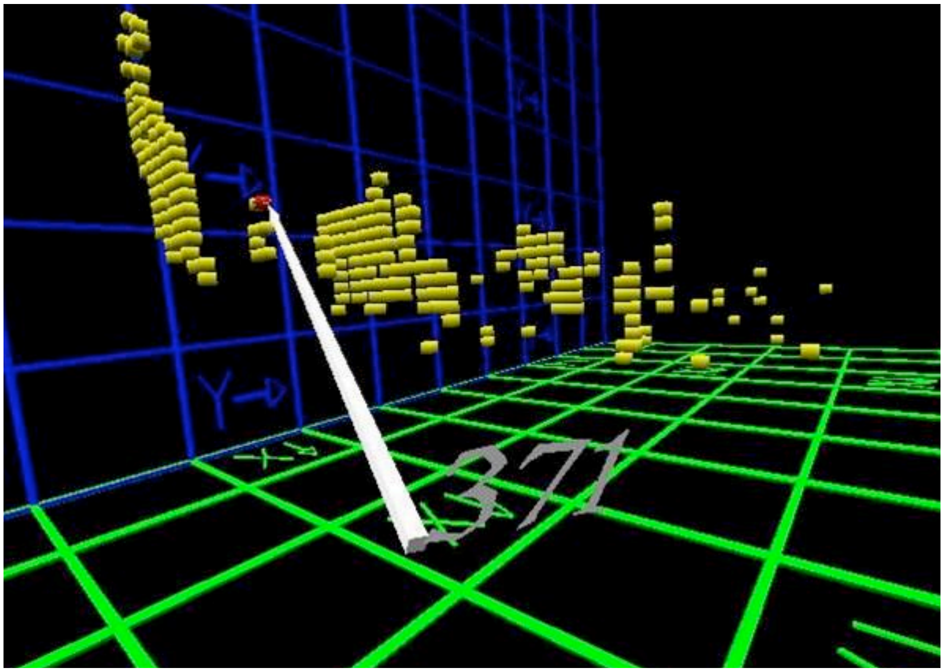

FIGURE 3.1 - A 3D SCATTER PLOT IN CAVEDATAVIEW

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.

Our CAVE is a four-screen 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. The user points the wand in the direction desired, then pushes forward on the joystick to move in that direction. The buttons to the left and right of the joystick rotate the dataset about the *Z-axis* counter-clockwise and clockwise, respectively. 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 3.2 shows a user standing in the CAVE while CaveDataView is running.

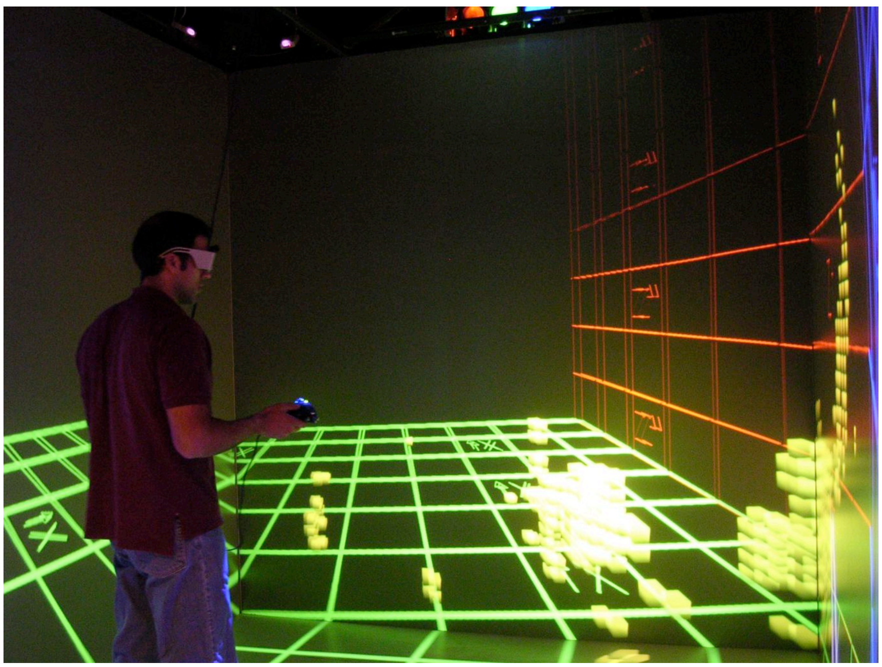

FIGURE 3.2 - CAVEDATAVIEW IN THE CAVE

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 3.4). 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.

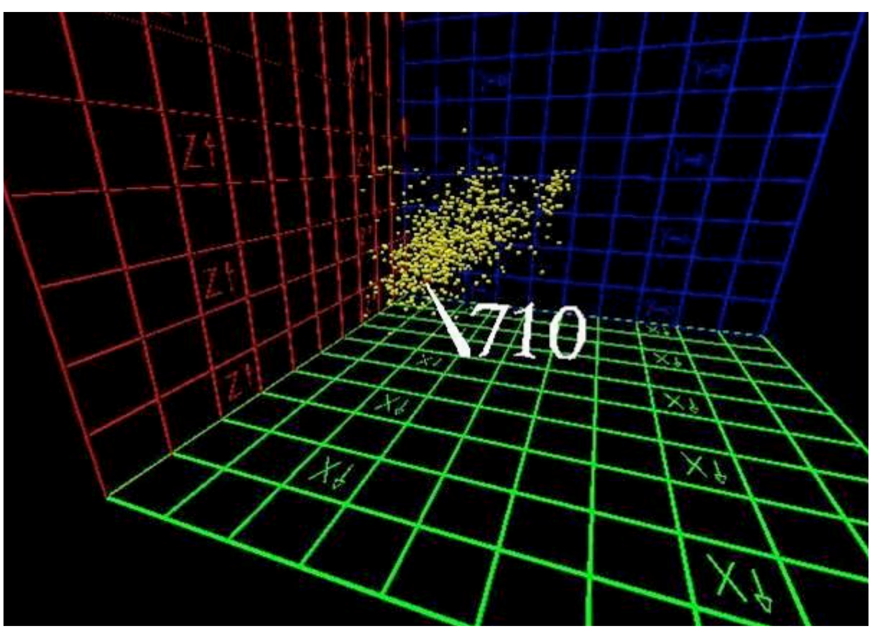

FIGURE 3.3 - RAY-CASTING AND DATA POINT LABELING IN CAVEDATAVIEW

3.2 CaveEquationView

CaveEquationView is the second VE application we developed to test our hypotheses. Similar to CaveDataView, CaveEquationView was developed in C/C++ using DIVERSE. As mentioned previously, a powerful feature and significant benefit regarding development with DIVERSE is that a desktop machine can be used to create the application, freeing valuable time on the CAVE for other purposes.

Unlike CaveDataView, CaveEquationView was developed using DGL (not DPF). DGL is the DIVERSE interface to OpenGL, allowing great flexibility in the application as the code is written in low-level OpenGL commands. We tested CaveEquationView under DGL 2.4.1 and DTK 2.4.14, which run on Mac OS X, Linux, and Irix operating systems. This application also makes use of DTK to encapsulate tracking information, and uses the DSO architecture to seamlessly migrate to different operating systems and environments.

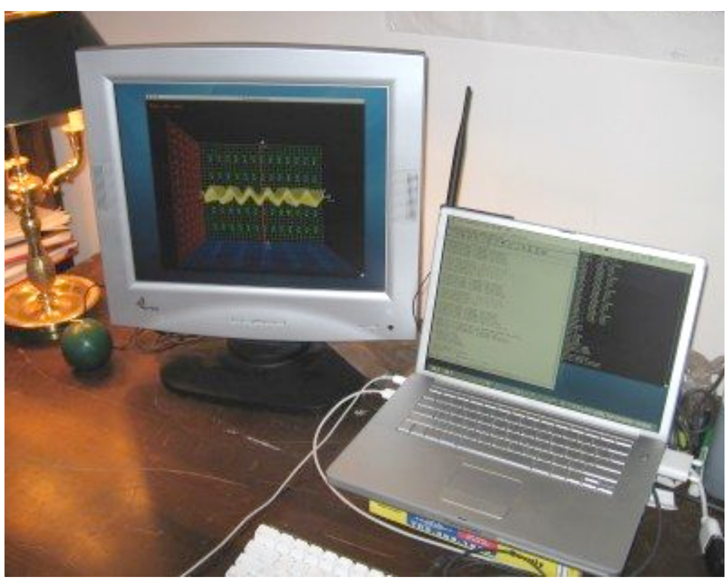

FIGURE 3.4 - CAVEEQUATIONVIEW DESKTOP DEVELOPMENT WITH DGL AND DIVERSE ON MAC OS X

Initially developed to display the surface plots of various mathematical equations, CaveEquationView prompts the user to enter an equation, which the application then parses, computes, and displays. During the course of more exploratory work, it was quickly determined that the need to display more complex surfaces necessitated the adaptation of CaveEquationView to read in pre-computed points from data files. In that respect it is quite similar to CaveDataView as both applications read in tab-delimited files of three column data. The data is then scaled to fit in the VE (or desktop window), and the resultant surface plot is shown. A grid and variable axis are also drawn in the application to aid in user orientation. Axis labeling, along with minimum and maximum points of each variable are included with the plot to assist the user with pertinent information (see Figure 3.8). Lighting and Gouraud shading are utilized in the plots, which show surface contours, slope, and other plot details (see Figure 3.6).

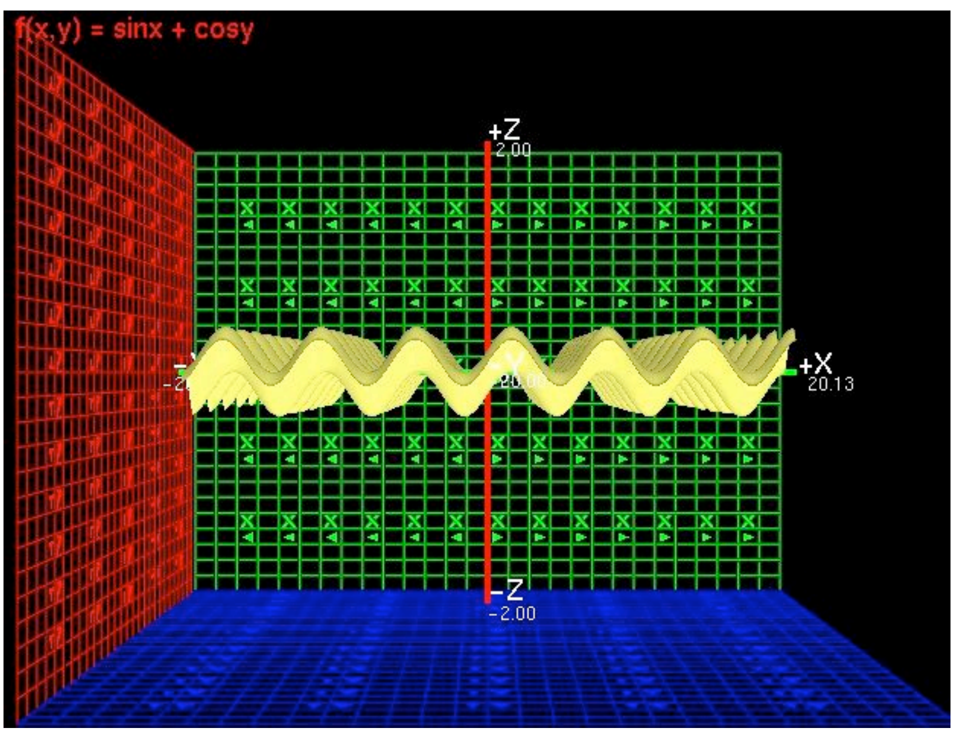

FIGURE 3.5 – CAVEEQUATIONVIEW DISPLAYING F(X,Y)=SIN(X) + COS(Y)

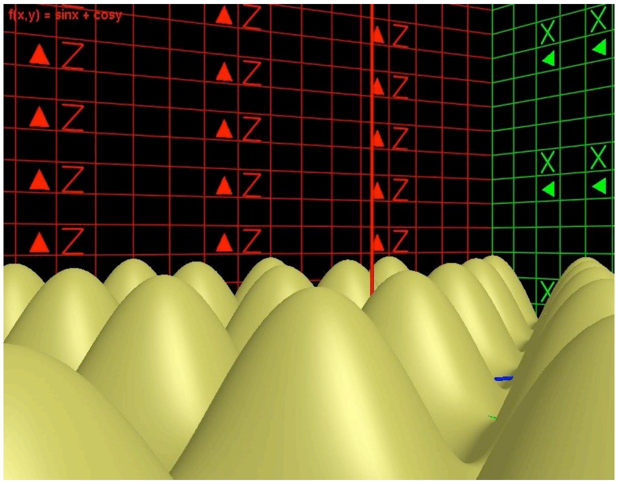

FIGURE 3.6 - IN THE MIDDLE OF THE PLOT SHOWN IN FIGURE 3.5, WITH SHADING AND LIGHTING IN CAVEEQUATIONVIEW REVEALING PLOT CONTOURS

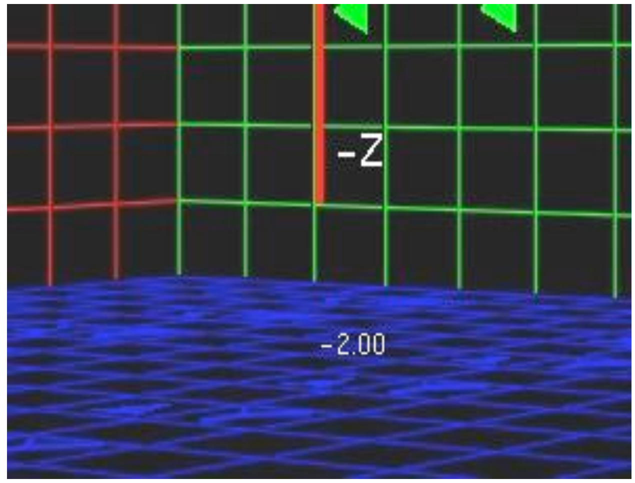

FIGURE 3.7 - AXIS AND MAX/MIN LABELING

The description of the CAVE stated above in section 3.1 could be reprinted here to describe the setup used with CaveEquationView noting one exception. When CaveDataView was developed, the machine that rendered the images on the CAVE walls was an SGI supercomputer circa 1997. Shortly after the first exploratory pilot study, a Linux cluster was developed and built at Virginia Tech that could also be used to drive the images in the CAVE. Since the performance of the cluster was significantly greater than the SGI, the number of data points used in the surface plots was increased from 40,000 to 360,000 with no frame rate degradation discernable to the user. Both the number of points in the plot, as well as the frame rate, are quite important when running subjects in an experimental setting. The frame rate must be rapid enough so that when the user is navigating through the dataset, no delay (or jerkiness) is apparent. If the movement is not smooth, the experience is very distracting and may detract from the immersive aspects of the CAVE, and break the user's sense of presence. While not as important as the frame rate may be, the number of points in the plot are useful because as the number of points are increased, the curvature of various parts of the plot appear smoother. This effect is noticed to a much greater extent in the CAVE than on the desktop. In the CAVE, the user can be completely immersed in the plot and examine edges that appear a few inches in front of her face. With smoother contours, it may be possible to see properties of the plot that were not as apparent with jagged edges.

Navigation through the dataset is performed exactly as described in section 3.1. The wand has a joystick and four small buttons, two on the left and right sides of the joystick. The user points the wand in the direction desired, and intuitively pushes forward on the joystick to move in the desired direction.

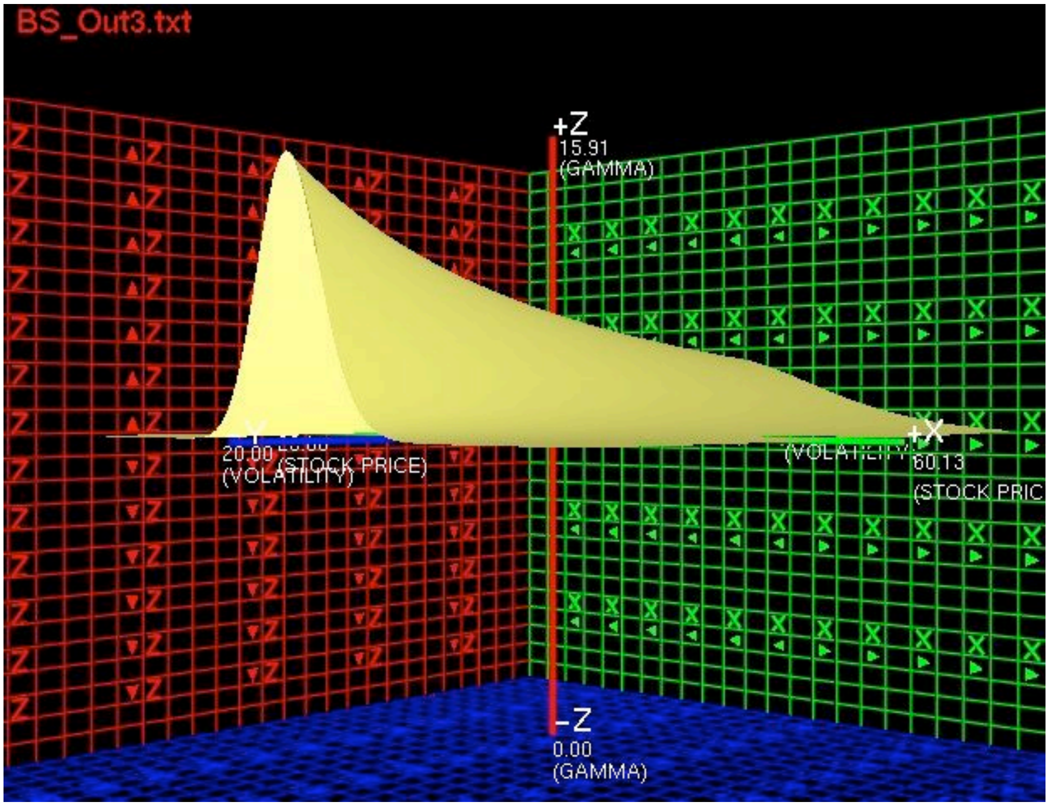

FIGURE 3.8 - PRE-COMPUTED DATA POINTS (STOCK OPTIONS COMPONENTS) DISPLAYED IN CAVEEQUATIONVIEW

As will be explained in Chapter 5, it was decided to use pre-computed points for surface plots in the second pilot study. Figure 3.8 shows a plot of several components that comprise a stock option based on the Black-Scholes options model (for more information regarding stock options and the pricing model used for computations please see Appendix D).

3.2.1 Changes Implemented in CaveEquationView for Final Experiment

Based on several suggestions from knowledgeable subjects explained in detail in Chapter 5, we decided to modify CaveEquationView in two ways after the second pilot study. The first modification was to add a height (*z-axis*) dependent contour shading scheme to the application.

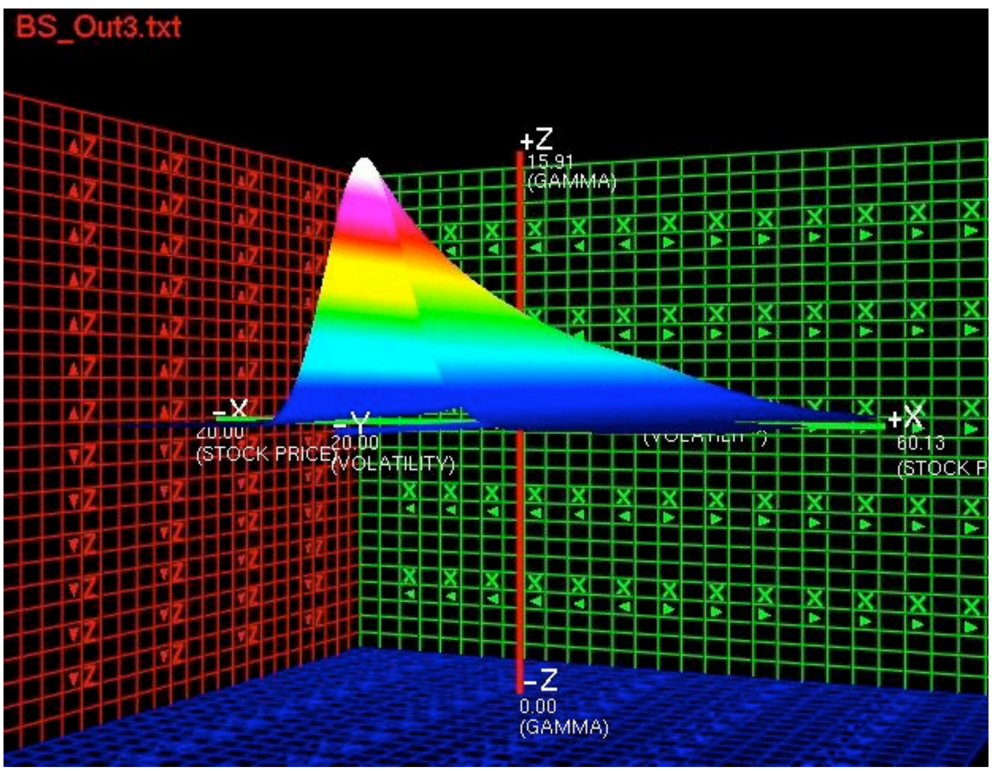

FIGURE 3.9 – THE SAME PLOT SHOWN IN FIGURE 3.8 WITH HEIGHT-DEPENDENT CONTOUR SHADING

The second modification involved the navigation methods used to interact with the plot. As mentioned previously, there are four buttons on the wand, two on the left, and two on the right side of the joystick. Previously when these buttons were pressed, the speed at which the user traveled through and around the dataset was increased or decreased. For the last experiment (Chapter 6), pressing either of the two buttons to the right of the joystick rotates the surface plot clockwise about the *z-axis*. Pressing either one of the two buttons to the left of the joystick rotates the surface plot counter-clockwise about the *z-axis*. This navigation modification was helpful in making data manipulation and understanding more rapid and intuitive than the previous method.

Chapter 4: Pilot Study (Scatter Plots with CaveDataView)

4.1 Experimental Design

In the CAVE, we separated the two components of immersion examined, FOR and HBR, and analyzed their effects independently. This led to our experimental design, shown in table 4.1. In the high FOR condition, we use all four walls of the CAVE, while in the low FOR 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 high FOR, 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 low FOR, 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 Based Rendering	No Head Based Rendering
Low FOR	One screen CAVE	One wall CAVE, no HBR
High FOR	Four screen CAVE	Four wall CAVE, no HBR

TABLE 4.1 – EXPERIMENTAL CONDITIONS

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

4.2 Task Description

We developed six 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.

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 scatter plot. We felt this was an important basic task in that it is a key component in gaining understanding of a single data point.

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.

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.

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.

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).

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 six tasks once, and saw all four conditions at least once. Table 4.2 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 4.2 – SUBJECT TASK ORDER

4.4 Results

This section is organized by exploring the results pertaining to the first hypothesis regarding FOR, followed by the second hypothesis regarding head based rendering.

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

Condition 1 = four walls with head based rendering

Condition 2 = four walls with no head based rendering

Condition 3 = one wall with head based rendering

Condition 4 = one wall with no head based rendering

4.4.1 FOR

In examining FOR, 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 FOR 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 high FOR 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 FOR in completing the tasks.

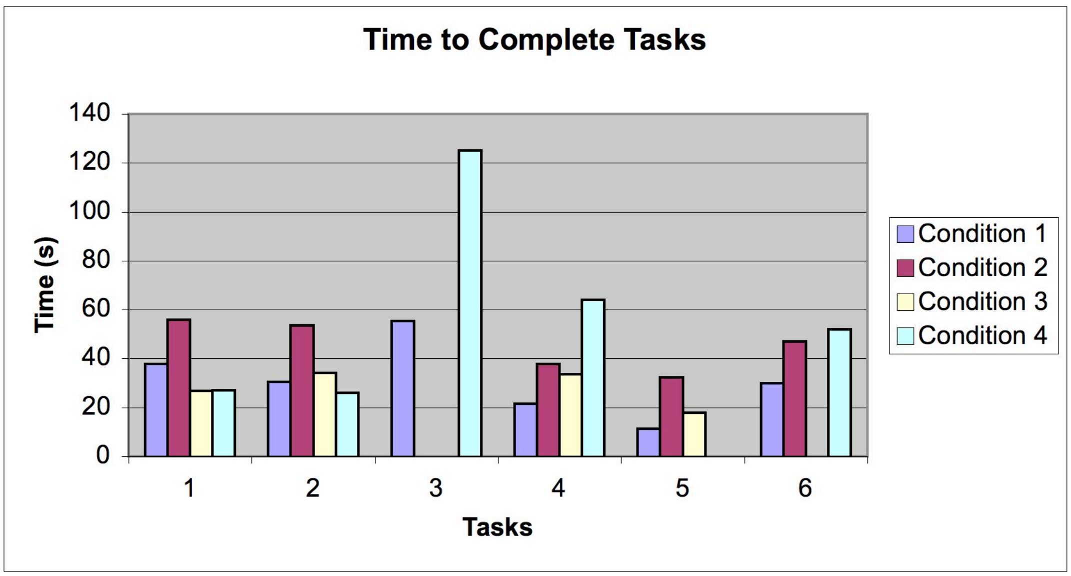

FIGURE 4.1 – TIME TO COMPLETE THE TASKS IN THE FOUR EXPERIMENTAL CONDITIONS

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 FOR. 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-Based Rendering

For head based rendering, 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 based rendering 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 based rendering. There was, however, one subject out of the four who did not find head based rendering useful.

4.5 Conclusions

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 FOR resulted in generally lower times than a low degree of FOR.
- Head based rendering 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 FOR and head based rendering 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 FOR and head based rendering overall were perceived to be useful and beneficial, especially when viewing large datasets.

We believe this pilot study to be an important first step toward proving tangible benefits of immersion in information visualization applications. Our novel strategy for separating the effects from the components of immersion is used for the remainder of experiments conducted in this body of research.

For the more formal combined study to follow, 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 based rendering then those currently used. The majority of the current set of tasks may be accomplished by viewing the datasets from afar, where head based rendering is less significant.

Several changes in the interaction methods might also prove beneficial:

- 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.

Chapter 5: Pilot Study (Surface Plots with CaveEquationView)

5.1 Experimental Design

In the previous pilot study, we discuss how the components of immersion are separated and later examined and analyzed. In this pilot study, we continue along the same lines of thinking with respect to design and separation of components. However, besides examining FOR and head based rendering we add a third component, stereoscopic viewing. This leads to a total of eight conditions as initially described in section 1.7, depicted in the following table:

High FOR				**Low FOR**		
	Stereo	**No Stereo**			**Stereo**	**No Stereo**
Head Based Rendering	4 screen CAVE, HBR, Stereo	4 screen CAVE, HBR, No Stereo		**Head Based Rendering**	1 screen CAVE, HBR, Stereo	1 screen CAVE, HBR, No Stereo
No Head Based Rendering	4 screen CAVE, No HBR, Stereo	4 screen CAVE, No HBR, No Stereo		**No Head Based Rendering**	1 screen CAVE, No HBR, Stereo	1 screen CAVE, No HBR, No Stereo

TABLE 5.1 – EXPERIMENTAL CONDITIONS FOR SURFACE PLOT PILOT STUDY

Tasks involving eight different datasets for the eight conditions were performed, and the times to complete these tasks were recorded. Subjects provided usefulness and difficulty ratings on a seven-point scale after the tasks were performed for each condition. These ratings were not asked after each task like the previous pilot study, but rather after the tasks were completed to gauge the subject's opinion of the usefulness and difficulty of performing the specified tasks in each condition.

We decided for this pilot study that in order to maximize the utility of CaveEquationView in a formal experiment, the surface plots should represent real world functions, and not simply abstract mathematical representations. In this manner, we felt that giving the tasks a concrete meaning and purpose would improve accuracy of the results and facilitate the analysis of the data. Furthermore, we felt that suggestions for future improvements in CaveEquationView would be readily forthcoming if the subjects could relate to the plots and tasks with a good degree of familiarity.

Since real-world applicability was a primary factor in deciding which domain area was to be the basis for displaying the surface plots, this computer science master's thesis student combined several long-term interests and professional experience in formulating this surface plot pilot study with the plots of stock option data.

Stock options were introduced in the US markets for listed trading in 1973. Fischer Black and Myron Scholes formulated the theory behind the pricing of these securities a few years before the introduction of trading. This work appeared in a seminal publication [Black73]. The Black-Scholes pricing model has become the most widely used method to determine the value of these derivatives around the world by traders, theorists and other finance practitioners. Incidentally, this work later was later recognized with a Nobel Prize in Economics some 24 years after the publication. As this student's background was in options trading, we thought displaying surface plots that contain varied information related to the Black-Scholes model could be an effective domain in which to run a pilot study. Please see Appendix D for a basic primer on stock options, and the underlying mathematical variables and formulae related to the Black-Scholes model.

5.2 Task Description

The tasks for this pilot study were designed to incorporate a combination of spatial understanding/awareness of the dataset and a command of the underlying theory and components of stock options. These tasks may be placed into one of the following general categories (for the specific questions and tasks performed, please see Appendix B.3):

Trend Determination

As described similarly in the first exploratory pilot study, this task requires the subject to obtain a general sense of orientation and understanding of the dataset in order to identify trends within the plot. An example of a question in this category could be "As time increases, what happens to the price of the option?" This type of task is fundamental to any sort of analysis with 3D graphs in information visualization.

Maximum/Minimum Point Estimation

Subjects were asked to find the maximum or minimum "height" of the surface plot, and to name the corresponding coordinates of that max/min point. We felt this was an important and relevant task as one might often want to know exactly where these points are in order to gain a greater insight into the dataset, especially with financial applications.

Range Estimation

In this type of task, subjects were asked to define a range of values for which a certain variable or component exhibited a behavior. For example, a question that might typify this task could be, "What range of values in the stock produce an increase in the delta?"

Interpretive

With an interpretive task, the subject was asked to examine certain aspects of the surface plot, and make inferences regarding the importance or influence of these components in relation to one another. While this task presented some difficulty in being able to objectively measure by a timing metric, this type of task is crucially important to gaining an understanding of the dataset and then translating that knowledge to further an educational insight in the behavior of the components of the option model. This master's candidate often wished that he had viewed surface plots similar to those displayed in this pilot study before trading, as some behavior exhibited in the pricing components is definitely not intuitive.

Slope Comparison

This kind of task involved the comparison of the slope of the surface plot at one specified point, to the slope of the plot at another point. We feel this task is important and has significant real-world relevance. It requires the subject to rapidly ascertain an understanding of the dataset and presents an ideal manner in which to compare the effects of the components of immersion.

5.3 Subjects

Due to our decision to use a real-world application requiring domain expertise, we recruited eight subjects with knowledge of stock options. Seven of the subjects were male, and all had some familiarity with the Black-Scholes pricing model. Several were finance professors at Virginia Tech, and the others were either graduate students in Finance, or were MBA candidates.

We used a within-subjects design, where each of the eight subjects performed tasks in all eight conditions. Each subject viewed all eight stock options surface plots once in each of the eight possible experimental conditions listed in table 5.1. Table 5.2 describes and numbers the conditions. Table 5.3 lists the order of the subjects, conditions, and plots viewed.

Condition	**Description**
1	4 walls, HBR, Stereo
2	4 walls, HBR, No Stereo
3	4 walls, No HBR, Stereo
4	4 walls, No HBR, No Stereo
5	1 wall, HBR, Stereo
6	1 wall, HBR, No Stereo
7	1 wall, No HBR, Stereo
8	1 wall, No HBR, No Stereo

TABLE 5.2 – DESCRIPTION OF EXPERIMENTAL CONDITIONS

Subject #		Order of Conditions							
	Plot #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
1		6	1	3	5	2	4	7	8
2		2	4	7	8	5	6	3	1
3		8	2	1	6	7	3	4	5
4		4	6	5	3	8	7	1	2
5		3	5	6	4	1	8	2	7
6		5	7	4	2	6	1	8	3
7		1	3	8	7	4	2	5	6
8		7	8	2	1	3	5	6	4

TABLE 5.3 – SUBJECTS AND ORDER OF CONDITIONS

5.4 Results

This section is organized in a manner similar to section 4.4. The results pertaining to the first hypothesis regarding FOR are examined first. The second hypothesis with respect to head based rendering follows, and finally the third hypothesis regarding stereoscopic viewing is examined in the context of these results. Accuracy of the responses was not recorded due to the exploratory nature of this study.

5.4.1 FOR

To compare high FOR with low FOR, conditions 1 vs. 5, 2 vs. 6, 3 vs. 7, and 4 vs. 8 were compared for each of the task types performed. The completion time of the particular task type was averaged in the high FOR, and low FOR conditions. In three out of the five task types, the high level of FOR resulted in markedly lower task completion times. Those times are reported in Table 5.4

Task Type	HFOR	LFOR
Trend	20.8	44.4
Slope	65.1	88.9
Max/min	54.2	71.0
Range	103.1	93.4
Interpretive	48.9	41.5

TABLE 5.4 – COMPLETION TIMES FOR THE TASK TYPES IN THE HIGH FOR VS. LOW FOR CONDITIONS

We hypothesized that the higher level of FOR would help to create reduced task completion times with regard to trend determination. This could be explained by the user rapidly turning her head to gain an understanding of relationships in the data, rather than by navigating when one wall was used. The reduced times for the slope and max/min tasks may be explained similarly. While the range and interpretive question response times are higher in the high FOR condition, the times are relatively close to one another, and looks like the difference is not large enough to be meaningful. Therefore, it is difficult to state why these time differences occurred.

The disorientation ratings for the high FOR vs. the low FOR conditions were 1.4 and 1.7, respectively (scale 1-7, 7 equivalent to most disorienting). The condition usefulness ratings were 5.8 for the high FOR conditions, and 5.0 for the low FOR (scale 1-7, 7 equivalent to most useful). These ratings indicate a preference for a high FOR.

All users verbally expressed their preference for all screens being used in the CAVE. One subject continued to look for the extension of the surface plot on the side screens, even though they were not projected in that condition. Interestingly, although the subjects strongly expressed the view that four screens were helpful and that they would prefer to view the surface plots in the high physically immersive conditions, this was not shown by lower times across all task types.

5.4.2 Head-Based Rendering

To compare the conditions with head based rendering to those without, conditions 1 vs. 3, 2 vs. 4, 5 vs. 7, and 6 vs. 8 were examined. The completion time of the particular task type was averaged in the conditions where HBR was used, and also in the conditions where it was not used. Please refer to Table 5.5 for the task completion times.

Task Type	HBR	non-HBR
Trend	34.3	30.9
Slope	79.8	74.2
Max/min	64.4	60.8
Range	93.6	102.9
Interpretive	37.0	53.4

TABLE 5.5 – COMPLETION TIMES FOR THE TASK TYPES IN THE NBR VS. NON-HBR CONDITIONS

In three out of the five task types, the conditions with HBR were slower than the conditions without HBR. These times are all quite close to one another (with the exception of the interpretive question), so it is not clear even to state if there is a trend toward HBR, or non-HBR. In the case of the interpretive question type, we think that the HBR helped to create a stronger spatial understanding of the plot, which assisted the subject in responding to the task more quickly.

The disorientation ratings for the HBR vs. non-HBR conditions were 1.7 and 1.8, respectively (scale 1-7, 7 equivalent to most disorienting). The condition usefulness ratings were 5.1 for

HBR, and 5.1 without (scale 1-7, 7 equivalent to most useful). These ratings do not clearly indicate a preference for either condition.

When asked about head based rendering one subject responded, "I don't think I ever really used this." Another subject felt that with the surface plots it was not very helpful. However, it was noted that two subjects were moving their heads quite a bit to obtain different viewpoints when asked to find maximum points.

5.4.3 Stereoscopic Viewing

To examine any potential effects of stereoscopic viewing, conditions 1 vs. 2, 3 vs. 4, 5 vs. 6, and 7 vs. 8 were compared. The completion time of the particular task type was averaged in the conditions where stereo was used, and also in the conditions where it was not used. Please refer to Table 5.6 for the task completion times.

Task Type	Stereo	non-Stereo
Trend	30.6	34.6
Slope	95.9	58.1
Max/min	60.4	64.8
Range	95.0	101.5
Interpretive	53.3	37.1

TABLE 5.6 – COMPLETION TIMES FOR THE TASK TYPES IN THE STEREO VS. NON-STEREO CONDITIONS

In the slope and interpretive tasks, the times were substantially higher with stereo on. The other task times were close to one another. In the slope determination tasks, we hypothesized that the task time would be reduced with stereo. A possible explanation for this outcome is that the stereo with resultant depth cues, made judging the slopes of regions more difficult, especially if one were close to the plot.

The disorientation ratings for the stereo vs. non-stereo conditions were 1.6 and 1.9, respectively (scale 1-7, 7 equivalent to most disorienting). The condition usefulness ratings were 5.2 for stereo, and 5.3 without (scale 1-7, 7 equivalent to most useful). These ratings are too close to indicate a preference for either condition.

When queried about stereoscopic viewing, subjects were not able to recall the conditions where stereo was used and where it was not used. We thought that stereoscopic viewing would be less apparent in surface plots than viewing in scatter plots, but in this pilot study it was a bit of a surprise to discover that no subject could differentiate between stereo and non-stereo conditions.

5.4.4 Interaction Between the Components of Immersion

We also decided to look at the interaction between combinations of components of immersion. With our previously defined components, we have the following combinations: field of regard and head based rendering, field of regard and stereoscopic viewing, head based rendering and stereoscopic viewing, and field of regard, head-based rendering, and stereoscopic viewing taken together. Please see Figures 5.1-5.5 for the average task type completion time in each of the eight experimental conditions. Also see Appendix B.4 for the raw data.

For the trend completion task, it appears that the high FOR + Stereo condition yielded the fastest average time. This makes sense as the high FOR conditions were much faster in the trend task, and the added depth cue from stereo made it easier to determine relationships between the variables. In the slope comparison task, high FOR + HBR – Stereo was the quickest. In determining slopes, if one is very close to the plot, the stereo might cause some confusion due to the images that each eye sees not fusing properly.

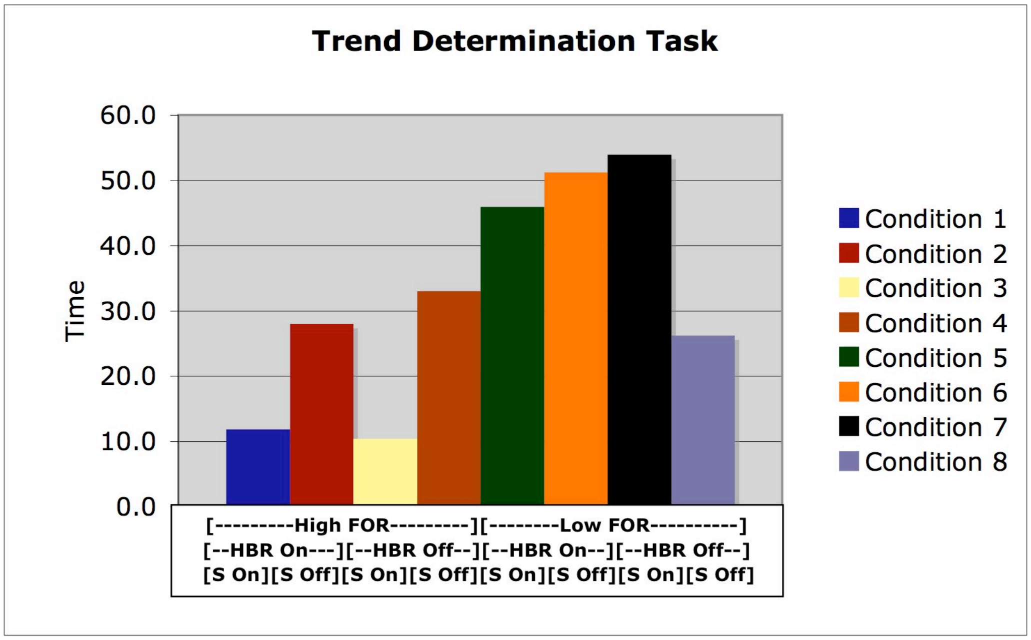

FIGURE 5.1 – TREND TASK COMPLETION TIMES IN EACH EXPERIMENTAL CONDITION

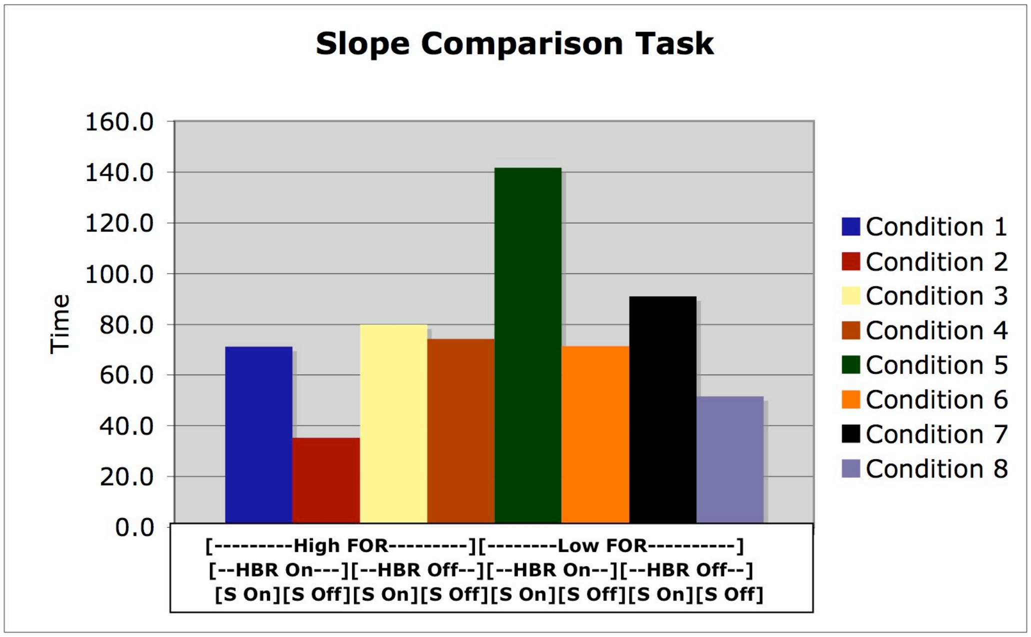

FIGURE 5.2 – SLOPE TASK COMPLETION TIMES IN EACH EXPERIMENTAL CONDITION

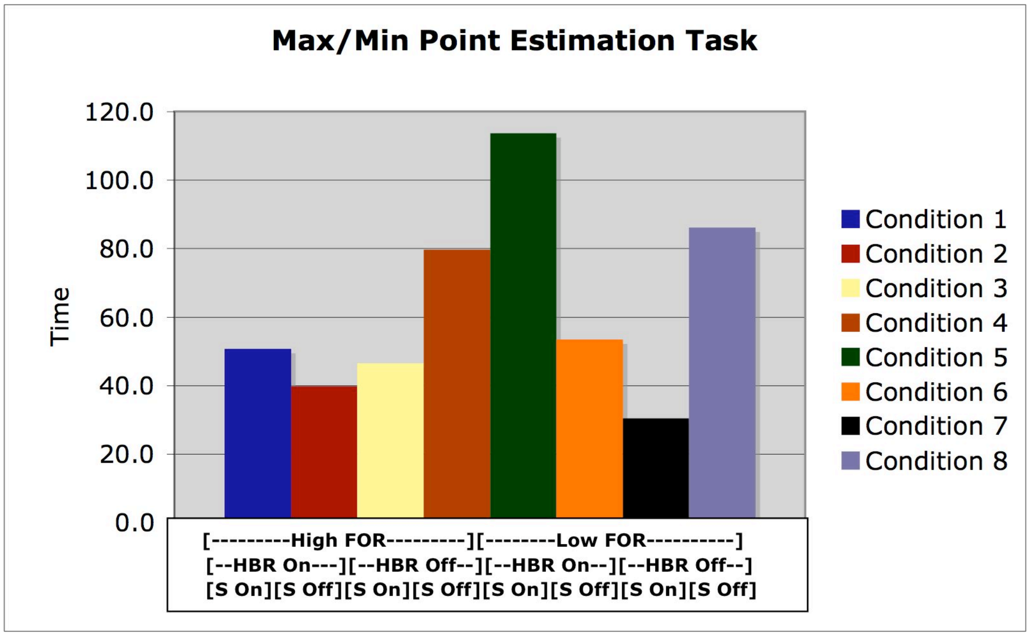

FIGURE 5.3 – MAX/MIN TASK COMPLETION TIMES IN EACH EXPERIMENTAL CONDITION

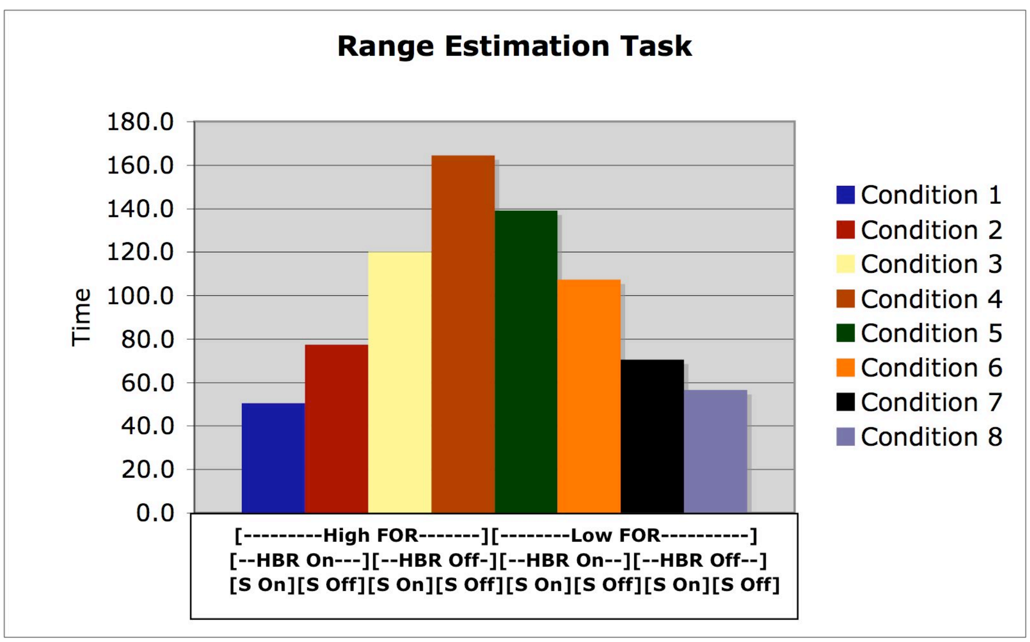

FIGURE 5.4 – RANGE TASK COMPLETION TIMES IN EACH EXPERIMENTAL CONDITION

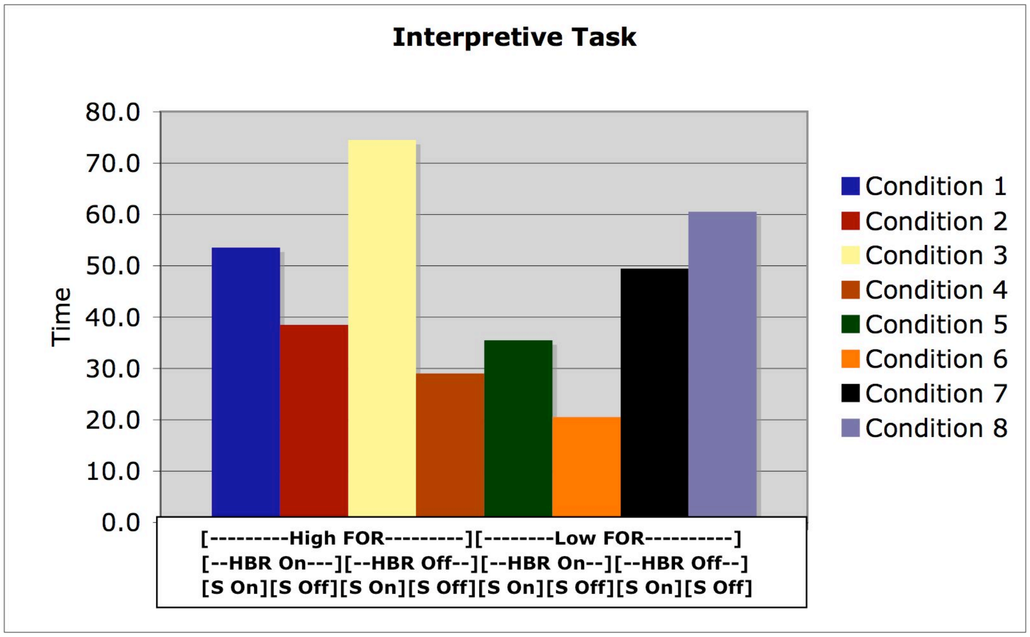

FIGURE 5.5 – INTERPRETIVE TASK COMPLETION TIMES IN EACH CONDITION

In the max/min point estimation task, low FOR + Stereo was the fastest. We think that the depth cues provided by stereo helped with this task, and the low FOR caused the subject to focus more intently on a small section of the graph, creating lower response times. In the range estimation task, the most immersive condition was quickest, which was high FOR + HBR + Stereo. For this task, viewing large sections of the graph were important, along with having strong depth cues. Hence, the most immersive condition was fastest. In the interpretive task, low FOR + HBR yielded the fastest average times. Here, viewing a smaller section of the graph from an intuitive perspective was thought to be most beneficial, and the data trend toward this contention.

5.5 Conclusions

Due to the experimental design and relatively small number of subjects recruited, detailed statistical analysis was not performed on the data collected in this pilot study. From the observed patterns in the data, it appears as though there are trends in favor of high FOR in the trend determination, slope comparison, and max/min point estimation tasks. The other two task types, range estimation and interpretive, were asked twice each. Therefore, since there were far fewer times these tasks were performed (only 4/20), differences in the times between conditions may not be as representative as the other task times.

One factor that may have had an effect on the outcome was the expected level of domain expertise and potential confusion with some of the questions. A few of the subjects were experts and could answer several questions without looking at the graphs, while others had difficulty interpreting the tasks to be performed. As recruiting subjects with this knowledge in the Blacksburg area is difficult, we felt our sample subject population though imperfect, to be well suited for this pilot study.

Some suggestions to improve understanding related to viewing the plots in CaveEquationView were implemented for the next experiment. As described in section 3.2.1, we implemented a height dependent shading model. We feel that his feature would make a substantial impact on the subject's ability to rapidly garner information. See Figure 5.1, and 5.2 for before and after comparisons with the height dependent feature. The second improvement modified the navigation and interaction of the dataset with the wand. In the old method, subjects had to fly through the dataset to obtain various viewpoints. With the new method, the user can still fly through the dataset, but now is able to rotate the dataset about the *z-axis* quickly and efficiently, merely by pressing buttons to the left or right of the joystick on the wand.

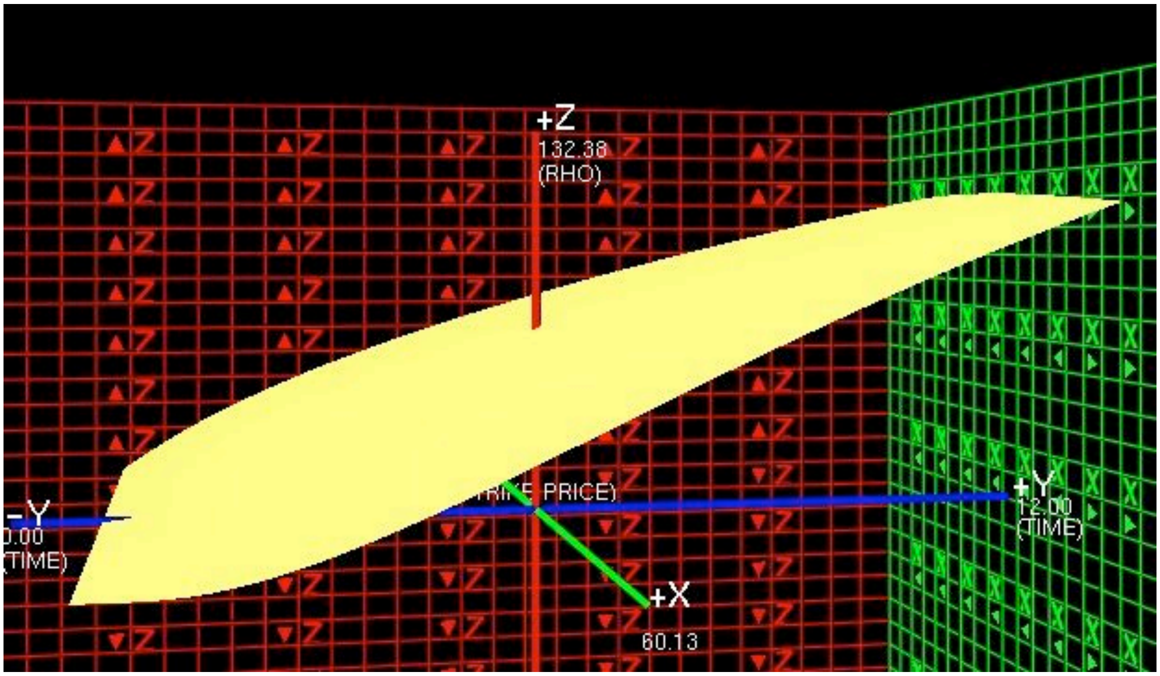

FIGURE 5.6 -- A SURFACE PLOT OF OPTIONS COMPONENTS USED IN THE PILOT STUDY

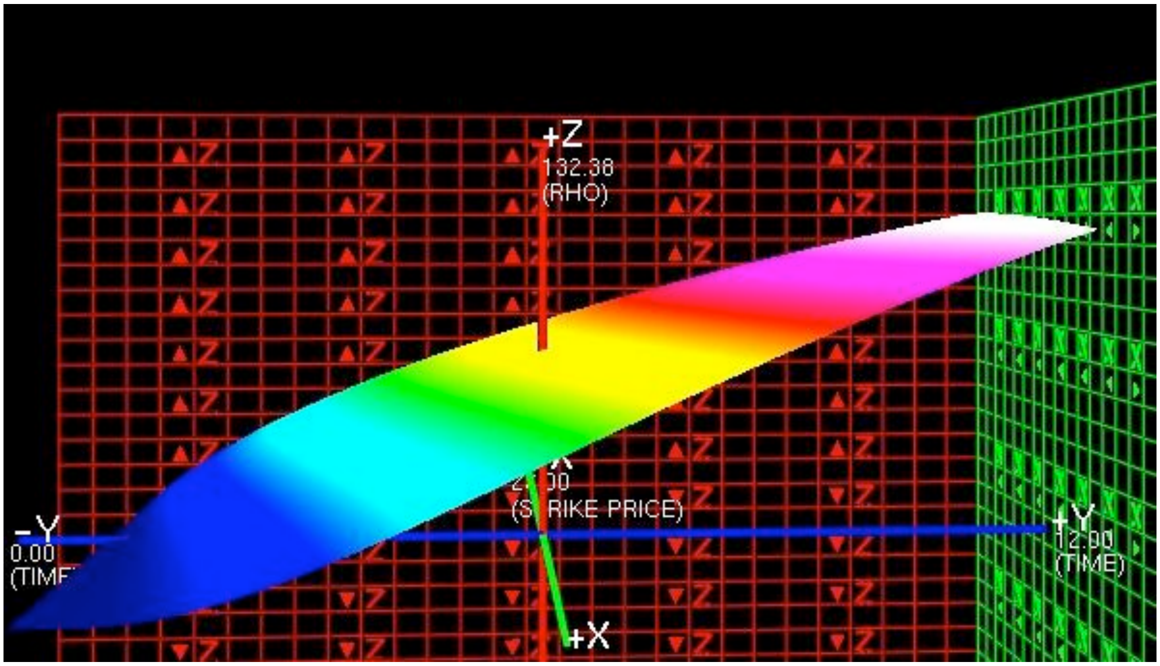

FIGURE 5.7 – THE SAME PLOT AS FIGURE 5.6. NOTE: THE PLOT'S CURVATURE IS EASILY SEEN HERE WITH CONTOUR SHADING, BUT NOT SEEN READILY IN FIGURE 5.6

Chapter 6: Combined Experiment with Scatter Plots and Surface Plots

6.1 Experimental Design

Similar to the previous study, we examined three components of immersion. They are FOR, head-based rendering, and stereoscopic viewing. We randomly divided the subjects such that each performed three tasks in the scatter plots, and three tasks in the

High FOR				Low FOR		
	Stereo	**No Stereo**			**Stereo**	**No Stereo**
Head Based Rendering	4 screen CAVE, HBR, Stereo	4 screen CAVE, HBR, No Stereo		**Head Based Rendering**	1 screen CAVE, HBR, Stereo	1 screen CAVE, HBR, No Stereo
No Head Based Rendering	4 screen CAVE, No HBR, Stereo	4 screen CAVE, No HBR, No Stereo		**No Head Based Rendering**	1 screen CAVE, No HBR, Stereo	1 screen CAVE, No HBR, No Stereo

TABLE 6.1 – EXPERIMENTAL CONDITIONS FOR COMBINED EXPERIMENT

surface plots. Each subject performed the tasks in a total of four (out of a possible eight) conditions. Hence, the components of immersion examined are between-subjects independent variables. It was thought that this design, by allowing the subjects to perform tasks in varying levels of FOR, would provide the most accurate results since FOR was considered to be quite important for several tasks in both types of plots. Please review Tables 6.1 through 6.6 for a description of the conditions, and the subject ordering for the scatter and surface plots.

We determined that in order to fully test our hypotheses with maximum effectiveness, the experiment should include two scatter plots and two surface plots where all of the tasks were completed in each of the experimental conditions eight times for each plot type. The time taken to complete the tasks was recorded and subjects were also asked to provide ratings on a seven-point scale based on their perception of disorientation and ease of completing the tasks in the specified conditions.

Based on the time that the surface plot pilot study subjects used to complete the tasks, we decided to shorten the experiment to reduce fatigue and potential errors. Our design combined and refined elements from both previous studies, but necessitated the recruitment of far more subjects to achieve statistical power.

6.2 Procedure

Each subject was trained with a sample scatter plot, and a surface plot similar to the ones used in the experiment. The subject was given ample time to become familiar with the look and feel of the plots, the navigation methods, and the rotation of the datasets. The navigation methods were first demonstrated and explained, with guidance provided to the subject until a comfortable level of proficiency was reached. The subjects were also instructed in how to use HBR by turning one's head, crouching, moving about, etc.

The subjects started in the center of the CAVE in both plot types. The scatter plot was shown at its' origin where very little of the dataset was viewable initially. The surface plot was shown initially such that the entire plot was viewable on the front CAVE wall.

Each task was stated verbally, and the time recorded commenced as soon as the task was finished being stated. The time measurement ended when the subject reported an answer, or stated that they were not sure of the correct one. In this study, only the side walls in the CAVE were used since the floor projector was malfunctioning (three walls used instead of the typical four). The implications of the floor projector not being used will be explained in the results discussion section.

6.3 Task Description

The tasks we chose in this experiment were the ones from the previous studies we felt were most effective in eliciting a sense of spatial awareness and understanding of the dataset. The tasks performed in this experiment can be placed in the following general categories (definitions are reprinted here from the previous sections for your convenience):

Single Point Search

The Single Point Search task had subjects locate a differently colored point in a densely packed group of points in the scatter plots. 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).

Trend Determination

As described similarly in the previous pilot studies, this task requires the subject to obtain a general sense of orientation and understanding of the dataset in order to identify trends within the plot. An example of a question in this category could be "As X increases, what happens to Z?" This type of task is fundamental to any sort of analysis with 3D graphs in information visualization and was performed in both types of plots.

Maximum/Minimum Point Estimation

Subjects were asked to find the maximum or minimum "height" of the surface plot, or the data point in the scatter plot with the max "height." In the surface plot, the corresponding coordinates of that max/min point were a requirement for task completion. We felt this was an important and relevant task as one might often want to know exactly where these points are in order to gain a greater insight into the dataset, especially with financial applications.

Slope Comparison

This kind of task involved the comparison of the slope of the surface plot at one specified point, to the slope of the plot at another point. We feel this task is important and has significant real-world relevance. It requires the subject to rapidly ascertain an understanding of the dataset and presents an ideal manner in which to compare the effects of the components of immersion.

Please see Appendix C.3 for a complete list of questions and tasks.

6.4 Subjects

Since the length of the experiment was shortened (reducing the number of conditions that may be tested at one time), and the need for statistical verification of trends was of paramount importance, we recruited far more subjects than in the previous studies. Thirty-two individuals participated in the experiment ranging in age from 20 to 59, with an average of 27. Eight were female and all were familiar with programs like Matlab, Mathematica, or the simple graphing functions in Excel. They were not recruited on the basis of domain expertise or experience with VEs, but were required to possess a basic understanding of graphs and plots. The requirements for the subject pool were deemed reasonable as the users of such scatter and surface plots are those that do not necessarily have strictly technical backgrounds. Although the majority had computer science/engineering backgrounds (69%), the remainder came from a broad range of backgrounds like those from the social sciences who have manipulated large amounts of data in various ways.

High FOR				Low FOR		
	Stereo	No Stereo			Stereo	No Stereo
Head Based Rendering	1, 9, 17, 25	2, 10, 18, 26		Head Based Rendering	5, 13, 21, 29	6, 14, 22, 30
No Head Based Rendering	3, 11, 19, 27	4, 12, 20, 28		No Head Based Rendering	7, 15, 23, 31	8, 16, 24, 32

TABLE 6.2 – SUBJECT NUMBER WHO PERFORMED TASKS IN EACH CONDITION FOR SCATTER PLOT 1

High FOR				Low FOR		
	Stereo	No Stereo			Stereo	No Stereo
Head Based Rendering	2, 10, 18, 26	1, 9, 17, 25		Head Based Rendering	6, 14, 22, 30	5, 13, 21, 29
No Head Based Rendering	4, 12, 20, 28	3, 11, 19, 27		No Head Based Rendering	8, 16, 24, 32	7, 15, 23, 31

TABLE 6.3 – SUBJECT NUMBER WHO PERFORMED TASKS IN EACH CONDITION FOR SCATTER PLOT 2

High FOR				Low FOR		
	Stereo	No Stereo			Stereo	No Stereo
Head Based Rendering	5, 13, 21, 29	6, 14, 22, 30		Head Based Rendering	1, 9, 17, 25	2, 10, 18, 26
No Head Based Rendering	7, 15, 23, 31	8, 16, 24, 32		No Head Based Rendering	3, 11, 19, 27	4, 12, 20, 28

TABLE 6.4 – SUBJECT NUMBER WHO PERFORMED TASKS IN EACH CONDITION FOR SURFACE PLOT 1

High FOR				Low FOR		
	Stereo	No Stereo			Stereo	No Stereo
Head Based Rendering	8, 16, 24, 32	7, 15, 23, 31		Head Based Rendering	4, 12, 20, 28	3, 11, 19, 27
No Head Based Rendering	6, 14, 22, 30	5, 13, 21, 29		No Head Based Rendering	2, 10, 18, 26	1, 9, 17, 25

TABLE 6.5 – SUBJECT NUMBER WHO PERFORMED TASKS IN EACH CONDITION FOR SURFACE PLOT 2

High FOR				Low FOR		
	Stereo	No Stereo			Stereo	No Stereo
Head Based Rendering	1	2		Head Based Rendering	5	6
No Head Based Rendering	3	4		No Head Based Rendering	7	8

TABLE 6.6 – CONDITION NUMBER WITH COMPONENTS OF IMMERSION USED

6.5 Results

In this section we examine the results from the combined experiment in a manner similar to the results sections in the previous two experiments. The components of immersion that were examined are discussed in their respective sections with regard to the tasks performed in the scatter and surface plots. A general linear model was constructed to test for statistical significance of any effects.

6.5.1 FOR

Since the three tasks the users performed were distinct, the effect on the individual tasks is examined with respect to the plot types.

Scatter Plots

The average response time for the single point search task in the high FOR conditions v. the low FOR conditions was 34.8 v. 25.3sec with $F_{1,30} = 1.949$; $p = .168$. For the trend determination task these times were 38.4 v. 50.4sec, with $F_{1,30} = 3.84$; $p = .055$. The min/max estimation task times were 16.5 v. 19.8sec with $F_{1,30} = .859$; $p = .358$. Accuracy for these tasks was 100%. There was a tendency toward a significant effect on the trend determination task.

The average disorientation for the high FOR conditions was 1.8 v. 1.9 (scale of 1-7, 7 being most disorienting) for the low FOR ones. The average ease of task completion rating was 5.7 (scale of 1-7, 7 being "easiest") for both conditions. On the trend determination task, the high FOR appears to have reduced the task completion time.

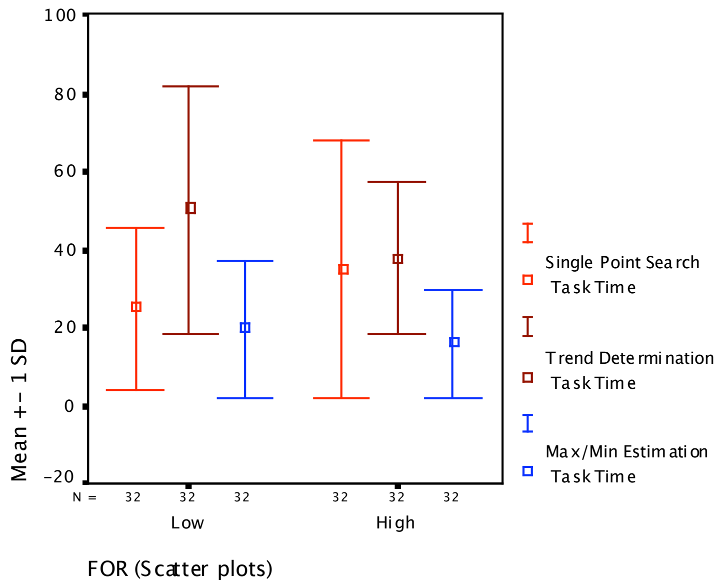

FIGURE 6.1 – AVERAGE TASK COMPLETION TIMES IN THE HIGH FOR V. THE LOW FOR CONDITIONS IN THE SCATTER PLOTS WITH STANDARD DEVIATION SHOWN FOR EACH TASK

Surface Plots

The average max/min estimation task times for the surface plots in the high FOR vs. low were 97.1 vs. 65.4sec, with $F_{1,30} = 3.99$; p = .051. For the trend determination task these times were 18 vs. 17.5sec with $F_{1,30} = .017$; p = .898, and for the slope comparison task, 35.3 vs. 36.9, with $F_{1,30} = .044$; p = .835. The max/min task accuracy for the high and low FOR conditions are respectively, 83.9 and 91.8% with an overall mean of 87.8%. There was a significant effect on the max/min estimation task time with $F_{1,30} = 3.99$; p = .051. However, there was a subject whose time was nearly five standard deviations from the mean. If we exclude that subject, the effect is no longer significant.

The average disorientation ratings were 1.7 vs. 2.2 (scale 1-7, 7 most disorienting) and the ease of task completion ratings were 4.1 vs. 4.9 (scale 1-7, 7 being "easiest"). Subjects felt they were less disoriented in the high FOR condition, yet thought the same condition made the task more difficult to complete.

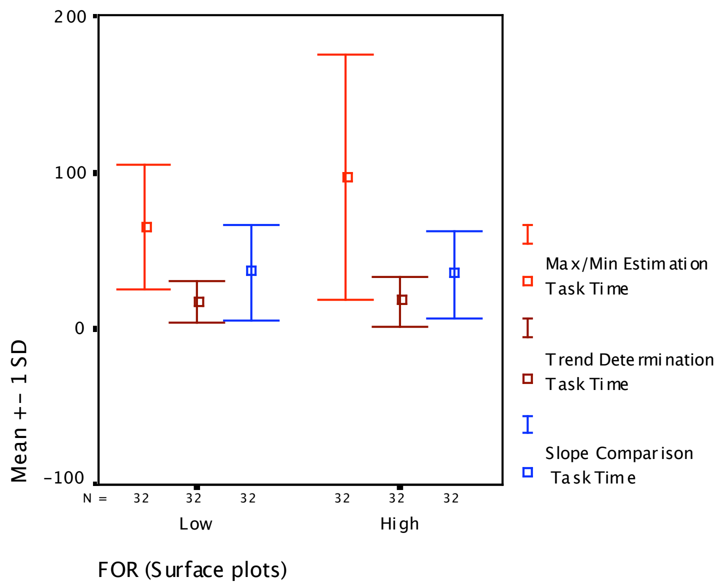

FIGURE 6.2 – AVERAGE TASK COMPLETION TIMES IN THE HIGH FOR V. THE LOW FOR CONDITIONS IN THE SURFACE PLOTS WITH STANDARD DEVIATION SHOWN FOR EACH TASK

6.5.2 Head-Based Rendering

Scatter Plots

The average single point search task completion times were not significant for the HBR v. the non-HBR conditions, and were 32.4 vs. 27.6sec with $F_{1,30} = .542$; $p = .465$. For the trend task they were 40.4 vs. 48.4sec with $F_{1,30} = 1.833$; $p = .181$, and for the max/min task they were 17.2 vs. 19.1sec with $F_{1,30} = .360$; $p = .551$. Accuracy was 100% for these tasks.

The average disorientation ratings were 1.8 vs. 1.9 (1-7, 7 most disorienting) and the average ease of task completion ratings was 5.7 (1-7, 7 being "easiest") in both cases.

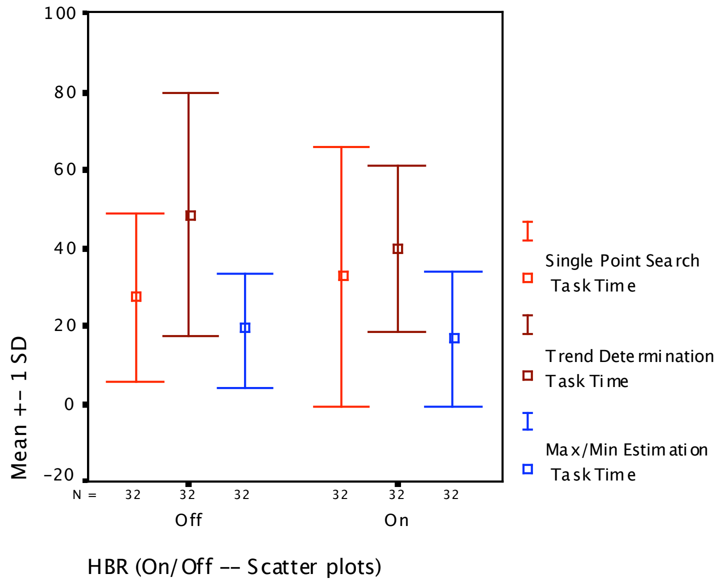

FIGURE 6.3 – AVERAGE TASK COMPLETION TIMES IN THE HBR V. NON-HBR CONDITIONS IN THE SCATTER PLOTS WITH STANDARD DEVIATION SHOWN FOR EACH TASK

Surface Plots

The max/min estimation task times for HBR vs. non HBR were 77.2 v. 85.3sec with $F_{1,30}$ = .265; p = .609. The trend determination task times were 16.9sec vs. 18.7 with $F_{1,30}$ = .241; p = .625, and the slope comparison task times were 30.6 vs. 41.6sec with $F_{1,30}$ = 2.082; p = .155. Accuracy was 91.9 v. 83.8 %. Though all task times were faster and accuracy was improved in the HBR conditions, none were statistically significant.

The average disorientation ratings were 1.9 (scale 1-7, 7 most disorienting) in both HBR and non-HBR conditions. Ease of task completion ratings were 4.7 vs. 4.3 (1-7, 7 being “easiest”).

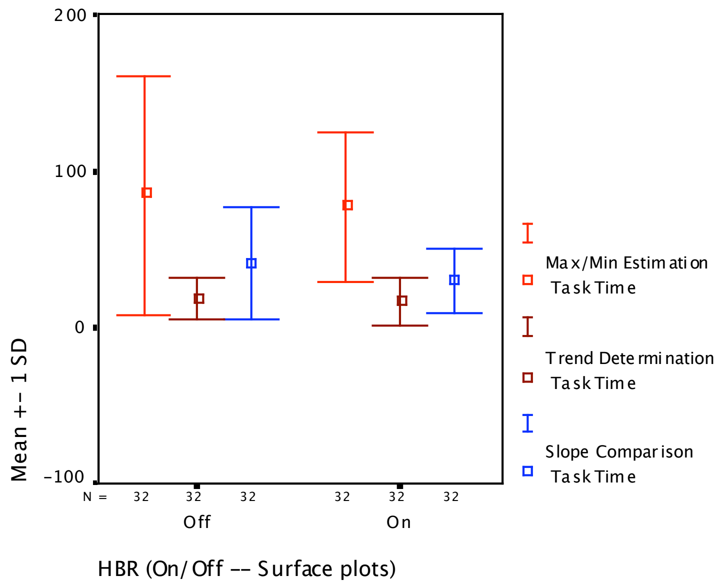

FIGURE 6.4 – AVERAGE TASK COMPLETION TIMES IN THE HBR V. NON-HBR CONDITIONS IN THE SURFACE PLOTS WITH STANDARD DEVIATION SHOWN FOR EACH TASK

6.5.3 Stereoscopic Viewing

Scatter Plots

In the stereo vs. non stereo conditions the average task completion times were 25.6 vs. 34.3 with $F_{1,30} = 1.412$; $p = .240$ in the single point search task. In the trend determination task they were 42.4 vs. 46.4 with $F_{1,30} = .561$; $p = .457$, and in the max/min estimation task they were 16.7 vs. 19.6 with $F_{1,30} = .694$; $p = .408$. None of these time differences were statistically significant. Accuracy was 100% for these tasks.

The average disorientation ratings in the stereo vs. non stereo conditions were: 1.65 v. 2.1 (scale 1-7, 7 most disorienting), and the ease of task completion ratings were: 6.0 vs 5.4 (scale 1-7, 7 being "easiest"). By all these metrics, stereoscopic viewing has a clear trend toward providing benefit, but none of these data are statistically significant.

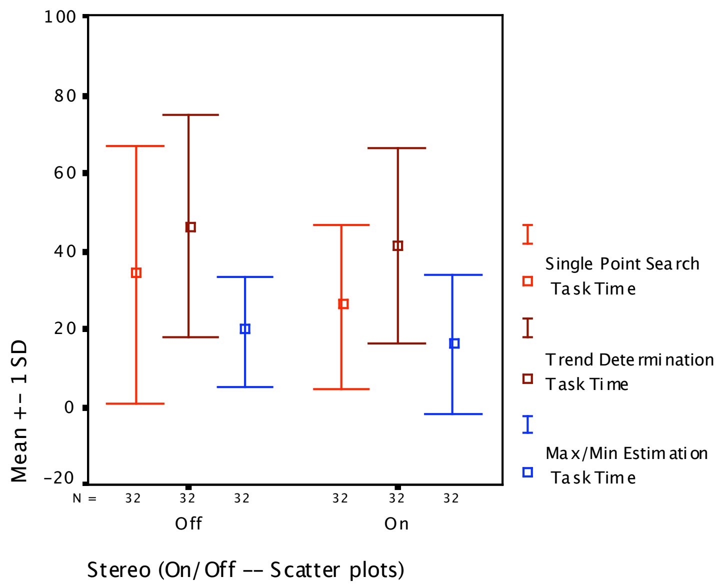

FIGURE 6.5 – AVERAGE TASK COMPLETION TIMES IN THE STEREO V. NON-STEREO CONDITIONS IN THE SCATTER PLOTS WITH STANDARD DEVIATION SHOWN FOR EACH TASK

Surface Plots

The average task completion times in the stereo v. non-stereo conditions were 82.9 vs. 79.5 sec for max/min estimation with $F_{1,30} = .046$; $p = .831$. For the trend determination task these were 19.4 vs. 16.1 with $F_{1,30} = .819$; $p = .369$, and in the slope comparison task these times were 35.1 vs. 37.1 with $F_{1,30} = .067$; $p = .797$. Min/max task accuracy was 87.7% vs. 88% in the stereo v. non-stereo conditions.

The average disorientation ratings were: 1.8 vs. 2.1, and the ease of task completion ratings were: 4.7 vs. 4.4. No clear trend was discernable for the stereo v. non-stereo conditions in the surface plots.

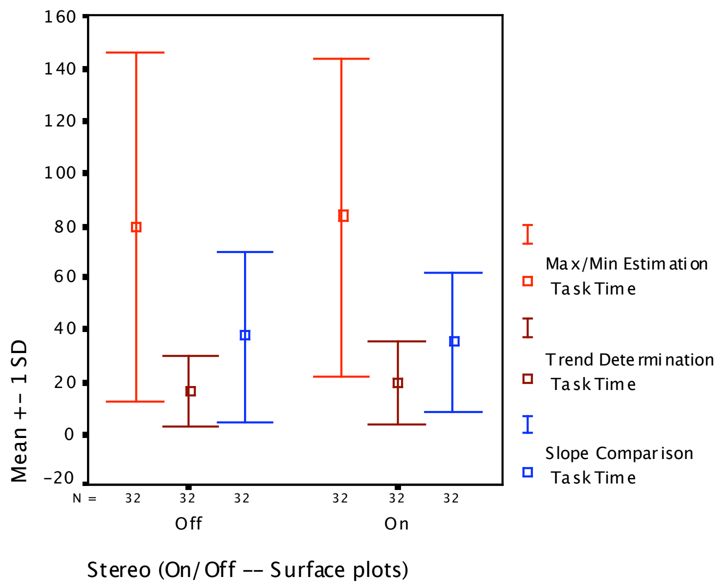

FIGURE 6.6 – AVERAGE TASK COMPLETION TIMES IN THE STEREO V. NON-STEREO CONDITIONS IN THE SURFACE PLOTS WITH STANDARD DEVIATION SHOWN FOR EACH TASK

6.5.4 Interaction Between the Components of Immersion

We examined the interaction between combinations of components of immersion. With our previously defined components, we have the following combinations: field of regard and head based rendering, field of regard and stereoscopic viewing, head based rendering and stereoscopic viewing, and field of regard, head-based rendering, and stereoscopic viewing taken together. After analyzing the data, we found an interesting and statistically significant interaction effect $F_{1,30} = 4.03$; $p = .05$ in the trend determination task of the scatter plots between field of regard and head-based rendering.

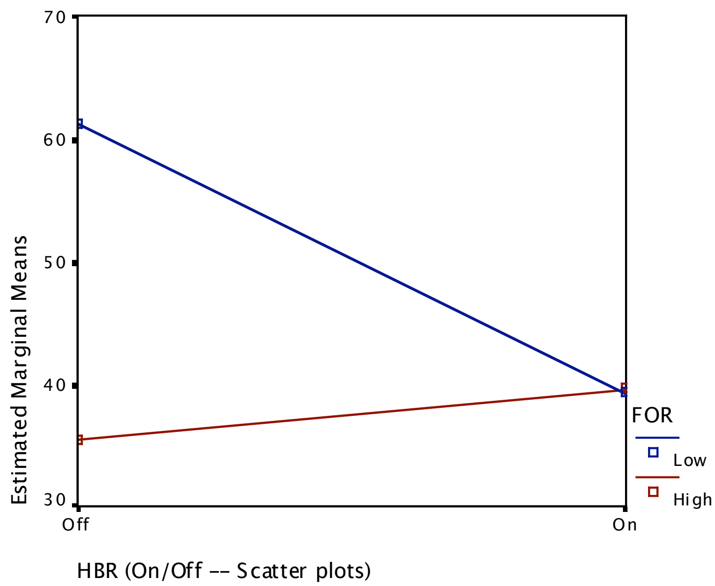

FIGURE 6.7 – INTERACTION EFFECT ON TREND DETERMINATION TASK TIMES WITH FOR AND HBR IN SCATTER PLOTS

In the high FOR condition when HBR is off, this combination yields the lowest task response time and is nearly equivalent to HBR being on with either the high or low FOR condition (Figure 6.1). However, when the FOR is low and HBR is off, this results in a dramatically longer response time than the other possible combinations.

In the surface plots, there was a similar, but more pronounced interaction (Figure 6.8) in the trend determination task between FOR and stereoscopic viewing, but this was not significant. It could be said to be close, however, with $F_{1,30} = 3.69$; $p = .06$

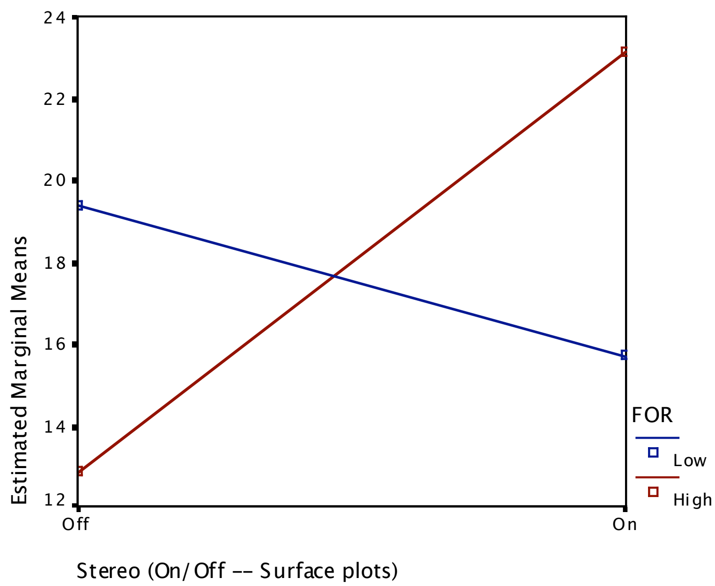

FIGURE 6.8 – INTERACTION EFFECT ON TREND DETERMINATION TASK TIMES WITH FOR AND STEREO IN THE SURFACE PLOTS

6.6 Discussion

In two of the three tasks in the scatter plots, the high FOR condition has yielded lower task response times, and one of the tasks was close to being significant. This task involved determining the trend exhibited in the scatter plots. This makes logical sense and agrees with our hypothesis that the high FOR condition would be beneficial for users to determine trends in the data. When observing the subjects, it was quite clear that many turned to the side walls in the CAVE to quickly glean information. When the side walls were turned off in the low FOR condition and the subjects had previously experienced the high FOR condition while practicing, subjects were observed turning to the side walls even when they were turned off, and expressed frustration that the screens were blank.

The surface plots exhibited a significant negative effect on completion times in the high FOR condition. However, as noted in the results section if an individual who had a time nearly five times the standard deviation away from the mean was excluded from the study, there would be no significant FOR effect. Given that the subject deviates so greatly from the mean, it is a reasonable step to exclude those times and conclude that FOR did not have a significant effect on task completion times in the surface plots. The disorientation and ease of task completion ratings meanwhile do exhibit a trend toward the high FOR condition. Although this individual's exclusion from the data changed the significance of the negative influence of the high FOR condition, the accuracy rating implied that the low FOR condition yielded more better responses. One reason for this could be that in the low FOR conditions, this forced the subjects to focus on a smaller part of the surface plot, which made estimation more accurate.

HBR yielded faster times in two out of three tasks in the scatter plots though these differences were not statistically significant. The first task, to find the pink colored cube in a densely packed dataset, was the task that produced a slower time in the HBR condition. One reason for this might be that as the user zeroed in on the colored cube, small head movements could occlude the view of the target cube as the user was navigating through and trying to aim the virtual ray at the cube in order to read back the ID number.

In the surface plots, all times were faster in the HBR conditions than without. Although not statistically significant, this is interesting because it was noted that many users in the surface plots appeared to keep their heads still. Many subjects in the scatter plots however, were extensively moving their heads to gain a better perspective and escape occluding data points. Accuracy was better in the HBR condition; perhaps by permitting a more natural perspective, the user could more intuitively find a viewpoint that provided a more accurate response. In the surface plots, the disorientation and ease of task completion times favored the HBR condition.

Stereoscopic viewing produced faster task completion times in all three tasks, than in the conditions without stereo in the scatter plots. We believed that stereo exhibits strong depth cues, which can be quite important in judging distances, trends, and selecting certain data points. The disorientation and ease of task ratings also seem to indicate a trend for stereo in the scatter plots, but no data here was found to be statistically significant.

In the surface plots, stereoscopic viewing yielded slightly longer task times than in the conditions without. Accuracy was similar in both stereo and non-stereo conditions. While these data are not significant, it could be possible that the stereo condition introduced some effect that caused increased difficulty judging distances and determining slopes in the surface plots.

There is a very interesting interaction effect between FOR and HBR. When HBR is on, times do not change much when the FOR level changes. However, when HBR is turned off and the high FOR condition is present, task times decrease drastically. The task affected is the trend determination task in the scatter plots. The users tended to back up a bit from the dataset, while using the buttons on the wand to rotate the dataset and turning their heads simultaneously to obtain a better idea of trends. This could explain why there is not much difference in times when HBR is on. When HBR is off in the low wall condition, the users must navigate more to obtain a better understanding of the dataset, hence increasing the task completion times. When the HBR is turned off in the high FOR condition, the user could back up a bit, use the buttons for rotating the dataset, and obtain a good understanding of the trends without much navigation. This would explain the decrease in times when HBR is off, when increasing the level of FOR.

There is another interesting interaction effect between FOR and stereoscopic viewing in the surface plots for the trend determination task. As FOR is increased, with stereo, the average task time increases. When stereo is not used and the FOR level is low, the time decreases markedly. When stereo is not used, the image of the surface plot appears flatter, and it is suspected that this forces the subject to obtain trend cues from the height dependent contour shading on the surface plots rather than a purely visual estimation. When a high FOR, the user could easily obtain views of the plot by looking at the left and right screens, which without stereo made determining trends less time consuming.

The trends found in our scatter plot pilot study favoring high FOR, HBR, and Stereo were largely confirmed in the formal combined study. In the trend determination task performed in the combined study, this time difference tended toward significance favoring the high FOR condition. HBR yielded faster times in two out of the three tasks, and all tasks were faster in the stereo condition in the formal study with scatter plots. This was the general trend in the scatter plot pilot study.

For the surface plot pilot study, we saw a trend for faster times in the high FOR condition. This trend was not found in the formal experiment with the surface plots as times between high and low FOR conditions were close when the outlier was excluded in the formal experiment. For HBR in both studies concerning the surface plots, no trend was detectable in the pilot study, but there was a trend in favor of HBR in the formal study. For stereo in the pilot study, no trend was found, but in the formal study all tasks were faster though not statistically significant.

In the formal study, one of the projectors (floor) in the CAVE was not working. While this did not affect the FOR, we feel that if it had been functioning properly, the results may have been stronger in the fully immersive condition. Several users commented that it would have been nice to have the floor in order to obtain a top down view of the dataset, particularly when performing the min/max task in the surface plots. This would have made selection of the corresponding points easier, as both X and Y-axes could be viewed easily at the same time. In the trend determination task in both plots, quickly looking at the plot from the top down rapidly provides the subject with information that had to be obtained by rotating the datasets when only the three walls were used. We strongly suspect the floor projector would enhance the spatial understanding of both plot types more rapidly than without. The missing floor also discourages people from navigating into the datasets, which is where the beneficial influences of HBR and stereo would be more pronounced. We look forward to others continuing similar work in the CAVE, and anxiously await the repair of our system.

Chapter 7: Conclusions

This body of research is aimed at exploring if there are any tangible benefits of immersion in 3D information visualization applications using the CAVE. We posited several hypotheses regarding the components of immersion, and developed a method in which to quantify those effects. We also developed two info viz applications, and conducted three separate studies in order to answer some research questions posed and to test our hypotheses.

To summarize our findings in each of the experiments and relate them to the research questions posed and hypotheses stated, a brief recap of the studies, procedures, and outcomes follow.

In the first exploratory pilot study, we examined the effects of FOR and HBR. Using the CAVE, and our scatter plot CaveDataView application we asked four subjects to perform several tasks in each of the four conditions. Although a statistical analysis was not performed due to the nature of the study and the few number of subjects used, we found that there were trends favoring high FOR and HBR.

In the pilot study using our surface plot application CaveEquationView, we asked eight knowledgeable subjects to perform tasks in eight different plots, in eight of the possible conditions with FOR, HBR, and stereoscopic viewing. Here no statistical analysis was performed on the data as there were not enough data points to achieve a meaningful analysis. There was a slight trend in favor of high FOR, but this cannot be stated with any degree of certainty. None of the other components or user ratings metrics proved definitive. We also examined the interaction between the components, but no promising trends were discovered.

In the combined experiment, we had 32 subjects perform tasks in both the scatter and surface plots with the effects of FOR, HBR, and stereoscopic viewing explored and analyzed. In the scatter plots, the high level FOR was found to have a nearly statistically beneficial effect on one of the tasks. HBR showed some nice trends in favor of its use, but the data were not statistically significant. Stereo also showed some nice trends in the scatter plots, but had a slight negative effect in the surface plots. In the surface plots, stereo was found to have nearly the opposite effect. The data concerning stereo also were not statistically significant. We also examined the interaction effects and found that FOR and HBR had a statistically significant effect in the scatter plots. Specifically, our hypotheses were:

1. *A high degree of FOR will allow higher levels of task performance and greater user satisfaction/ease of task completion when visualizing datasets represented by 3D scatter plots and 3D surface plots.*

While the scatter plot pilot study results showed some promise, the pilot study was inconclusive. The combined study however showed that there were some benefits to the high FOR condition in scatter plots where trend determination is concerned. Often times the users remarked how nice the high FOR condition was in both types of plots, but the data show greater significance to the beneficial effect in the scatter plots.

2. *Head based rendering will allow higher levels of task performance and greater user satisfaction/ease of task completion when visualizing datasets represented by 3D scatter plots and 3D surface plots. This effect will be less pronounced on surface plots, than in scatter plots.*

The data in all studies we conducted seem to indicate positive trends for the use of HBR, but the data do not show a statistically significant trend in the use of HBR. However, in the formal study HBR on was very beneficial in the low FOR condition for the trend determination task.

3. *Stereoscopic viewing will allow higher levels of task performance and greater user satisfaction/ease of task completion when visualizing datasets represented by 3D scatter plots and 3D surface plots. However, stereoscopic viewing will have a greater beneficial effect on 3D scatter plots than it will have on 3D surface plots.*

In the first exploratory study, stereoscopic viewing was not examined. In the pilot study, no trend was found to indicate a preference for stereo. In the combined experiment, stereo was found to have beneficial trends in the scatter plots. However, in the surface plots, the trend was slightly negative. In this sense, our hypothesis was correct in that it provided a benefit in the scatter plot, and this was greater than the beneficial effect found in the surface plot. Stereo on was beneficial, however, in the low FOR condition in the trend determination task.

We had thought, based on the previous studies, that the level of FOR would be the most important element in the components of immersion for the given metrics we used. It was important, as we had expected with trend determination in the scatter plots, but not as important as we had expected in the surface plots.

While we found some advantages to analyzing datasets in immersive conditions, the results do not strongly suggest that immersion would be widely beneficial in the information visualization domain. We also think our results would have been more definitive if the CAVE were at full functionality, but do not believe that our studies show an overarching need for immersive systems in data analysis.

We have produced empirical evidence to show that immersive qualities are beneficial in some circumstances. We have also demonstrated a method using the CAVE to separate and examine components of immersion. The preceding statements were prime motivating factors in conducting these studies, and though the results are not as convincing as we might have hypothesized, are nonetheless a part of the tangible, successful outcomes of this research.

7.1 Contributions

We feel that we have made several contributions to this area of research. Namely, they are:

1. A detailed study of the primary components of immersion's effect on 3D information visualization

Not much research has been conducted in this area to explain potential benefits of immersion. While our studies concentrate on 3D information visualization, the benefit of such immersive systems remains an extremely interesting and important question that is left essentially unanswered. Our research is a step in the direction of addressing the issue as to whether these expensive and complicated immersive systems are providing any real benefit beyond the "wow" factor.

2. Two information visualization applications for use in the CAVE

We developed CaveDataView and CaveEquationView. They both provide a unique way of looking at data, since no applications like these have previously been developed specifically for use in the CAVE. Parts of CaveDataView have been used as a basis for another application studying immersive effects using a different display environment. CaveEquationView can display surface plots of both pre-computed data or mathematical equations on the fly. It may be used in the future as a platform for further immersion experiments and as a learning tool for students of various disciplines to study surface plots.

3. A method for quantifying the effects of immersion using the CAVE

As mentioned earlier in this paper, our novel method provides a simple and intuitive manner in which to measure the effects of various components of immersion. Components of immersion and other variables are not confounded with our design, which should lead to more accurate results than previously achieved with subsequent studies.

7.2 Future Work

Though we found a high degree of FOR to tend toward a significant beneficial effect on scatter plots, and HBR and stereoscopic viewing to exhibit some promising trends, there are many unanswered questions, which necessitate further study. The following are some suggestions that some in the future may wish to carry out to further study this fascinating topic.

Modification of experimental design

It would be interesting to see if different datasets in the scatter plots produce a similar outcome. Datasets that are more enveloping can and should be used for further study. Simple pre-tests of users ability to interpret basic graphical functions may help to reduce variance. Changing the FOR level to a between subjects independent variable, while HBR and stereoscopic viewing are within subjects might reduce variance as the subject may react adversely to a drastic change such as the level of FOR. User tasks that also make better use of the components of immersion should be used for additional studies. This would be crucial in determining effects, as our tasks perhaps did not thoroughly test the components as comprehensively as possible.

Modification of current info viz application

CaveEquationView might be modified to toggle through flat shading, contour shading, and a wire mesh view while studying the effects of immersion with more complicated mathematical functions.

Expand immersion testing domain

An educational application where some students see a plot in certain conditions (for example, students in a derivatives class) and are then tested on their knowledge retention might be another domain in which to conduct further study.

While our study showed promising trends for some components of immersion with certain tasks in particular types of plots, no unifying benefit of immersion may be surmised from our study. This is still an important research question that remains unanswered definitely. We hope to have laid the groundwork for further, probing, and meaningful research in this field.

References

[Ahlberg95] Ahlberg, C. and Wistrand, E. (1995) "An Information Visualization & Exploration Environment." Proceedings of IEEE Information Visualization, pp. 66-73.

[Anupam95] Anupam, V., Dar, S., Leibfried, T., and Petajan, E., (1995) "DataSpace: 3-D Visualizations of Large Databases." Proceedings of IEEE Information Visualization, pp. 82-89.

[Arns99] Arns, L., Cook, D., Cruz-Neira, C. (1999) "The Benefits of Statistical Visualization in an Immersive Environment." Proceedings of IEEE Virtual Reality, pp. 88-95.

[Arsenault04] Arsenault, R. and Ware, C. (2004) "The Importance of Stereo and Eye Coupled Perspective for Eye-Hand Coordination in Fish Tank VR."

[Belcher03] Belcher, D., Billinghurst, M., Hayes, S.E., and Stiles, R. (2003) "Using Augmented Reality for Visualizing Complex Graphs in Three Dimensions." Proceedings of ISMAR

[Black73] Black, F. and Scholes, M. (1973) "The Pricing of Options and Corporate Liabilities." The Journal of Political Economy, Vol. 81, No. 3, pp. 637-654.

[Bowman05] Bowman, D., Kruiff, E., LaViola, Jr., J.J., and Poupyrev, I. (2005) "3D User Interfaces" Addison-Wesley

[Bowman04] Bowman, D., and Raja, D. (2004) "A Method for Quantifying the Benefits of Immersion Using the CAVE." Presence-Connect, 5

[Boyd03] Boyd, D.R.S., Walton, J.P.R.B., and Gallop, J.R. (2003) "Putting You in the Picture: Enhancing Visualization with a Virtual Environment."

[Brooks99] Brooks, F. (1999) "What's Real About Virtual Reality" IEEE Computer Graphics and Applications, 19(6), 16-27

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

[Datey01] Datey, A. (2001) "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/

[Dede96] Dede, C., Salzman, M. C., Loftin, R. B. (1996) "ScienceSpace: Virtual Realities for Learning Complex and Abstract Scientific Concepts." Proceedings of IEEE VRAIS, pp. 246-252.

[Feiner90] Feiner, S., and Beshers, C., (1990) "Worlds within Worlds: Metaphors for Exploring *n*-Dimensional Virtual Worlds." Proceedings UIST, Snowbird, UT, October 3-5, pp. 76-83.

[Gruchalla04] Gruchalla, K. (2004) "Immersive Well-Path Editing: Investigating the Added Value of Immersion." Proceedings of IEEE Virtual Reality, pp. 157-164.

[Hodges96] Hodges, L. F., Rothbaum, B., Watson, B., Kessler, G. Drew., Opdyke, D. (1996) "A Virtual Airplane for Fear of Flying Therapy." Proceedings of IEEE VRAIS,

[Hollands95] Hollands, J.G., Pierce, B., Magee, L., (1995) "Displaying Quantitative Information in Two and Three Dimensions." Proceedings of Human Factors and Ergonomics Society 39th Annual Meeting, pp. 1425-1429.

[Hull03] Hull, J.C., (2003) "Options, Futures, and Other Derivatives." 5th Edition, Prentice-Hall

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

[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 IEEE Virtual Reality, pp. 164-168.

[Mine95] Mine, M., (1995) "Virtual Environment Interaction Techniques." TR95-018, UNC Chapel Hill, Department of Computer Science

[Natenberg88] Natenberg, S., (1988) "Option Volatility and Pricing Strategies." Probus Publishing

[Palais99] Palais, R.S., (1999) "The Visualization of Mathematics: Towards a Mathematical Exploratorium." Notices of the AMS, June/July, pp. 647-649

[Parker99] Parker, G., Franck, G., Ware, C., (1999) "Visualization of large Nested Graphs in 3D: Navigation and Interaction

[Pausch97] Pausch, R., Proffitt, D., and Williams, G. (1997) "Quantifying Immersion in Virtual Reality." Proceedings of ACM SIGGRAPH.

[Robertson93] Robertson, G., Card, S., and Mackinlay, J. (1993) "Information Visualization Using 3D Interactive Animation." Communications of the ACM, 36(4).

[Raja04] Raja, D., Bowman, D., Lucas, J., and North, C. (2004) "Exploring the Benefits of Immersion in Abstract Information Visualization." Proceedings of 8th Annual Immersive Projection Technology Conference

[Ross88] Ross, S., (1988) "A First Course in Probability" Macmillan Publishing Company

[Salzman96] 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).

[Simon04] Simon, A., Smith, R., and Pawlicki, R. (2004) "OmniStereo for Panoramic Virtual Environment Display Systems" Proceedings of IEEE Virtual Reality, pp. 67-74.

[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 ACM VRST, 163-172.

[Slater03] Slater, M., (2003) "A Note on Presence Terminology." www.presence-connect.com, January

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

[Ware96] Ware, C. and Franck, G., (1996) "Evaluating Stereo and Motion Cues for Visualizing Infomation Nets in Three Dimensions." ACM Transactions on Graphics, 15, 121-140.

[Ware05] Ware, C., and Mitchell, P. (2005) "Reevaluating Stereo and Motion Cues for Visualizing Graphs in Three Dimensions." Proceedings of ACM?

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

Appendix A: Pilot Study Forms (Scatter Plots with CaveDataView)

Appendix A.1: Instructions

General Instructions for Pilot Study

Thank you for volunteering to participate in our study. For this study we are evaluating the usefulness of Virtual Environments for analyzing statistical data. In this study you will be asked to perform some visualization tasks using the CAVE. During the experiment we will be writing down notes and recording time data. We are not evaluating your performance; we are only evaluating the usefulness of the VE for analyzing statistical data.

There will be two phases to the study as we vary different conditions. For each phase we will ask you to perform two tasks on different sets of data. The instructions for each phase will be explained to you before you begin. Please let us know if you have any questions or if you are unclear about the instructions so that we can clarify the task for you. If you have any comments about the system or the tasks you are asked to perform please voice them aloud so that we may record them. If at any time you need to take break or feel dizzy or uncomfortable let us know and we will stop the experiment.

The CAVE is a 4-screen projection display for virtual environments. The CAVE uses stereo optic viewing and head tracking to make the experience look more realistic. You will wear a pair of head-tracked glasses to put the two stereo images together and to track your movement in the environment.

To fly through the visualization you will use a tracked wand device with a joystick. To fly, point the wand in the direction you wish to move and push forward on the joystick. Pushing backward on the joystick also moves you backward opposite of the direction you are pointing. To rotate your view left and right push left and right on the joystick.

Attached to the joystick is a virtual arrow that you manipulate by pointing in the direction you choose. Use the virtual arrow to select data points. When selected, a data point will show a number on all sides and the number will also appear near your wand. During the experiment you will be asked to locate data points by their location and report back their number. Once you locate a data point, point the virtual arrow at it and call out the number to the experimenter.

The first few minutes will allow you to become acquainted with the CAVE, the wand and the virtual arrow. You will be able to practice using a dataset similar to the ones where you will perform the tasks for the experiment. Please use this time to become fully acquainted with the system and it's operation. Feel free to ask any questions, and remember to tell us about any comments you have about the system and the experiment.

Thank you for participating,
Dheva and John

Appendix A.2: Pre-questionnaire:

Please help us to categorize our user population by completing the following items.

1) Gender (circle one): Male Female

2) Age: _____________

3) Do you wear glasses or contact lenses (circle one)?

No Glasses Contact Lenses

4) Occupation (if student, indicate graduate or undergraduate):

5) Major / Area of specialization (if student):

6) How often do you use computers...

...for work? (circle the best answer)	...for fun? (circle the best answer)
a. not at all	a. not at all
b. once a month	b. once a month
c. once a week	c. once a week
d. several times a week	d. several times a week
e. daily	e. daily

7) Does the "fun" use include a substantial amount of video game play
(circle one)

Yes No

8) Please rate your experience level with Virtual Environments

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

9) Have you ever used a virtual reality (VR) system? If so, please describe it (what type of display was used, what kind of application (e.g. game, architectural walk-through) was running, how did you interact with the system, etc.).

10) Rate your familiarity with Information Visualization: (circle one)

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

11) Please rate your experience level with 2D Scatterplots

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

12) Please rate your experience level with 3D Scatterplots

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

13) Which tools have you used to visualize data? (Excel, Spotfire, etc.)

Appendix A.3: Tasks

Tasks:

1. Find cube (which represents a data point) with highest Y value
2. Find cube that has the lowest X and lowest Y value
3. Do the points represent a trend?
4. How many clusters exist with greater than 20 cubes?
5. Find the pink colored cube
6. Find the two points furthest away from the dataset

Appendix A.4: Raw Data From Pilot Study

	Task					
Condition	**1**	**2**	**3**	**4**	**5**	**6**
1	37.7	30.5	55.4	21.4	11.34	30
2	55.9	53.38	0	37.7	32.28	47
3	26.85	34.1	-	33.6	17.8	-
4	27	25.98	125	64	-	23.19 & 52
	Task					
Subject	**1**	**2**	**3**	**4**	**5**	**6**
1	39.1	30.5	-	37.7	17.8	23.19
2	26.85	34.1	55.4	21.4	32.28	52
3	27	25.98	-	33.6	11.34	47
4	55.9	53.38	125	64	46.59	30

				Condition		
		1	**2**	**3**	**4**	
Task	**Task 4**	21.4	37.7	33.6	64	

Appendix B: Pilot Study Forms (Surface Plots with CaveEquationView)

Appendix B.1: Instructions

General Instructions for Pilot Study

Thank you for volunteering to participate in our study. For this study we are evaluating the usefulness of virtual environments (VE) for analyzing options pricing data displayed in surface plots. In this study you will be asked to perform some visualization tasks using the CAVE. During the experiment we will be writing down notes and recording time data. We are not evaluating your performance; we are only evaluating the usefulness of the VE for analyzing statistical data.

The CAVE is a 4 screen projection display for virtual environments. The CAVE uses stereo optic viewing and head tracking to make the experience look more realistic. You will wear a pair of head-tracked glasses to put the two stereo images together and to track your movement in the environment.

You will see a series of various plots that contain options pricing information. There will be a few questions asking you to identify relationships between the variables. After each plot is viewed, a few very short questions will be asked about your perceived feelings regarding the conditions.

Please let me know if you have any questions or if you are unclear about the instructions so that we can clarify the task for you. If you have any comments about the system or the tasks you are asked to perform please voice them aloud so that we may record them. If at <u>any</u> time you need to take break or feel dizzy or uncomfortable let us know and we will stop the experiment.

To fly through the visualization of the surface plot, you will use a tracked wand device with a joystick. To fly, point the wand in the direction you wish to move and push forward on the joystick. Pushing backward on the joystick also moves you backward opposite of the direction you are pointing. To rotate your view left and right push left and right on the joystick.

The first few minutes will allow you to become acquainted with the CAVE, the wand and navigation. You will be able to practice using a surface plot similar to the ones where you will perform the tasks for the experiment. Please use this time to become fully acquainted with the system and its operation. Feel free to ask any questions, and remember to tell us about any comments you have about the system and the experiment.

Thank you for participating,
Dheva

Appendix B.2: Pre-questionnaire

Please help us to categorize our user population by completing the following items.

1) Gender (circle one): Male Female

2) Age: ______________

3) Do you wear glasses or contact lenses (circle one)?

No Glasses Contact Lenses

4) Occupation (if student, indicate graduate or undergraduate):

5) Major / Area of specialization (if student):

6) How often do you use computers...

...for work? (circle the best answer)	...for fun? (circle the best answer)
a. not at all	a. not at all
b. once a month	b. once a month
c. once a week	c. once a week
d. several times a week	d. several times a week
e. daily	e. daily

7) Does the "fun" use include a substantial amount of video game play (circle one)

Yes No

8) Please rate your experience level with Virtual Environments

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

9) Have you ever used a virtual reality (VR) system? If so, please describe it (what type of display was used, what kind of application (e.g. game, architectural walk-through) was running, how did you interact with the system, etc.).

10) Rate your familiarity with Information Visualization: (circle one)

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

11) Please rate your experience level with 3D scatter plots

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

12) Please rate your experience level with 3D surface plots

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

13) Which tools have you used to visualize data? (Excel, Mathematica, MatLab, Spotfire, etc.)

Appendix B.3: Tasks

Tasks for 3D Options Pricing Surface Plots

BS_Out1.txt

Price as a function of stock price v. time

1. As the stock price increases, what happens to the price of the option?
2. What does this graph tell you about the relative increase in price that time has on the option, i.e., the rate of increase in price as the stock rises when time to expiration is short as opposed to long?
3. At which stock price(approx) does time seem to have the greatest effect on price increase?

BS_Out2.txt

Theta as a function of strike price v. time

1. Approx what strike price produces the maximum amount of decay?
2. What can one say about the decay of options with strike prices below the maximum decay, as opposed to the options of strike prices above the point of max decay when time remaining to expiration is short?
3. What interesting generalization might one be able to make about the decay across all strike prices when time remaining to expiration is great?

BS_Out3.txt

Gamma as a function of stock price v. volatility

1. What effect on gamma does increasing volatility generally have on the strike prices with greatest gamma when volatility is relatively low?
2. What does the graph tell you about the range of stock prices that possess gamma when volatility is increased?
3. As volatility increases, does gamma actually increase for some stock prices?

BS_Out4.txt

Delta as a function of stock price v. volatility

1. What effect does increasing volatility have on the rate of change of the delta?
2. For some stock prices increasing volatility increases the delta, and for other stock prices increasing volatility decreases the delta. When volatility is relatively low, what stock price does this shift occur?

BS_Out5.txt

Vega as a function of stock price v. time

1. At what approx stock price does vega appear to increase most rapidly as time increases?
2. At what approx point (stock price and time to expiration) is vega at an absolute max?

BS_Out6.txt

Rho as a function of strike price v. time

1. At what approx strike price does rho appear to increase most rapidly as time increases?
2. At what approx point (strike price and time to expiration) is rho at an absolute max?

BS_Out7.txt

Delta as a function of stock price v. time

1. At what point (stock price and time) is delta at a maximum?
2. For what range of time does an increase in the stock price produce sharp increases in the delta?

BS_Out8.txt

Gamma as a function of stock price v. time

1. What is the approximate range of stock prices where gamma is significant and time to expiration is short?
2. What general assumptions might one make about the time frame in which gamma is critically important?
3. When time to expiration is great, what happens to gamma as the stock price increases?

Appendix B.4: Raw data for pilot study with surface plots

Conditions

Plot 1	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Task 1	1	7	5	12	13	3	16	9
Task 2	11	63	27	150	144	8	30	10
Task 3	17	38	105	69	240	29	305	77
Disorientation	1	1	3	1	3	2	1	3
Usefulness	6	7	5	6	5	6	5	5
AvgTaskTime	10	36	46	77	132	13	117	32

Plot 2	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Task 1	35	29	40	92	275	25	25	224
Task 2	78	7	80	103	77	185	37	26
Task 3	5	77	4	26	139	133	100	65
Disorientation	1	2	1	3	2	1	3	1
Usefulness	6	5	6	6	6	5	6	5
AvgTaskTime	39	38	41	74	164	114	54	105

Plot 3	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Task 1	16	16	10	27	58	10	60	46
Task 2	64	35	71	159	153	148	102	75
Task 3	33	37	18	63	19	101	79	2
Disorientation	2	3	1	5	1	5	2	1

Usefulness	6	2	6	4	5	5	4	5
AvgTaskTime	38	29	33	83	77	86	80	41

Plot 4	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Task 1	100	30	36	79	39	34	27	18
Task 2	65	55	115	23	43	19	5	109
Disorientation	1	3	1	3	1	1	1	1
Usefulness	3	6	6	6	3	5	6	6
AvgTaskTime	83	43	76	51	41	27	16	64

Plot 5	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Task 1	85	27	196	35	239	82	77	165
Task 2	112	77	102	220	95	74	72	35
Disorientation	1	1	3	1	1	4	1	1
Usefulness	7	5	4	5	2	5	5	6
AvgTaskTime	99	52	149	128	167	78	75	100

Plot 6	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Task 1	136	46	35	10	111	90	70	13
Task 2	20	16	31	1	44	95	21	56
Disorientation	3	1	1	1	1	1	1	1
Usefulness	6	6	5	6	3	5	6	7
AvgTaskTime	78	31	33	6	78	93	46	35

Plot 7	1	2	3	4	5	6	7	8
Task 1	36	37	13	6	41	20	4	30
Task 2	37	120	169	170	125	67	39	38
Disorientation	1	4	1	2	1	1	1	2
Usefulness	6	5	5	4	6	3	3	6
AvgTaskTime	37	79	91	88	83	44	22	34

Plot 8	1	2	3	4	5	6	7	8
Task 1	26	45	36	22	8	45	61	6
Task 2	42	22	34	35	28	22	94	12
Task 3	4	3	15	37	1	9	15	9
Disorientation	1	1	2	1	1	1	5	1
Usefulness	6	6	7	5	5	6	5	5
AvgTaskTime	24	23	28	31	12	25	57	9

AvgDis(Cond)	1.4	2.0	1.6	2.1	1.4	2.0	1.9	1.4
AvgUse(Cond)	5.8	5.3	5.5	5.3	4.4	5.0	5.0	5.6

Appendix C: Combined Experiment Forms

Appendix C.1: Instructions

General Instructions for Experiment

Thank you for volunteering to participate in this study. In this experiment we are evaluating the usefulness of virtual environments (VE) for analyzing 3D data displayed in two formats; surface plots and scatter plots. In this study you will be asked to perform some visualization tasks using the CAVE. During the experiment I will be writing down notes and recording time data. Your performance is not being evaluated--only the usefulness of the VE for analyzing data is being evaluated, so please do not feel pressured in any manner.

The CAVE is a 4-screen projection display for virtual environments. The CAVE uses stereo optic viewing and head tracking to make the experience look more realistic. You will wear a pair of head-tracked glasses to put the two stereo images together and to track your movement in the environment.

You will see two scatter plots and two surface plots. There will be two questions asking you to perform tasks and an interpretive question for each plot. After each plot is viewed, your perceived feelings regarding the conditions and tasks will be recorded.

To fly through the visualization of either plot, you will use a tracked wand device with a joystick. To fly, point the wand in the direction you wish to move and push forward on the joystick. Pushing backward on the joystick moves you opposite of the direction you are pointing. To change your heading push left and right on the joystick. There are four buttons on the wand, two on the right side of the joystick, and two on the left. Push on either of the buttons to the right to rotate the plot clockwise, or push on either of the buttons to the left of the joystick to rotate counter-clockwise.

The first few minutes will allow you to become acquainted with the CAVE, the wand and navigation. You will be able to practice using a scatter plot and surface plot similar to the ones used in the experiment. Please use this time to become fully acquainted with the system, its operation and feel free to ask any questions.

Please let me know if you have any concerns or if the instructions are not clear. If you have any comments about the system or about the tasks you are asked to perform please voice them aloud so that they may be recorded. If at any time you need to take break, if you feel dizzy or uncomfortable let me know and the experiment will be paused.

Thank you for participating,

Dheva

Appendix C.2: Pre-questionnaire

Please help us to categorize our subject population by completing the following items.

1) Gender (circle one): Male Female

2) Age: _____________

3) Do you wear glasses or contact lenses (circle one)?

No Glasses Contact Lenses

4) Occupation (if student, indicate graduate or undergraduate):

5) Major / Area of specialization:

6) How often do you use computers...

...for work? (circle the best answer)	**..for fun?**
a. not at all	a. not at all
b. once a month	b. once a month
c. once a week	c. once a week
d. several times a week	d. several times a week
e. daily	e. daily

7) Does the "fun" use include a substantial amount of video game play (circle one)

Yes No

8) Please rate your experience level with Virtual Environments

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

9) Have you ever used a virtual reality (VR) system? If so, please describe it (what type of display was used, what kind of application (e.g. game, architectural walk-through) was running, how did you interact with the system, etc.).

10) Rate your familiarity with Information Visualization: (circle one)

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

11) Please rate your experience level with 3D scatter plots

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

12) Please rate your experience level with 3D surface plots

No Experience 1-------2-------3-------4-------5-------6-------7 Expert

13) Which tools have you used to visualize data? (Excel, Mathematica, MatLab, Spotfire, etc.)

Appendix C.3: Tasks

Scatter Plot: Collegedata.txt and Financialdata.txt

Tasks/Questions:

1. Find the pink colored cube.
2. Do the points indicate a trend? If so, what is the relationship between the variables?
3. Find the cube with the greatest Z value (height).

Surface Plot: BS_Out1.txt and BS_Out6.txt

Tasks/Questions:

1. When Z is at a max, which X and Y values correspond to this point?
2. When X and Y are at the midpoint of the graph (intersection of green and blue axis lines), what happens to Z (inc or dec) as X increases?
3. As Y increases, which region of X produces the steepest slope up in the +Z direction (lower third, mid, or upper)?

Appendix C.4: Data from Combined Experiment

Descriptive Statistics

	N	Minimum	Maximum	Mean	Std. Deviation
Gender	32	0	1	.75	.440
Age	32	20	59	27.19	7.494
Computer Use (Work)	32	4	5	4.91	.296
Computer Use (Fun)	32	1	5	4.25	1.320
Video Game Player	32	0	1	.44	.504
VE Experience	32	1	7	2.87	1.963
Info Viz Experience	32	1	6	3.12	1.718
Scatter Plot Experience	32	1	6	2.66	1.494
Valid N (listwise)	32				

Descriptive Statistics

	Physical Immersion Level	Head Based Rendering (On/Off)	Stereoscopic Viewing (On/Off)	Mean	Std. Deviation	N
Scatter Plots Task 1 Time	Low	Off	Off	35.75	34.317	8
			On	20.13	12.088	8
			Total	27.94	26.132	16
		On	Off	23.38	17.663	8
			On	22.00	9.196	8
			Total	22.69	13.622	16
		Total	Off	29.56	27.129	16
			On	21.06	10.421	16
			Total	25.31	20.672	32
	High	Off	Off	28.75	18.266	8
			On	25.75	14.888	8
			Total	27.25	16.172	16
		On	Off	49.13	50.155	8
			On	35.57	39.386	7
			Total	42.80	44.403	15
		Total	Off	38.94	37.952	16
			On	30.33	28.309	15
			Total	34.77	33.365	31
	Total	Off	Off	32.25	26.802	16
			On	22.94	13.419	16
			Total	27.59	21.380	32
		On	Off	36.25	38.682	16
			On	28.33	27.500	15
			Total	32.42	33.425	31
		Total	Off	34.25	32.799	32
			On	25.55	21.224	31
			Total	29.97	27.840	63
Scatter Plots Task 2 Time	Low	Off	Off	65.00	39.663	8
			On	57.63	30.720	8
			Total	61.31	34.482	16
		On	Off	43.13	31.697	8
			On	35.75	17.686	8
			Total	39.44	25.086	16
		Total	Off	54.06	36.477	16
			On	46.69	26.720	16
			Total	50.37	31.676	32
	High	Off	Off	33.38	18.500	8
			On	37.63	25.707	8
			Total	35.50	21.747	16
		On	Off	44.25	11.424	8
			On	38.14	21.248	7

		Total	Off	54.06	36.477	16
			On	46.69	26.720	16
			Total	50.37	31.676	32
	High	Off	Off	33.38	18.500	8
			On	37.63	25.707	8
			Total	35.50	21.747	16
		On	Off	44.25	11.424	8
			On	38.14	21.248	7
			Total	41.40	16.392	15
		Total	Off	38.81	15.880	16
			On	37.87	22.891	15
			Total	38.35	19.257	31
	Total	Off	Off	49.19	34.067	16
			On	47.63	29.248	16
			Total	48.41	31.243	32
		On	Off	43.69	23.024	16
			On	36.87	18.746	15
			Total	40.39	21.001	31
		Total	Off	46.44	28.738	32
			On	42.42	24.932	31
			Total	44.46	26.792	63
Scatter Plots Task 3 Time	Low	Off	Off	20.50	12.375	8
			On	14.13	6.446	8
			Total	17.31	10.084	16
		On	Off	23.13	18.489	8
			On	21.25	28.034	8
			Total	22.19	22.961	16
		Total	Off	21.81	15.259	16
			On	17.69	19.992	16
			Total	19.75	17.620	32
	High	Off	Off	22.00	16.274	8
			On	19.75	20.464	8
			Total	20.88	17.899	16
		On	Off	12.62	5.097	8
			On	11.00	4.435	7
			Total	11.87	4.704	15
		Total	Off	17.31	12.616	16
			On	15.67	15.435	15
			Total	16.52	13.837	31
	Total	Off	Off	21.25	13.988	16
			On	16.94	14.942	16
			Total	19.09	14.405	32
		On	Off	17.88	14.179	16
			On	16.47	20.722	15
			Total	17.19	17.362	31
		Total	Off	19.56	13.961	32
			On	16.71	17.666	31
			Total	18.16	15.828	63
Scatter Plots Disorientation Rating	Low	Off	Off	2.63	1.598	8
			On	1.38	.744	8
			Total	2.00	1.366	16
		On	Off	2.00	.926	8
			On	1.75	1.035	8
			Total	1.88	.957	16
		Total	Off	2.31	1.302	16
			On	1.56	.892	16
			Total	1.94	1.162	32
	High	Off	Off	1.88	.991	8
			On	1.75	1.165	8
			Total	1.81	1.047	16
		On	Off	1.88	.641	8
			On	1.71	.488	7
			Total	1.80	.561	15
		Total	Off	1.88	.806	16
			On	1.73	.884	15
			Total	1.81	.833	31
	Total	Off	Off	2.25	1.342	16
			On	1.56	.964	16
			Total	1.91	1.201	32
		On	Off	1.94	.772	16
			On	1.73	.799	15
			Total	1.84	.779	31
		Total	Off	2.09	1.088	32
			On	1.65	.877	31
			Total	1.87	1.008	63
Scatter Plots Ease of Task Completion Rating	Low	Off	Off	5.00	.535	8
			On	5.88	1.356	8
			Total	5.44	1.094	16
		On	Off	5.88	.991	8
			On	5.88	1.246	8
			Total	5.88	1.088	16
		Total	Off	5.44	.892	16
			On	5.88	1.258	16
			Total	5.66	1.096	32
	High	Off	Off	5.88	1.356	8
			On	6.00	.535	8
			Total	5.94	.998	16
		On	Off	5.00	.756	8
			On	6.14	.378	7
			Total	5.53	.834	15

Correlations

		Gender	Age	Computer Use (Work)	Computer Use (Fun)	Video Game Player	VE Experience	Info Viz Experience	Scatter Plot Experience	Surface Plot Experience	Scatter Plots Task 1 Time	Scatter Plots Task 2 Time	Scatter Plots Task 3 Time	Scatter Plots Disorientation Rating	Scatter Plots Ease of Task Completion Rating
Gender	Pearson Correlation	1	-.298*	.309*	.500**	.509**	.224	.000	.159	.073	-.031	-.144	.014	-.023	.261*
	Sig. (2-tailed)	.	.017	.013	.000	.000	.075	1.000	.208	.569	.806	.257	.910	.857	.039
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Age	Pearson Correlation	-.298*	1	.168	-.341**	-.270*	-.014	.143	.084	.127	.154	-.055	-.042	-.132	-.096
	Sig. (2-tailed)	.017	.	.184	.006	.031	.914	.258	.511	.317	.223	.665	.743	.298	.454
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Computer Use (Work)	Pearson Correlation	.309*	.168	1	.144	.068	.201	.404**	.362**	.461**	.095	-.145	-.016	-.104	.172
	Sig. (2-tailed)	.013	.184	.	.255	.596	.111	.001	.003	.000	.453	.253	.901	.414	.177
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Computer Use (Fun)	Pearson Correlation	.500**	-.341**	.144	1	.412**	.436**	.085	.192	.208	-.075	.124	-.163	-.054	.273*
	Sig. (2-tailed)	.000	.006	.255	.	.001	.000	.502	.128	.099	.555	.328	.198	.674	.030
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Video Game Player	Pearson Correlation	.509**	-.270*	.068	.412**	1	.383**	.307*	.292*	.358**	-.018	-.021	-.158	.035	.133
	Sig. (2-tailed)	.000	.031	.596	.001	.	.002	.013	.019	.004	.886	.867	.212	.783	.300
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
VE Experience	Pearson Correlation	.224	-.014	.201	.436**	.383**	1	.646**	.755**	.652**	-.210	-.061	.080	-.298*	.232
	Sig. (2-tailed)	.075	.914	.111	.000	.002	.	.000	.000	.000	.095	.633	.528	.017	.067
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Info Viz Experience	Pearson Correlation	.000	.143	.404**	.085	.307*	.646**	1	.771**	.652**	-.007	-.171	.076	-.177	.193
	Sig. (2-tailed)	1.000	.258	.001	.502	.013	.000	.	.000	.000	.958	.176	.552	.161	.129
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Scatter Plot Experience	Pearson Correlation	.159	.084	.362**	.192	.292*	.755**	.771**	1	.796**	-.048	-.284*	.042	-.226	.228
	Sig. (2-tailed)	.208	.511	.003	.128	.019	.000	.000	.	.000	.704	.023	.744	.073	.072
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Surface Plot Experience	Pearson Correlation	.073	.127	.461**	.208	.358**	.652**	.652**	.796**	1	-.096	-.024	-.031	-.220	.073
	Sig. (2-tailed)	.569	.317	.000	.099	.004	.000	.000	.000	.	.452	.853	.809	.080	.569
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Scatter Plots Task 1 Time	Pearson Correlation	-.031	.154	.095	-.075	-.018	-.210	-.007	-.048	-.096	1	-.189	-.110	.291*	-.274*
	Sig. (2-tailed)	.806	.223	.453	.555	.886	.095	.958	.704	.452	.	.135	.387	.020	.029
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Scatter Plots Task 2 Time	Pearson Correlation	-.144	-.055	-.145	.124	-.021	-.061	-.171	-.284*	-.024	-.189	1	.141	.130	-.356**
	Sig. (2-tailed)	.257	.665	.253	.328	.867	.633	.176	.023	.853	.135	.	.265	.306	.004
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Scatter Plots Task 3 Time	Pearson Correlation	.014	-.042	-.016	-.163	-.158	.080	.076	.042	-.031	-.110	.141	1	.013	-.196
	Sig. (2-tailed)	.910	.743	.901	.198	.212	.528	.552	.744	.809	.387	.265	.	.917	.124
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Scatter Plots Disorientation Rating	Pearson Correlation	-.023	-.132	-.104	-.054	.035	-.298*	-.177	-.226	-.220	.291*	.130	.013	1	-.498**
	Sig. (2-tailed)	.857	.298	.414	.674	.783	.017	.161	.073	.080	.020	.306	.917	.	.000
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	63
Scatter Plots Ease of Task Completion Rating	Pearson Correlation	.261*	-.096	.172	.273*	.133	.232	.193	.228	.073	-.274*	-.356**	-.196	-.498**	1
	Sig. (2-tailed)	.039	.454	.177	.030	.300	.067	.129	.072	.569	.029	.004	.124	.000	.
	N	63	63	63	63	63	63	63	63	63	63	63	63	63	63

*. Correlation is significant at the 0.05 level (2-tailed).

**. Correlation is significant at the 0.01 level (2-tailed).

Descriptive Statistics

	Physical Immersion Level	Head Based Rendering (On/Off)	Stereoscopic Viewing (On/Off)	Mean	Std. Deviation	N
Surface Plots Task 1 Time	Low	Off	Off	73.25	36.955	8
			On	50.13	28.412	8
			Total	61.69	34.009	16
		On	Off	58.25	41.952	8
			On	80.00	49.926	8
			Total	69.12	45.942	16
		Total	Off	65.75	38.969	16
			On	65.06	42.165	16
			Total	65.41	39.940	32
	High	Off	Off	113.13	114.633	8
			On	104.75	87.382	8
			Total	108.94	98.561	16
		On	Off	73.50	39.170	8
			On	96.88	58.119	8
			Total	85.19	49.376	16
		Total	Off	93.31	85.247	16
			On	100.81	71.806	16
			Total	97.06	77.626	32
	Total	Off	Off	93.19	84.815	16
			On	77.44	68.817	16
			Total	85.31	76.396	32
		On	Off	65.88	39.991	16
			On	88.44	53.061	16
			Total	77.16	47.619	32
		Total	Off	79.53	66.687	32
			On	82.94	60.704	32
			Total	81.23	63.281	64
Surface Plots Task 2 Time	Low	Off	Off	23.63	19.018	8
			On	15.63	7.444	8
			Total	19.63	14.550	16
		On	Off	15.13	13.527	8
			On	15.75	10.512	8
			Total	15.44	11.707	16
		Total	Off	19.38	16.536	16
			On	15.69	8.799	16
			Total	17.53	13.164	32
	High	Off	Off	14.13	9.538	8
			On	21.25	15.691	8
			Total	17.69	13.073	16
		On	Off	11.63	6.968	8
			On	25.00	24.219	8
			Total	18.31	18.550	16
		Total	Off	12.88	8.172	16
			On	23.13	19.809	16
			Total	18.00	15.789	32
	Total	Off	Off	18.88	15.340	16
			On	18.44	12.215	16
			Total	18.66	13.642	32
		On	Off	13.38	10.551	16
			On	20.37	18.658	16
			Total	16.88	15.328	32
		Total	Off	16.13	13.249	32
			On	19.41	15.544	32
			Total	17.77	14.422	64
Surface Plots Task 3 Time	Low	Off	Off	47.75	44.819	8
			On	42.38	36.742	8
			Total	45.06	39.688	16
		On	Off	34.13	19.082	8
			On	23.25	14.548	8
			Total	28.69	17.327	16
		Total	Off	40.94	34.013	16
			On	32.81	28.745	16
			Total	36.88	31.251	32
	High	Off	Off	37.13	42.259	8
			On	39.00	21.193	8
			Total	38.06	32.310	16
		On	Off	29.25	19.129	8
			On	35.75	29.334	8
			Total	32.50	24.158	16
		Total	Off	33.19	31.948	16
			On	37.38	24.779	16
			Total	35.28	28.205	32
	Total	Off	Off	42.44	42.437	16
			On	40.69	29.028	16
			Total	41.56	35.776	32
		On	Off	31.69	18.629	16
			On	29.50	23.281	16
			Total	30.59	20.770	32
		Total	Off	37.06	32.698	32
			On	35.09	26.501	32
			Total	36.08	29.541	64
Surface Plots Disorientation Rating	Low	Off	Off	2.63	1.685	8
			On	1.75	1.165	8

			On	35.09	26.501	32
			Total	36.08	29.541	64
Surface Plots Disorientation Rating	Low	Off	Off	2.63	1.685	8
			On	1.75	1.165	8
			Total	2.19	1.471	16
		On	Off	1.88	1.356	8
			On	2.50	1.195	8
			Total	2.19	1.276	16
		Total	Off	2.25	1.528	16
			On	2.13	1.204	16
			Total	2.19	1.355	32
	High	Off	Off	2.13	1.356	8
			On	1.38	.518	8
			Total	1.75	1.065	16
		On	Off	1.63	.744	8
			On	1.75	1.389	8
			Total	1.69	1.078	16
		Total	Off	1.88	1.088	16
			On	1.56	1.031	16
			Total	1.72	1.054	32
	Total	Off	Off	2.38	1.500	16
			On	1.56	.892	16
			Total	1.97	1.282	32
		On	Off	1.75	1.065	16
			On	2.12	1.310	16
			Total	1.94	1.190	32
		Total	Off	2.06	1.318	32
			On	1.84	1.139	32
			Total	1.95	1.227	64
Surface Plots Ease of Task Completion Rating	Low	Off	Off	4.50	1.069	8
			On	5.25	1.035	8
			Total	4.88	1.088	16
		On	Off	5.00	1.195	8
			On	5.00	1.069	8
			Total	5.00	1.095	16
		Total	Off	4.75	1.125	16
			On	5.13	1.025	16
			Total	4.94	1.076	32
	High	Off	Off	4.38	1.598	8
			On	3.25	1.389	8
			Total	3.81	1.559	16
		On	Off	3.63	1.302	8
			On	5.13	1.642	8
			Total	4.37	1.628	16

Correlations

		Gender	Age	Computer Use (Work)	Computer Use (Fun)	Video Game Player	VE Experience	Info Viz Experience	Scatter Plot Experience	Surface Plot Experience	Surface Plots Task 1 Time	Surface Plots Task 2 Time	Surface Plots Task 3 Time	Surface Plots Disorientation Rating	Surface Plots Ease of Task Completion Rating	Surface Plots Task 1 Accuracy
Gender	Pearson Correlation	1	-.298*	.309*	.500**	.509**	.224	.000	.159	.073	.052	-.186	-.060	-.289*	.109	-.012
	Sig. (2-tailed)	.	.017	.013	.000	.000	.075	1.000	.208	.569	.686	.141	.638	.021	.390	.922
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Age	Pearson Correlation	-.298*	1	.168	-.341**	-.270*	-.014	.143	.084	.127	.170	.216	.228	.058	-.126	.077
	Sig. (2-tailed)	.017	.	.184	.006	.031	.914	.258	.511	.317	.180	.086	.070	.647	.323	.544
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Computer Use (Work)	Pearson Correlation	.309*	.168	1	.144	.068	.201	.404**	.362**	.461**	.136	.043	-.005	.032	-.035	.134
	Sig. (2-tailed)	.013	.184	.	.255	.596	.111	.001	.003	.000	.284	.733	.971	.804	.786	.291
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Computer Use (Fun)	Pearson Correlation	.500**	-.341**	.144	1	.412**	.436**	.085	.192	.208	.138	-.062	-.076	-.240	.041	.134
	Sig. (2-tailed)	.000	.006	.255	.	.001	.000	.502	.128	.099	.277	.624	.553	.057	.749	.292
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Video Game Player	Pearson Correlation	.509**	-.270*	.068	.412**	1	.383**	.307*	.292*	.358**	-.135	-.206	-.101	-.354**	-.055	-.070
	Sig. (2-tailed)	.000	.031	.596	.001	.	.002	.013	.019	.004	.288	.103	.426	.004	.668	.584
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
VE Experience	Pearson Correlation	.224	-.014	.201	.436**	.383**	1	.646**	.755**	.652**	-.079	-.076	-.201	-.202	-.063	-.050
	Sig. (2-tailed)	.075	.914	.111	.000	.002	.	.000	.000	.000	.533	.549	.110	.110	.623	.698
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Info Viz Experience	Pearson Correlation	.000	.143	.404**	.085	.307*	.646**	1	.771**	.652**	-.101	-.051	-.072	-.081	-.014	-.020
	Sig. (2-tailed)	1.000	.258	.001	.502	.013	.000	.	.000	.000	.425	.688	.570	.526	.913	.875
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Scatter Plot Experience	Pearson Correlation	.159	.084	.362**	.192	.292*	.755**	.771**	1	.796**	-.123	-.092	-.254*	-.053	-.141	-.089
	Sig. (2-tailed)	.208	.511	.003	.128	.019	.000	.000	.	.000	.334	.469	.043	.680	.266	.482
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Surface Plot Experience	Pearson Correlation	.073	.127	.461**	.208	.358**	.652**	.652**	.796**	1	-.221	-.122	-.256*	-.162	-.088	-.065
	Sig. (2-tailed)	.569	.317	.000	.099	.004	.000	.000	.000	.	.079	.336	.041	.202	.491	.610
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Surface Plots Task 1 Time	Pearson Correlation	.052	.170	.136	.138	-.135	-.079	-.101	-.123	-.221	1	.232	.267*	.345**	-.248*	.030
	Sig. (2-tailed)	.686	.180	.284	.277	.288	.533	.425	.334	.079	.	.065	.033	.005	.049	.812
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Surface Plots Task 2 Time	Pearson Correlation	-.186	.216	.043	-.062	-.206	-.076	-.051	-.092	-.122	.232	1	.411**	.229	-.048	-.207
	Sig. (2-tailed)	.141	.086	.733	.624	.103	.549	.688	.469	.336	.065	.	.001	.069	.704	.100
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Surface Plots Task 3 Time	Pearson Correlation	-.060	.228	-.005	-.076	-.101	-.201	-.072	-.254*	-.256*	.267*	.411**	1	.242	-.157	.150
	Sig. (2-tailed)	.638	.070	.971	.553	.426	.110	.570	.043	.041	.033	.001	.	.054	.216	.236
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Surface Plots Disorientation Rating	Pearson Correlation	-.289*	.058	.032	-.240	-.354**	-.202	-.081	-.053	-.162	.345**	.229	.242	1	-.068	.013
	Sig. (2-tailed)	.021	.647	.804	.057	.004	.110	.526	.680	.202	.005	.069	.054	.	.592	.921
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Surface Plots Ease of Task Completion Rating	Pearson Correlation	.109	-.126	-.035	.041	-.055	-.063	-.014	-.141	-.088	-.248*	-.048	-.157	-.068	1	.152
	Sig. (2-tailed)	.390	.323	.786	.749	.668	.623	.913	.266	.491	.049	.704	.216	.592	.	.230
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64
Surface Plots Task 1 Accuracy	Pearson Correlation	-.012	.077	.134	.134	-.070	-.050	-.020	-.089	-.065	.030	-.207	.150	.013	.152	1
	Sig. (2-tailed)	.922	.544	.291	.292	.584	.698	.875	.482	.610	.812	.100	.236	.921	.230	.
	N	64	64	64	64	64	64	64	64	64	64	64	64	64	64	64

*. Correlation is significant at the 0.05 level (2-tailed).

**. Correlation is significant at the 0.01 level (2-tailed).

Appendix C.5: Statistical Analysis for Combined Experiment

Scatter Plots

Tests of Between-Subjects Effects

Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	Scatter Plots Task 1 Time	5302.109[a]	7	757.444	.989	.448
	Scatter Plots Task 2 Time	7405.000[b]	7	1057.857	1.561	.166
	Scatter Plots Task 3 Time	1392.109[c]	7	198.873	.772	.613
Intercept	Scatter Plots Task 1 Time	58141.266	1	58141.266	75.946	.000
	Scatter Plots Task 2 Time	123904.000	1	123904.000	182.821	.000
	Scatter Plots Task 3 Time	20484.766	1	20484.766	79.497	.000
IMMERS	Scatter Plots Task 1 Time	1491.891	1	1491.891	1.949	.168
	Scatter Plots Task 2 Time	2601.000	1	2601.000	3.838	.055
	Scatter Plots Task 3 Time	221.266	1	221.266	.859	.358
HBR	Scatter Plots Task 1 Time	415.141	1	415.141	.542	.465
	Scatter Plots Task 2 Time	1242.563	1	1242.563	1.833	.181
	Scatter Plots Task 3 Time	92.641	1	92.641	.360	.551
STEREO	Scatter Plots Task 1 Time	1080.766	1	1080.766	1.412	.240
	Scatter Plots Task 2 Time	380.250	1	380.250	.561	.457
	Scatter Plots Task 3 Time	178.891	1	178.891	.694	.408
IMMERS * HBR	Scatter Plots Task 1 Time	1711.891	1	1711.891	2.236	.140
	Scatter Plots Task 2 Time	2730.063	1	2730.063	4.028	.050
	Scatter Plots Task 3 Time	848.266	1	848.266	3.292	.075
IMMERS * STEREO	Scatter Plots Task 1 Time	1.266	1	1.266	.002	.968
	Scatter Plots Task 2 Time	100.000	1	100.000	.148	.702
	Scatter Plots Task 3 Time	9.766	1	9.766	.038	.846
HBR * STEREO	Scatter Plots Task 1 Time	19.141	1	19.141	.025	.875
	Scatter Plots Task 2 Time	175.563	1	175.563	.259	.613
	Scatter Plots Task 3 Time	15.016	1	15.016	.058	.810
IMMERS * HBR * STEREO	Scatter Plots Task 1 Time	582.016	1	582.016	.760	.387
	Scatter Plots Task 2 Time	175.563	1	175.563	.259	.613
	Scatter Plots Task 3 Time	26.266	1	26.266	.102	.751
Error	Scatter Plots Task 1 Time	42871.625	56	765.565		
	Scatter Plots Task 2 Time	37953.000	56	677.732		
	Scatter Plots Task 3 Time	14430.125	56	257.681		
Total	Scatter Plots Task 1 Time	106315.000	64			
	Scatter Plots Task 2 Time	169262.000	64			
	Scatter Plots Task 3 Time	36307.000	64			
Corrected Total	Scatter Plots Task 1 Time	48173.734	63			
	Scatter Plots Task 2 Time	45358.000	63			
	Scatter Plots Task 3 Time	15822.234	63			

a. R Squared = .110 (Adjusted R Squared = -.001)

b. R Squared = .163 (Adjusted R Squared = .059)

c. R Squared = .088 (Adjusted R Squared = -.026)

Surface Plots

Tests of Between-Subjects Effects

Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	Surface Plots Task 1 Time	27486.359[a]	7	3926.623	.978	.456
	Surface Plots Task 2 Time	1323.109[b]	7	189.016	.899	.514
	Surface Plots Task 3 Time	3204.984[c]	7	457.855	.495	.834
	Surface Plots Task 1 Accuracy	3767.437[d]	7	538.205	1.594	.156
Intercept	Surface Plots Task 1 Time	422337.516	1	422337.516	105.210	.000
	Surface Plots Task 2 Time	20199.516	1	20199.516	96.022	.000
	Surface Plots Task 3 Time	83304.391	1	83304.391	90.108	.000
	Surface Plots Task 1 Accuracy	493857.563	1	493857.563	1463.049	.000
IMMERS	Surface Plots Task 1 Time	16033.891	1	16033.891	3.994	.051
	Surface Plots Task 2 Time	3.516	1	3.516	.017	.898
	Surface Plots Task 3 Time	40.641	1	40.641	.044	.835
	Surface Plots Task 1 Accuracy	992.250	1	992.250	2.940	.092
HBR	Surface Plots Task 1 Time	1064.391	1	1064.391	.265	.609
	Surface Plots Task 2 Time	50.766	1	50.766	.241	.625
	Surface Plots Task 3 Time	1925.016	1	1925.016	2.082	.155
	Surface Plots Task 1 Accuracy	1040.063	1	1040.063	3.081	.085
STEREO	Surface Plots Task 1 Time	185.641	1	185.641	.046	.831
	Surface Plots Task 2 Time	172.266	1	172.266	.819	.369
	Surface Plots Task 3 Time	62.016	1	62.016	.067	.797
	Surface Plots Task 1 Accuracy	1.000	1	1.000	.003	.957
IMMERS * HBR	Surface Plots Task 1 Time	3890.641	1	3890.641	.969	.329
	Surface Plots Task 2 Time	92.641	1	92.641	.440	.510
	Surface Plots Task 3 Time	467.641	1	467.641	.506	.480
	Surface Plots Task 1 Accuracy	1296.000	1	1296.000	3.839	.055
IMMERS * STEREO	Surface Plots Task 1 Time	268.141	1	268.141	.067	.797

	1 Accuracy	1296.000	1	1296.000	3.839	.055
IMMERS * STEREO	Surface Plots Task 1 Time	268.141	1	268.141	.067	.797
	Surface Plots Task 2 Time	777.016	1	777.016	3.694	.060
	Surface Plots Task 3 Time	606.391	1	606.391	.656	.421
	Surface Plots Task 1 Accuracy	264.063	1	264.063	.782	.380
HBR * STEREO	Surface Plots Task 1 Time	5871.391	1	5871.391	1.463	.232
	Surface Plots Task 2 Time	221.266	1	221.266	1.052	.309
	Surface Plots Task 3 Time	.766	1	.766	.001	.977
	Surface Plots Task 1 Accuracy	169.000	1	169.000	.501	.482
IMMERS * HBR * STEREO	Surface Plots Task 1 Time	172.266	1	172.266	.043	.837
	Surface Plots Task 2 Time	5.641	1	5.641	.027	.871
	Surface Plots Task 3 Time	102.516	1	102.516	.111	.740
	Surface Plots Task 1 Accuracy	5.063	1	5.063	.015	.903
Error	Surface Plots Task 1 Time	224797.125	56	4014.234		
	Surface Plots Task 2 Time	11780.375	56	210.364		
	Surface Plots Task 3 Time	51771.625	56	924.493		
	Surface Plots Task 1 Accuracy	18903.000	56	337.554		
Total	Surface Plots Task 1 Time	674621.000	64			
	Surface Plots Task 2 Time	33303.000	64			
	Surface Plots Task 3 Time	138281.000	64			
	Surface Plots Task 1 Accuracy	516528.000	64			
Corrected Total	Surface Plots Task 1 Time	252283.484	63			
	Surface Plots Task 2 Time	13103.484	63			
	Surface Plots Task 3 Time	54976.609	63			
	Surface Plots Task 1 Accuracy	22670.437	63			

a. R Squared = .109 (Adjusted R Squared = -.002)

b. R Squared = .101 (Adjusted R Squared = -.011)

c. R Squared = .058 (Adjusted R Squared = -.059)

d. R Squared = .166 (Adjusted R Squared = .062)

Appendix D: A Basic Primer on Stock Options and the Black-Scholes Formula

This guide is intended as a very basic introduction to stock options, and the associated terminology that were used in the surface plot pilot study. For the interested reader, have a look at a good introductory text (applied to commodities options) Natenberg's "Option Volatility and Pricing Strategies." For a more comprehensive treatment, read "Options, Futures, and Other Derivatives," by John C. Hull. If the reader is interested in a mathematically based presentation, explore "Arbitrage Theory in Continuous Time," by Tomas Bjork. There are a huge number of books and materials on the subject, but these mentioned above are among my favorites.

The values calculated in the surface plots were generated using the Black-Scholes equation for pricing European options on non-dividend paying stocks.

Listed options are typically "vanilla" *calls* and *puts*. A *call option* is a derivative (a derivative is a security whose value is *derived* from an underlying asset—stock, bond, commodity, etc.) with the right, but not the obligation, to buy the asset at a certain price, at a particular date in the future. The particular date is termed the *expiration date* of the option, and the certain price is termed the *strike* or *exercise price*.

The general behavior of a call option may be illustrated by the following example. Suppose an interesting energy company exists named *Enbomb*. The stock seems to have been doing well lately and a savvy investor, Herbert Niedermeier Clueless, thinks the stock will move sharply upward soon. Instead of purchasing the stock directly, he decides to buy 1 *Enbomb* call option for $5 (the call price or *premium*) with an exercise price of $80, which has an expiration date of three months in the future. *Enbomb* stock is currently trading at $70. We will equate the price of *Enbomb* stock three months in the future (which we do not know now, of course) to S_T. Mr. Clueless decides to hold his call options until they expire in three months. What is the value of the options and what does he do? That really depends on S_T.

In three months, *Enbomb* stock is worth S_T.

- If S_T < $80, the call options expire worthless, and Mr. Clueless loses only the amount paid for the options. In this case, $5. The option payoff is zero, and the profit on the position is -$5.

- If S_T > $80, Mr. Clueless *exercises* the option to buy the stock. He pays $80 per share (the *exercise* price) to obtain *Enbomb* stock, and may sell it back into the market at S_T. Here the payoff of the option = S_T - $80, and the profit = S_T - $80 - $5.

The following figure shows the call payoff at expiration.

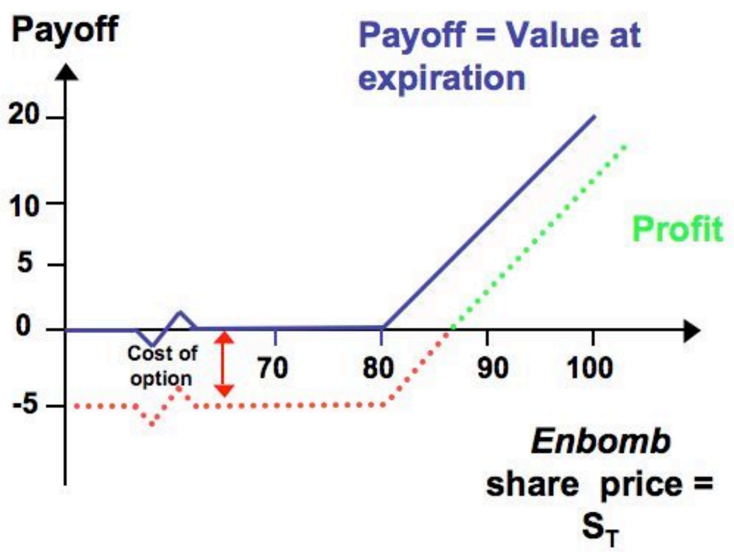

FIGURE D.1 – THE PAYOFF AND PROFIT OF AN ENBOMB CALL OPTION

The put option is similar to the call, except that the put option holder has the right, but not the obligation, to sell the underlying asset at a certain price, at a particular date in the future. Suppose an investor, Natasha Lookslikeamanskova, has heard some very negative rumors about Enbomb. She thinks the stock will tank, therefore she buys a put option for $5. In this case, assume *Enbomb* is currently trading at $90, with an exercise of $80, and an expiration date of three months in the future.

In three months, *Enbomb* stock is worth S_T.

- If S_T > $80, the call options expire worthless, and Ms. Lookslikeamanskova loses only the amount paid for the options. In this case, $5. The option payoff is zero, and the profit on the position is -$5.

- If S_T < $80, Ms. Lookslikeamanskova *exercises* the option to sell the stock. She pays $80 per share (the *exercise* price) to short sell (a sale which allows one to sell something they do not own to profit from a decline in price) *Enbomb* stock, and may buy it back in the market at S_T to close the position. Here the payoff of the option = $80 - S_T, and the profit = $80 - S_T - $5.

The following figure shows the put payoff at expiration:

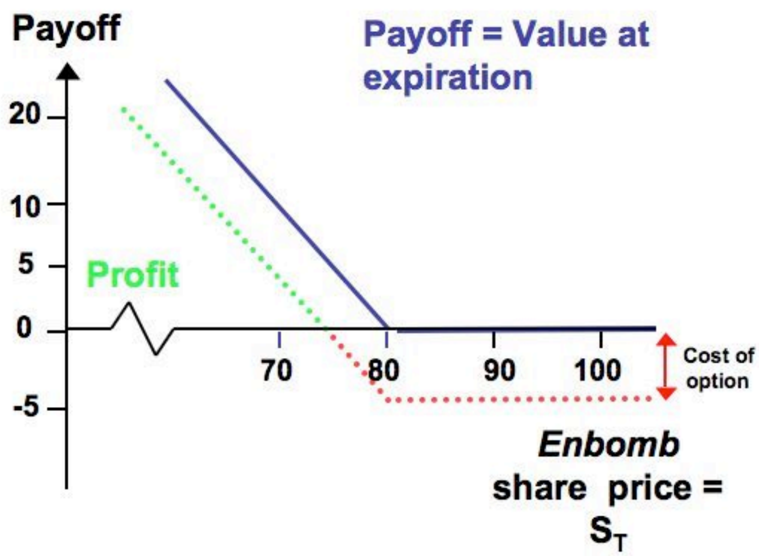

FIGURE D.2 – THE PAYOFF AND PROFIT OF AN ENBOMB PUT OPTION

So, now that the payoff at expiration is known, how is the value of the option determined throughout its life, that is, before it expires? This is where the Black-Scholes formula is useful, and is given by the equations below.

Call Option

$$c = S_o N(d_1) - Ke^{-rT} N(d_2)$$

Put Option

$$p = Ke^{-rT} N(-d_2) - S_o N(-d_1)$$

where

S_o = Current price of stock
K = Exercise price
r = Risk free interest rate, continuously compounded
T = Time remaining to expiration
σ = Volatility (annual) of underlying stock

$$d_1 = \frac{\ln(S_o/K)+(r+\sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = d_1 - \sigma\sqrt{T}$$

The function *N(x)* is the cumulative probability distribution function for a standardized normal distribution, ϕ(0, 1).

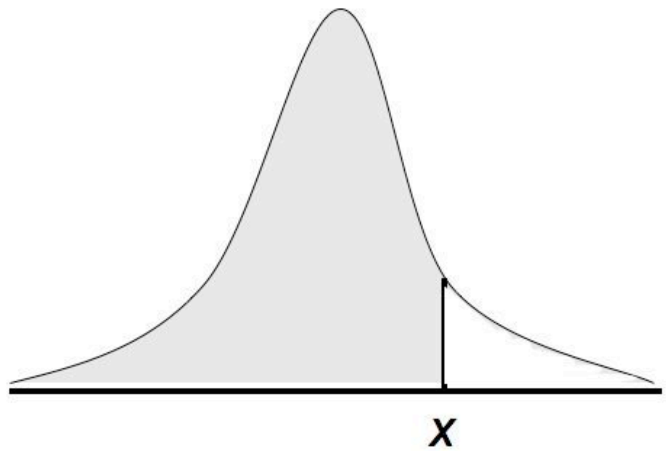

FIGURE D.3 – THE SHADED AREA UNDER THE CURVE TO THE LEFT OF X REPRESENTS N(X).

This value is given by the formula:

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-y^2/2} dy$$

A polynomial approximation used in the program written to calculate these values is:

$$N(x) = \begin{cases} 1 - N'(x)(a_1 k + a_2 k^2 + a_3 k^3 + a_4 k^4 + a_5 k^5) \rightarrow if(x) \geq 0 \\ 1 - N(-x) \rightarrow if(x) < 0 \end{cases}$$

$$k = \frac{1}{1 + \gamma x}, \gamma = 0.2316419$$

a_1 = 0.319381530, a_2 = -0.356563782, a_3 = 1.781477937, a_4 = -1.821255978, a_5 = 1.330274429

$$N'(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$$

The "Greeks"

The "greeks" measure various risk dimensions associated with options. They are crucial to a trader for managing risk. These components were included in the surface plots in our pilot study, and are described here for your convenience.

Delta

The *delta* is defined as the rate of change of the option price with respect to the price of the underlying asset (stocks in our cases). It can be shown that for a call option…

$$\Delta = N(d_1)$$

…and for a put option

$$\Delta = N(d_1) - 1$$

where d_1 is defined above.

Gamma

The *gamma* of an option is the rate of change of the delta with respect to the price of the underlying stock. It is the second partial derivative with respect to the stock price. It can be expressed for a European (non-dividend paying) call or put as …

$$\Gamma = \frac{N'(d_1)}{S_o \sigma \sqrt{T}}$$

where $N'(x)$ and d_1 defined above.

Theta

The *theta* of an option is the rate of change of the value with respect to the passage of time. For a call, this is expressed as…

$$\Theta = -\frac{S_o N'(d_1)\sigma}{2\sqrt{T}} - rKe^{-rT}N(d_2)$$

…for a put

$$\Theta = -\frac{S_o N'(d_1)\sigma}{2\sqrt{T}} + rKe^{-rT}N(-d_2)$$

where $N'(x)$, d_1, and d_2 are defined as above.

Vega

The *vega* of an option is the rate of change of the value of the portfolio with respect to the volatility of the underlying stock. For a call or put this is…

$$V = S_o\sqrt{T}N'(d_1)$$

where d_1 is defined previously.

Rho

The *rho* of an option is the rate of change of the value with respect to the interest rate. For a call this is given by…

$$\mathrm{P} = KTe^{-rT}N(d_2)$$

and for a put…

$$\mathrm{P} = -KTe^{-rT}N(-d_2)$$

where d_1 and d_2 defined previously.

I hope this guide has been helpful…if there are any errors please send me, Dheva Raja, an email at draja@vt.edu.

Vita

Dheva Raja was born on March 2, 1971. He holds a BA in Mathematics from New York University, an MBA from the Stockholm School of Economics, and an MS in Computer Science from Virginia Tech.

His professional life consisted of trading equity options on the floor of the American Stock Exchange. In his spare time, he enjoys cycling, auto/motorbike mechanics, reading, and taking care of his fiancee's horses.