Chapter VII

3D Interaction with Scientific Data Through Virtual Reality and Tangible Interfacing

Wen Qi
Eindhoven University of Technology, The Netherlands

Russell M. Taylor II
University of North Carolina–Chapel Hill, USA

Christopher Healey
North Carolina State University, USA

Jean-Bernard Martens
Eindhoven University of Technology, The Netherlands

ABSTRACT

Three-dimensional (3D) interaction with scientific data is still an immature topic. It involves studying visualization methods to faithfully represent data, on the one hand, and designing interfaces that truly assist users in the data analysis process, on the other hand. In this chapter, we study how the human computer interface influences performance in specific scientific visualization tasks. Although a wide range of virtual reality (VR) systems are in use today, there are few guidelines to help system and application developers in selecting the components most appropriate for the domain problem they are investigating. Using the results of an empirical study, we develop guidelines for the choice of display environment for four specific, but common, volume visualization tasks: identification and judgment of the size, shape, density, and connectivity of objects present in a volume. These tasks are derived from data analysis questions being asked by domain specialists studying Cystic Fibrosis (CF). We compared user performance in three different stereo VR systems: (1) a head-mounted display (HMD); (2) a fish tank

VR (fish tank); and (3) a fish tank VR augmented with a haptic device (haptic). HMD participants were placed inside the volume and walked within it to explore its structure. Fish tank and haptic participants saw the entire volume on-screen and rotated it to observe it from different perspectives. Response time and accuracy were used to measure performance. The results show that the fish tank and haptic groups were significantly more accurate at judging the shape, density, and connectivity of objects and completed the tasks significantly faster than the HMD group. Although the fish tank group was itself significantly faster than the haptic group, there were no statistical differences in accuracy between the two. Participants classified the HMD system as an inside-out display (looking outwards from inside the volume), and the fish tank and haptic systems as outside-in displays (looking inwards from outside the volume). Including haptics added an inside-out capability to the fish tank system through the use of touch. We recommend an outside-in system, since it offers both overview and context, two visual properties that are important for the volume visualization tasks we studied. In addition, based on the haptic group's opinion (80% positive) that haptic feedback aided comprehension, we recommend supplementing the outside-in visual display with inside-out haptics when possible. Based on the results from this user study, we further investigated the 3D interaction tasks from the design perspective of tangible interfaces. Since participants using the fish tank VR system performed better than the other groups in terms of time and accuracy, we asked the question whether or not the user performance could be further improved by adding tangible elements to the interface. In particular, we designed tangible interfaces for performing clipping-plane operations. Because of the dense nature of the data, we believe that adding a tangible clipping plane and an intersection image can help the user to better understand the complex data set. The computing platform and tangible interfaces are described to clarify the different design options. An experimental study is planned to quantitatively measure the added value of different aspects of the tangible interface.

INTRODUCTION

When talking about 3D interaction, people often think of 3D input devices, such as a 3D joystick, or 3D output devices, such as 3D stereo shutter glasses. However, 3D interaction should also be concerned about the activities that take place in the context of the 3D space that is being manipulated through these devices. The introduction of 3D interaction was driven by technological opportunities and by our desire to better exploit human familiarity with the 3D world that surrounds us daily. Interacting in 3D space has an intuitive feeling for a wide range of applications. In the early 1960s, Ivan Sutherland (1968) proposed his vision of using an immersive head-mounted-display-based computer system for 3D interaction. His work is generally recognized as the first 3D interface. Ever since, 3D interfaces and relevant interaction techniques have become increasingly interesting topics to study.

VR is the most popular approach towards 3D human-computer interfaces. Fred Brooks defines a VR experience as "any in which the user is effectively immersed in a responsive virtual world; this implies user dynamics that control the viewpoint" (Brooks, 1999, p.16). VR is an approach towards scientific visualization that makes multi-sensory 3D modeling of scientific data possible. While the emphasis is on visual representation, other senses, such as touch, can potentially complement and enhance what the scientist can visualize.

Although it is difficult to categorize all VR systems, this chapter separates them based on their display technology:

1. Projection-based VR systems, for example, CAVE (Cruz-Neira, 1993) or workbench (Kreuger, 1995).
2. Head-mounted display (HMD) VR systems (Sutherland, 1968).
3. Monitor-based desktop VR systems, for example, fish tank VR (Ware, 1993).

Visualization researchers increasingly use VR interfaces to build applications for domain scientists to display scientific data in 3D using a variety of visualization techniques (Hansen, 2004). However, there are currently few guidelines regarding which type of display system should be used, and even less evidence derived from qualitative and quantitative analysis. This can lead to the development of applications whose design may not use the most effective system to solve domain scientist's problems.

Tangible interfacing is another emerging interface perspective that is highly relevant for designing 3D interactions. One idea of tangible interfaces is that digital spaces have traditionally been manipulated with input devices, such as keyboard and mouse that were developed for traditional (2D) desktop activities. These input devices are ill suited to control and manipulate objects in 3D virtual worlds. Tangible interfaces are introduced to remove/decrease this discrepancy between input and output and are trying to open up new possibilities for interaction that more successfully blend the physical and digital worlds (Ullmer, 2001). Tangible user interfaces emphasize touch and physicality in both input and output. Often tangible user interfaces are coupled to the physical representation of actual objects, such as buildings in an urban planning application (Ishii, 2002), or wooden blocks for manipulation of online media (Ullmer, 1998).

In this chapter, we aim to better understand usability issues in 3D interaction through studying user performance within different 3D interfaces (including different VR systems and tangible interfaces) for four generic visualization tasks. The discussion of experimental work overlaps with an earlier publication on the use of different VR display systems for visualizing and manipulating volumes (Qi, 2006).

BACKGROUND

Related Work in VR Visualization

There has been a great deal of effort in the VR research community aimed at developing and integrating new devices and technologies to improve the usability of VR systems. Much work has investigated the usability and effectiveness of VR systems for simulating real-world scenarios. The study reported here attempts to validate the usefulness of three VR systems for a set of representative volume visualization tasks.

The case for stereo in scientific visualization is clear. Ware has shown that stereo viewing combined with motion parallax provided improved user performance in the 3D visualization of graphs, which argues for using VR rather than a traditional 2D projected display (Ware, 1996). A study by Arthur also demonstrated the advantages of a fish tank VR system for 3D tasks compared to 2D desktop images (Arthur, 1993). Of interest to us is which type of stereo VR system is most effective for scientific visualization of dense volume scalar fields.

The Effective Virtual Environments (EVE) group at UNC Chapel Hill has conducted presence, locomotion and re-directed walking studies within immersive HMD VR systems (Razzaque, 2001; Meehan, 2002). Immersive versus fish tank VR for searching and labeling has been studied by Demiralp (2003), who compared fish tank VR and CAVE displays for a visual search task. The results of their qualitative study showed that users preferred a fish tank display to the CAVE system for a scientific visualization application because of a perceived higher resolution, brightness, crispness and comfort of use. The results showed users perform an abstract visual search task significantly faster and more accurately in a fish tank environment, compared to the CAVE.

Navigation in HMD versus CAVE has been studied by Bowman (2002). He presented a preliminary experiment comparing human behavior and performance between a HMD and a four-sided spatially immersive display (SID). In particular, he studied users' preferences for real versus virtual turns in the virtual environment. The results indicated that participants have a significant preference for real turns in the HMD and for virtual turns in the SID. The experiment also found that females were more likely to choose real turns than males. This suggests that HMDs are an appropriate choice when users perform frequent turns and require spatial orientation.

Schulze et al. (2005) presented a user study comparing performance across multiple immersive environments for a counting task. They tested three VR displays: a CAVE-like environment, a single-wall display, and a desktop system (fish tank VR). Data they collected "led to four significant findings: (1) in the CAVE the participants preferred medium sized or large spheres over small spheres; (2) when only a few targets have to be marked, larger spheres were marked faster than smaller spheres; (3) large spheres are marked most accurately; and (4) performance for the wall display was not comparable to the fish tank VR display when the spheres were small. Additionally, occlusion and a larger field of view inhibited performance in the CAVE more than in the fish tank display when the task was dominated by visual search."

The scientific visualization community is continually developing better algorithms to represent data in a form suitable for comprehension. Traditional visualization schemes are entirely visually dependent. More and more VR systems for visualization applications incorporate haptic feedback. An early example of haptic representation of scientific data is found in the work of Brooks (1990). Users are assisted by a force reflective master manipulator during a complex molecular docking task. In this work, a force display is used to drive the system towards a local minimum and indicate tightness of fit. The nanoManipulator (nM) (Taylor, 1997) is a VR system that provides an improved, natural interface to scanning probe microscopy, including scanning tunneling microscopy and atomic force microscopy. The nM couples the microscope to a haptic VR interface that gives the scientist virtual telepresence on the surface, scaled by a factor of up to a million to one. The Visual Haptic Workbench (Brederson, 2000) is another testbed system for conducting research on the synergistic benefits of haptic displays using an integrated, semi-immersive virtual environment.

Several studies have measured the effects of a haptic display on human perception. Studies from Ernst have shown a clear influence of haptics on vision, demonstrating that vision does not necessarily completely capture haptics (Ernst, 2002). The human central nervous system seems to combine visual and haptic information in a fashion that is similar to a maximum-likelihood integrator. Visual dominance occurs only when the variance associated with visual estimation is lower than that associated with haptic estimation. Our study quantitatively investigates differences in user performance due to the presence or absence of haptic feedback for a visualization task.

Kosara (2003) suggested that user studies should be designed to evaluate visualization methods. This also applies to VR systems with visualization capabilities. Previous user studies have offered insight into the appropriate selection of VR systems for universal interaction and manipulation tasks such as rotation, navigation and sparse visual search. Our study extends this work to include several tasks specific to the visualization of dense volumetric data sets.

Related Work in Tangible Interfaces

Interaction in 3D space often requires a user with spatial reasoning and 3D perception skills. Researchers are also trying to tackle this 3D interaction problem from the perspective of interface design. Recent tangible interfaces that are based on more advanced tracking technologies can potentially improve the 3D interaction process, and several studies have already been undertaken to develop better interfaces for 3D interaction. Many of these studies have focused on more generic 3D manipulation tasks (Chen, 1988; Zhai, 1995; Hinckley, 1997).

The Passive Interface Props (PassProps) (Hinckley, 1994) was one of the first 3D interfaces to support continuous clipping interaction in 3D space. The PassProps was developed to allow surgeons to explore a patient's anatomy data by interactively generating cross-sections through the 3D data. The PassProps contains a head prop, a cutting-plane prop for creating intersections, and a pen-like prop for planning trajectories. The six degrees of freedom (DOF) that specify the position (i.e., translation and orientation) of each individual prop are tracked using (wired) magnetic trackers. Visual feedback of the user's actions is provided on a computer display in front of the user. The head prop is used to manipulate the orientation of the patient's anatomy. The rendering of the volumetric data on the screen follows the rotation of the head prop. The rendering is always positioned in the center of the screen, that is, it does not follow the translations of the head prop. The rendering scale (i.e., the zoom factor) is determined by the observer-to-object rendering distance, and is controlled by moving the head prop closer to or further away from the body. The user holds the cutting-plane prop relative to the head prop to specify the location and orientation of the slice through the 3D data. The generated intersection image is presented on the display, next to a (volume) rendering of the 3D model.

De Guzman et al. (2003) presented two tangible devices for navigating a slice through the human body. Interface A consisted of a 30-inch 2D model of a human body, together with a U-shaped fork at the end of an adjustable arm that could be rotated 180 degrees along the device's baseboard. Interface B consisted of a transparent 3D model of the human body and a free-moving hand-held fork. The fork in each case represented the intersection plane (window), and its position and orientation was used to generate an intersection image on a separate display.

The Cubic Mouse (CMouse) (Froehlich, 2000) was developed to support exploration of 3D geological data (seismic data) and car crash analysis data. The CMouse allows users to specify three orthogonal cutting planes and to perform so-called "chair cuts" through the data. The prop is a cube-shaped case with three perpendicular rods passing approximately through the centers of two parallel faces of the case. It is usually held in the non-dominant hand. The rods are used to control three orthogonal slices through the 3D data, that is, by pushing or pulling a rod, usually with the dominant hand, the corresponding intersection plane moves back and forth. The movement of a slice is hence constrained to the direction orthogonal to the slice. There is also a (wired) magnetic tracker embedded in the cube-shaped case. The tracked six DOF are used to translate and orient the data set in the virtual world, relative to the observer. The 3D data set and the orthogonal slices are visualized on a large stereo display in front of the user.

There are some limitations in the above systems that are likely to have an effect on their usability and user acceptance. First, because of the active tracking technology being used in these systems, the interaction elements need to be wired. Such wires obviously will have an effect on the freedom of movement, an effect that is seldom mentioned, let alone evaluated. Alternative techniques such as optical tracking allow for interaction elements that are passive and unwired, and are therefore likely to ameliorate this problem. Second, there is currently little insight into how different aspects of tangible interfaces, such as passive haptic feedback and enhanced perceptual feedback, assist users in their data analysis task in 3D space.

Scientific Problem

There are many medical applications that can benefit from using 3D interaction devices and techniques. For example, surgeons make use of 3D rendering and interaction to plan where to cut a patient because the body is 3D and the location of a tumor has a 3D location that is easier to understand. Our study on 3D interaction starts from the scientific problems asked by domain experts that are studying the structure of human lung mucus in both normal "wild-type" lungs and in the lungs of CF patients. This mucus is made up of a number of long polysaccharide molecules called mucins. It is known that there are a number of different types of mucin present in the mucus, and that the mucus is denser for CF patients than wild-type mucus. What is not known is how the different types of mucin are distributed in the mucus, and how particles can diffuse through it. The mucins may be uniformly distributed, or form distinct domains. There may be web-like superstructures formed by a subset of the mucins which contain clumps of other mucins. There may be large, small, or a variety of different sized water pockets surrounded by thin membranes. There may be continuous water paths within webs of mucins forming a lattice. Researchers are probing this by developing fluorescent dyes that attach differentially to the different mucin types, and by scanning the mucus with a confocal microscope to produce multiple 3D scalar fields, one for each dye. The resulting scalar fields in 3D are displayed to help them estimate sizes, distributions, and shapes of any resulting voids and structural elements. A virus, bacteria, or bacterial colony would traverse the mucus differently depending on its structure. The motion of such pathogens is of great interest to the study of CF, because lung infections are the source of many CF deaths. Researchers are probing this by placing small beads of various radii into the mucus and tracking the Brownian-driven motion of these beads over time to understand how they move through the mucus matrix. We wish to display the resulting motion paths in the presence of the above mesh structure to help our users correlate structure and density with bead motion paths.

AN EXPERIMENTAL STUDY OF VR SYSTEMS

We have designed different tools to solve the visualization and interaction problems described above with available VR technologies. In particular, we implemented three different VR systems for visualizing and interacting with a simulated data set that shares key properties with the real data. We designed a user study with this simulated data set to help determine which display and interaction system best supports the types of queries researchers are interested in without requiring our participants to be experts in CF. As described already, those experts are performing a diversity of tasks within dense volumetric scalar fields. Connectivity and relative density are of interest in addition to counting, shape, and size analysis. We aimed at providing similar tasks that were as generic as possible, so that the results of the user study could apply to other applications that explore dense 3D scalar fields looking for structure and pathways. We think that oil-field study and tumor segmentation might have similar needs for understanding complex dense data and for studying connectivity between portions of the data (oil reservoirs and blood vessels).

Figure 1. Head-mounted display system: (a) a user in the immersive HMD VR system; (b) head-mounted display with head sensor


(a) (b)

VR Systems for Visualization

All three systems display the volumetric data using the Visualization Toolkit (VTK), an open-source library that provides several different rendering algorithms (ray-casting, isosurface and 2D texture mapping) (Levoy, 1988; Schroeder, 2000). To enable real-time interaction, we chose Marching Cubes as the primary algorithm for rendering isosurfaces of the volumetric data. The standard structure of VTK does not provide a mechanism for integration with VR input devices, so we combined the VTK library with VRPN (Virtual Reality Peripheral Network) (Taylor, 2001) and UNC's Vlib (virtual-world library toolkit) to enable access to the visualization capabilities of VTK from our VR setups.

Immersive HMD VR System

The immersive VR system uses a V8 HMD from Virtual Research System. Each LCD provides a color VGA pixel resolution of 640 x 480 at a refresh rate of 60Hz. Head tracking is performed via a 3rdTech HiBall tracking system, a high-performance wide-area optical tracker that incorporates a six DOF sensor. The HMD/head tracking system consists of three main components. The outward-looking HiBall sensor is mounted on the back of the HMD (Figure1). The HiBall observes a subset of fixed-location infrared LEDs embedded in the ceiling. A tracking server coordinates communication and synchronization between the host computer and the HiBall and ceiling LEDs. Tracking data are transmitted through network switched Ethernet from the tracking server to a rendering computer via VRPN. We used a *DELL* Precision 530 (dual 2.8-GHz Xeon with 2GB RDRAM) and an NVidia *Quadro FX 1000* graphics card. The two VGA outputs from the graphics card are connected to the LCDs for each eye in the HMD via a video splitter to provide stereo-offset images. The working space for a user in this VR system is about 4.5 meters wide by 7 meters long by 4 meters tall (15 feet x 23 feet x 13 feet). A calibration procedure is used to calculate a precise transformation matrix between the sensor and the eyes. An additional hand sensor is also available for hand input, although it was not used during our experiments.

Fish-Tank VR

The second VR system is based on the concept of fish tank VR introduced by Colin Ware (1993). The computing platform of this VR system is identical to the HMD system with the following additional components:

1. A 17' CRT monitor with resolution of 1024 x 768 and a refresh rate of 100Hz to support stereo display, together with an infrared emitter and shutter stereo glasses from StereoGraphics Inc.
2. A PHANTOM Desktop™ haptic device for precise 6-DOF positioning and high fidelity 3-DOF force feedback output at 1kHz. In fish tank VR mode, the PHANTOM was used to rotate the volume around its center (additional operations were available during fish tank VR with haptics, as described below).
3. A DynaSight 3D optical tracker for measuring the 3D position of a target (reflective disc) attached to the front of the stereo glasses. When dynamic perspective is combined with stereoscopic viewing, a real-time 3D display appears that provides a virtual window into the computer-generated environment. Dynamic perspective eliminates the perceived image warping associated with static stereoscopic displays. An additional benefit of using the head to tune the perspective is that the hands are free to control the object being visualized, in our case with the PHANTOM.

The hardware components are organized to enable accurate and easy calibration. The tracker's control box is placed above the monitor on a metal plate supported by an arm (Figure 2). The arm's height guarantees continuous detection of the tracking and stereo signals. A cable between the infrared emitter for the stereo glasses and the control box for the head tracker synchronizes the devices. The real setup is shown in Figure 3.

Figure 2. A diagram of the fish tank VR system (with or without haptics)

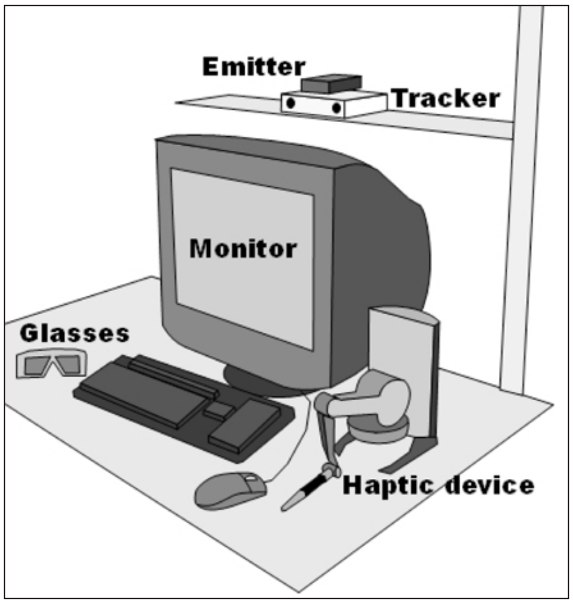

Figure 3. A snap shot of the fish tank VR system

Fish Tank with Haptics

Haptic visualization techniques have been developed for force feedback systems such as the PHANTOM. The fish tank VR with haptics prototype uses the same hardware setup as the fish tank VR system, except that the PHANTOM also provides force feedback, specifically a single point of haptic response, which is sufficient for our tasks. Although the stylus where force is applied is not visually located within the display volume (as compared to the Visual Haptic Workbench or the *ReachIn* systems), no users complained about the cognitive effort required to move the hand in one location while viewing another. An axis-aligned on-screen icon followed the stylus's motion in 3D, producing an effort similar to using a mouse to control the on-screen cursor. The haptic presentation of volumetric data employed different force models for different objects within the volume: viewers felt the outside surface of spheres and ellipsoids, but the inside of long curved tubes.

User Study

Our user study compares the three VR systems described above: VR, fish tank VR and fish tank VR with haptic feedback. Relative performance of these systems is measured over tasks involving the visualization of volumetric data.

Data and Task

Simulated volumetric data are generated to act as trials during our studies. A random number of two to four types of differently-shaped objects (sphere, ellipsoid, cylinder, and curved tube) are inserted with random positions (Figure 4). These objects may overlap with each other. The objects' properties (size, shape, and density) form experimental conditions that vary between trials. The bounding box of the volume is uniformly subdivided into eight sub-regions (a 2 x 2 x 2 array in the x, y, and z directions) within which object density may differ. Sub-regions are labeled with unique numbers (1 through 8) to enable participants to describe the paths of curved tubes within a volume and to indicate regions with the highest density.

There are always spheres and at least one curved tube within every volume. Trials may also contain ellipsoids, cylinders, and up to two additional curved tubes. Sphere sizes may vary between four possible radii ranging from six to twelve units. The density of objects within each sub-region is controlled to be sparse, medium, or dense. A single dense region (the "densest" region) exists within each volume. Sparse regions contain between 10%-60% of the number of objects in the dense region, while medium regions contain between 60%-90% of this number.

Participants are asked to complete fours tasks within each trial. Each task involves judging the properties of a specific object or of the overall volume, specifically:

Figure 4. An example trial from our experiment, showing a top-down view on a simulated volume with different experiment conditions like shape, size, density, and connectivity highlighted

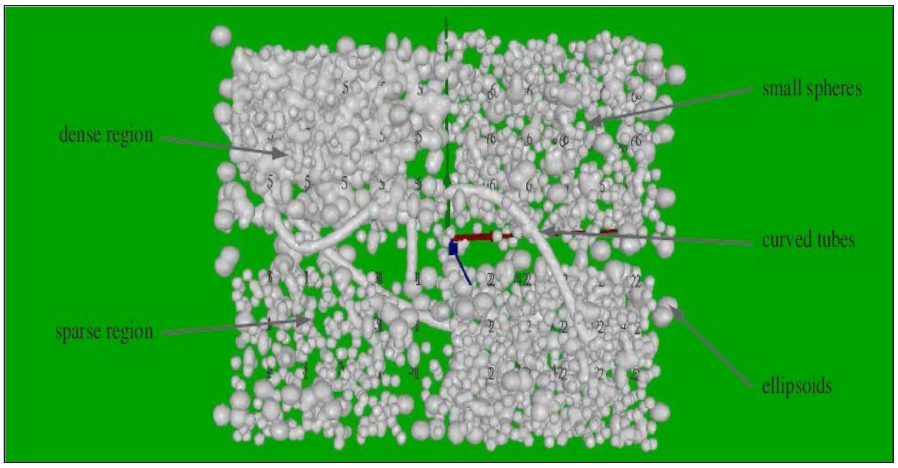

1. **Shape task:** Participants identify the number of differently-shaped objects within the volume and name the objects.
2. **Size task:** Participants report how many different sizes of spheres exist
3. **Density task:** Participants identify the densest sub-region in the volume
4. **Connectivity task:** Participants report how many curved tubes exist in the volume, and then determine which sub-region(s) the longest curved tube passes through. For example, Figure 4 shows two curved tubes

Participants are asked to give their answers as accurately as possible and to minimize response time. The size, density, and curve counting questions are presented in a multiple choice format. Participants are asked to describe the name of each kind of object for the shape question and all the sub-region numbers for the tube tracking question.

Experimental Procedure

A between-subject design was used, with VR system type as an independent factor: HMD, fish tank VR, and fish tank VR with haptics. Participants were randomly assigned into one of three groups. The HMD group wore the HMD and walked around within the tracked environment to observe the volumetric data as seen in Figure 5a. The fish tank group used the fish tank VR system and wore stereo shutter glasses to interact with volumetric data through the stylus of the PHANTOM as seen in Figure 5b. Although the stylus was tracked and displayed as an icon on the monitor, no force feedback was provided to this group. The haptics group added force feedback to the basic fish tank VR system.

Participants completed several steps during the experiment. As part of an initial interview session, they signed a consent form, answered basic demographic questions (age, gender, and occupation or major field of study), and identified their frequency of computer use and prior experience with any kind of VR system. A training session introduced the equipment and described the tasks to be performed. Next, the formal experiment session was conducted. Each experiment included 20 trials, with each trial containing a single volumetric data set. These twenty data sets were completely different from one another, and varied by object property (type,

Figure 5. Two views of volume data from an example experiment trial, (a) as seen in the HMD system, and (b) as seen in the fish tank and fish tank with haptics systems on the right.

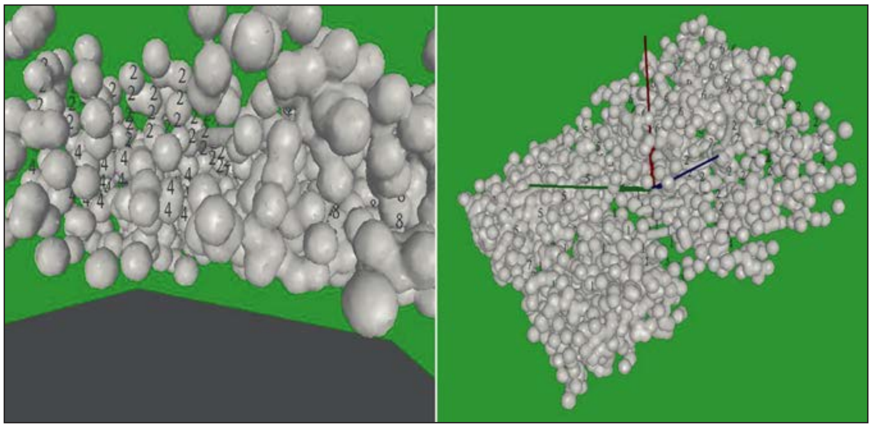
(a) (b)

size, position, and density). However, the same set of trials (20 data sets) in the same order was used for all three groups (HMD, fish tank, and fish tank with haptics).

Two dependent variables, the time taken to respond for each trial and the participant's accuracy (i.e., percentage correct) for each task were recorded. A short break was provided every half hour or whenever a participant asked for one. After completing the last trial in the formal experiment session, participants filled out a questionnaire describing their opinion about the system, any suggestions they had on how to improve the system, and so on (see Appendix). The study ended with a short debriefing during which the experimenter summarized the study goals. The participants were paid $9 for their participation.

Results

Forty participants volunteered for our experiment, 33 males and 7 females. The participants were randomly assigned into one of the three display system groups: 14 participants (12 males and 2 females) for the HMD group, 13 participants (11 males and 2 females) for the fish tank group, and 13 participants (10 males and 3 females) for the haptic group. The age of each participant and the frequency of computer use (on a scale from one to seven) were recorded before the experiment began. Average ages and frequencies of computer use were 23.2, 23, and 23.7, and 6.3, 6.0, and 5.6 for the HMD, fish tank and haptic groups, respectively. These data suggest we had similar ages and computer experience within each group.

Summary

Two kinds of measures of performance were derived for each trial a subject completed: response time *rt* and error rates P_e. A single *rt* value representing the total time in seconds needed to complete all four tasks was captured for each trial. We did not obtain the individual *rt*'s for each subtask since it was too difficult to record these separately. Four separate P_e values for the four tasks subjects completed were also obtained.

For the shape, size and density tasks, subjects' answers were coded as 1 for correct and 0 for incorrect. Then error rate P_e is defined as the proportion of wrong answers among all the answers. For the connectivity task, subjects' answers were coded in two observed parameters: the *false negative* and the *false positive*, as used in a Receiver Operating Characteristic curve (ROC).

For *rt* statistics, trials were divided by display system (HMD, fish tank, or fish tank with haptics). For P_e statistics, trials were divided by display system (HMD, fish tank, or fish tank with haptics) and task (shape, size, density, or connectivity). At times, more in-depth analyses on the data were performed when results obviously depended on other task parameters, such as in the case of counting sphere sizes, where performance obviously depended on the number of sizes present.

The shifts across conditions in average values of the logarithm of *rt* were studies using Analysis of Variance (*ANOVA*). A discussion on why analysis of lg(*rt*) should be preferred over analysis of *rt* itself can be found in a recent publication (Martens et al., 2007). The differences in error rates P_e were studied using chi-squared statistics.

In summary, the following significant differences in performance were identified:

1. The HMD group had the longest *rt*, followed by the fish tank with haptic groups. The fish tank group without haptics had the shortest *rt*.
2. For the shape task (counting the number of different shapes), the connectivity task (counting the number of curved tubes) and the density task (finding the densest subregion), the HMD group had higher P_e, compared to both fish tank groups (with and without haptics).
3. In the size task (counting the number of different sizes of spherical objects), none of the three groups is very accurate. The HMD

group made significantly more errors than both fish tank groups when only one size of sphere was present. This might be due to higher perspective distortion in the HMD case. When more than one size was present, subjects in all three groups tended to underestimate the number of different sizes.

4. In the connectivity task (identifying the sub-regions that the longest curved tube passes through), the HMD group produced more false negatives (missing the right sub-regions) and false positives (misjudging the wrong sub-regions) compared to both fish tank groups.

Detailed Analysis of Results

Performance Times

The response time *rt* needed to complete all four tasks during each trial was recorded during the formal experiment session. Subjects in the HMD group had significantly longer *rt* compared to the fish tank and the haptic groups. The *ANOVA* for the logarithm of *rt* was significant, $F(2, 165) = 40.058$; $p < 0.001$ (Figure 6). Post-hoc paired comparisons showed that the fish tank group was also significantly faster than the haptic group ($p < 0.001$). Overall, the HMD group spent 43% more time compared to the fish tank group, and the haptic group spent 23% more time compared to the fish tank group. Because of the high *rt* for the HMD group, we were forced to reduce the total number of trials for this system to 16. Because each trial tests all four tasks, this did not unbalance the experiment to favor certain conditions. Although subjects in the other two groups were able to finish all 20 trials within reasonable time, to maintain consistency we analyzed only the first 16 trials completed by each group.

Figure 6. Mean lg(rt) for the different experiment conditions, all results are divided by display system (HMD, fish tank, fish tank with haptics), error bars represent 95% confidence intervals.

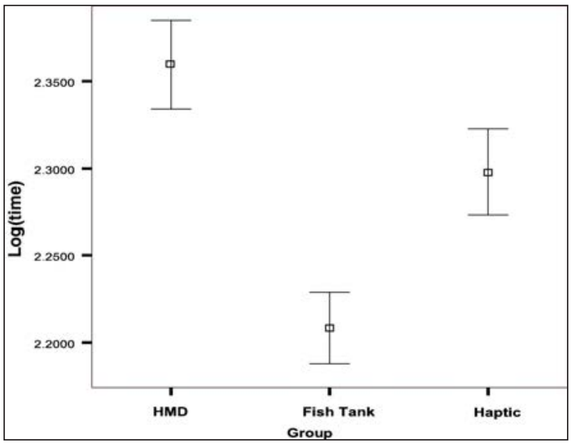

Accuracy in the Density Task

For the density task, the answers for every pairwise combination of groups are compared through a Chi-Squared test to find out whether or not there is an association between the error rates of finding the densest sub-region within a volume and the VR system used. The results are shown in Table 1 (The significant results where $p < 0.05$ are displayed in **boldface**, a convention that will also be used in the following tables) and can be summarized as follows.

The users in the HMD group produced significantly more errors then the users in both fish tank groups, while there was no significant difference between the two latter groups. In absolute terms, none of the three groups demonstrated very high accuracy, with P_e = 0.62, 0.38 and 0.43 for the

Table 1. Results of the Chi-squared analyses of overall error rate in the density task

	HMD	fish tank
fish tank	**$\chi^2 = 25.002$; $df = 1$; $p = 0.00 < 0.05$**	-
haptic	**$\chi^2 = 15.278$; $df = 1$; $p = 0.00 < 0.05$**	$\chi^2 = 1.206$; $df = 1$; $p = 0.272 > 0.05$

Figure 7. Mean Pe *values for the different experiment conditions, all results are divided by display system (HMD, fish tank, fish tank with haptics), error bars represent 95% confidence intervals: (a) mean* Pe *for the density task; (b) mean* Pe *for the shape task; (c) mean* Pe *for the size task; (d) mean* Pe *for counting the number of curved tubes in the connectivity task*

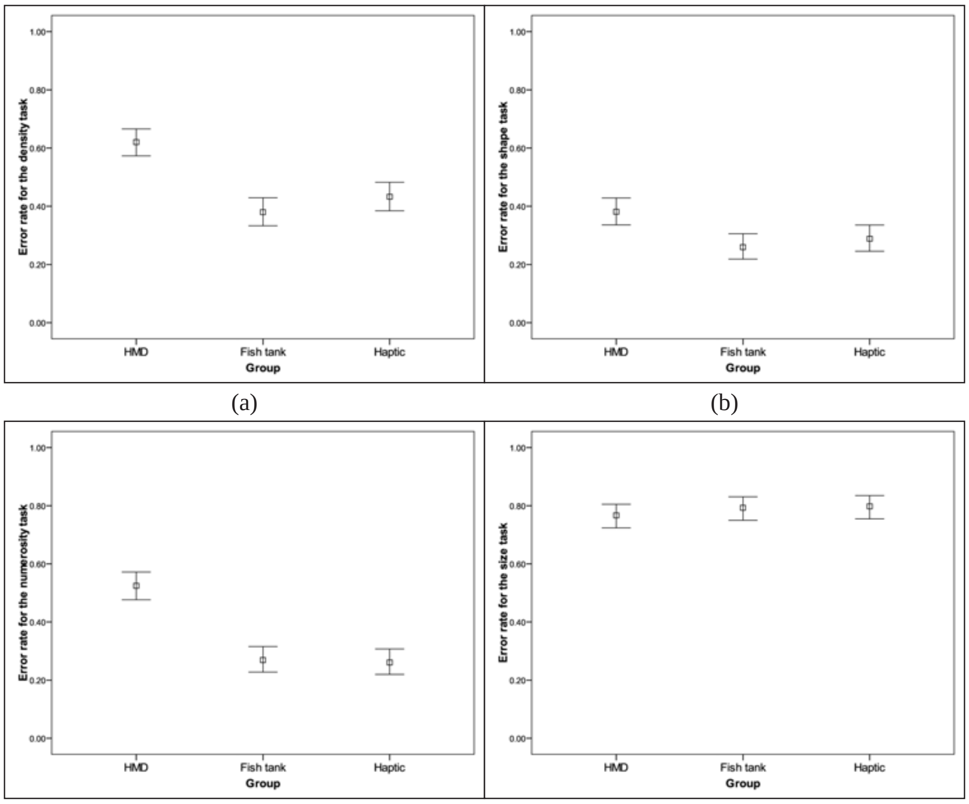

(a) (b)

(c) (d)

Table 2. Results of Chi-Square analyses of overall error rate in the shape task

	HMD	**fish tank**
fish tank	**$\chi^2 = 7.277$; $df = 1$; $p = 0.007 < 0.05$**	-
haptic	**$\chi^2 = 4.143$; $df = 1$; $p = 0.042 < 0.05$**	$\chi^2 = 0.435$; $df = 1$; $p = 0.510 > 0.05$

HMD, fish tank, and haptic groups, respectively (Figure 7a).

Accuracy in the Shape Task

The results of the Chi-Squared analysis for the shape task are shown in Table 2 and the conclusions as to the relative performance of all three systems are identical as in the case of the density task, that is, the HMD group is performing significantly worse than both fish tank groups.

In absolute terms, all three groups had reasonable accuracy, with $P_e = 0.38$, 0.26 and 0.29 for the HMD, fish tank, and haptic groups, respectively (Figure 7b). Further analysis indicated that the user performances of all three groups differed depending on the number of shapes present in a

Figure 8. Mean P_e with 95% confidence intervals for the shape task based on the number of shapes

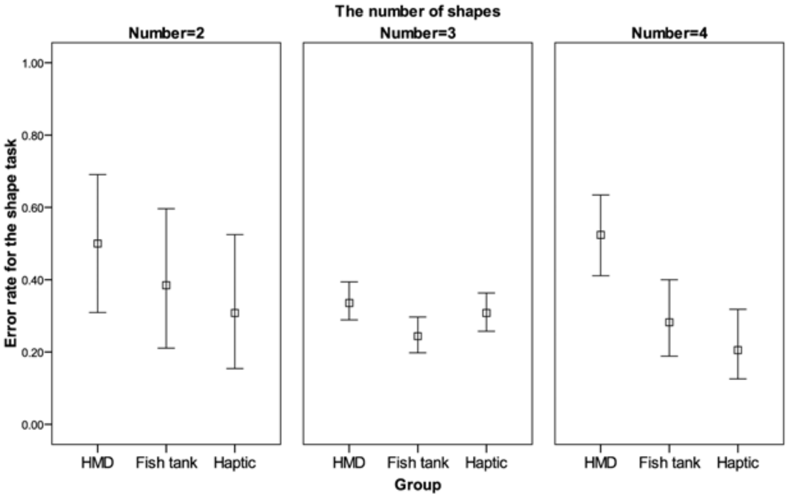

Table 3. Results of Chi-Squared analyses of overall error rate in the size task

	HMD	fish tank
fish tank	$\chi^2 = 0.438$; $df = 1$; $p = 0.508 > 0.05$	-
Haptic	$\chi^2 = 0.617$; $df = 1$; $p = 0.432 > 0.05$	$\chi^2 = 0.015$; $df = 1$; $p = 0.903 > 0.05$

displayed volume (Figure 8). When there are only two (sphere and curved tube) or all four kinds of shapes, the haptic group is more accurate than the other two groups (only the difference between HMD and haptic for the four shapes situation is statistically significant, though). The fish tank group is the most accurate when three kinds of shapes are presented (sphere, ellipsoid and curved tube or sphere, ellipsoid and cylinder or sphere, curved tube and cylinder). Regardless of the number of the shapes, the error rate is always the highest for the HMD group.

Accuracy in the Size Task

The results of the Chi-Squared analysis for the size task are shown in Table 3. No significant differences could be observed between the performances in the three groups.

In absolute terms, none of the three groups was accurate, with $P_e = 0.76$, 0.79 and 0.80 for the HMD, fish tank, and haptic groups, respectively. The error rates were all above 70%, although slightly fewer errors were made in the HMD group (Figure 7c). Further analysis based on the number of sizes showed that performance differences between systems varied (Figure 9).

When there is only one size or when there are two sizes of spheres, the haptic group is more accurate than the other two groups (although only the difference in case of one size between the HMD group and the haptics group is statistically significant). When there are three or four sphere sizes, the HMD group is somewhat more accurate than the other two groups (although this is only statistically significant in the case of four sizes). The only case where the error rate is below 50% (for all three groups) is when there is one size of sphere present. The chance of estimating the number of sizes correctly is even lower than guessing in case of the three- or four-sized condi-

Figure 9. Mean P_e for the size task based on the number of sizes

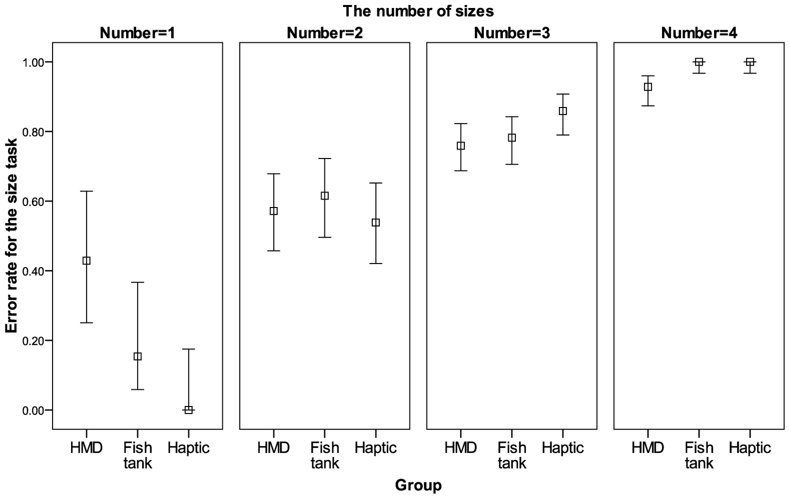

Table 4. Results of Chi-Squared analyses for the size task (including sign of estimation error)

	HMD	**fish tank**
fish tank	**$\chi^2 = 8.713$; $df = 2$; $p = 0.013 < 0.05$**	-
haptic	**$\chi^2 = 14.777$; $df = 2$; $p = 0.001 < 0.05$**	$\chi^2 = 1.347$; $df = 2$; $p = 0.51 > 0.05$

tion, which indicates that subjects significantly underestimate the number of different sizes in these latter cases.

In the previous analysis, user performance is based on the error rate, that is, the proportion of completely wrong answers among all the answers. When performing the size task, three cases can arise. The number of sizes can be estimated correctly, overestimated or underestimated. The resulting Chi-Squared analyses, based on three instead of two (right or wrong) categories, are reported in Table 4.

This more refined analysis reveals a significant difference between the HMD condition and the two fish tank conditions. In absolute terms, the proportion of underestimation was above 65% in all cases, although the proportion was lower during the HMD trials, with 0.68, 0.77 and 0.79 for the HMD, fish tank, and haptic groups, respectively. This reflects the fact that mistakes mainly originate from an underestimation of the number of different sizes for sphere objects, which is illustrated graphically in Figure 10.

Accuracy in the Connectivity Task

In the connectivity task, subjects answered two questions: the total number of curved tubes in a volume (numerosity question), and which subregions of the volume the longest tube passed through (spatial region question). The results of the accuracy analysis in case of the numerosity question are reported in Table 5. Similarly as in the above tasks, the HMD condition differed significantly in terms of accuracy from the two fish tank conditions, while the two latter conditions performed similarly.

In absolute terms, the accuracies were P_e = 0.52, 0.27 and 0.26 for the HMD, fish tank, and haptic groups, respectively (Figure 7d). Further analysis based on the task condition, that is, the number of curved tubes present, provides more insight into the performance differences (Figure

Figure 10. Proportions of answers in the size task, separated according to over- and underestimation

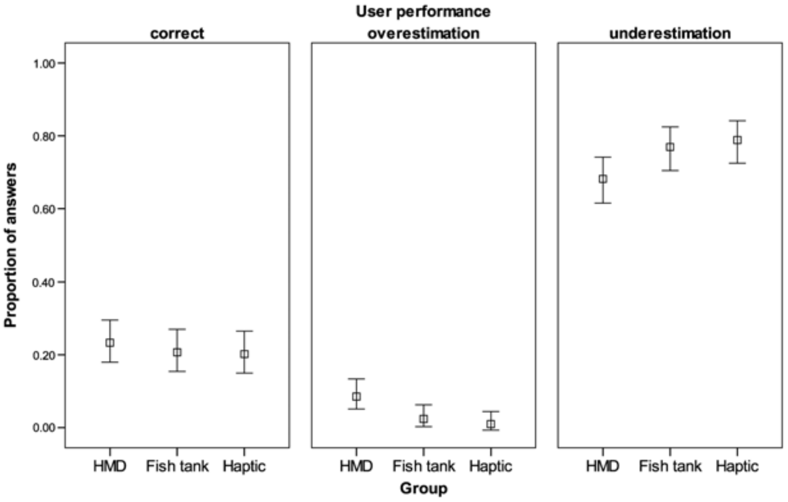

Table 5. Results of Chi-Squared analyses for counting the total number of curved tubes

	HMD	**fish tank**
fish tank	**$\chi^2 = 29.224$; $df = 1$; $p = 0.000 < 0.05$**	-
haptic	**$\chi^2 = 31.187$; $df = 1$; $p = 0.000 < 0.05$**	$\chi^2 = 0.037$; $df = 1$; $p = 0.847 > 0.05$

11). The users in the HMD group obviously have more problems distinguishing whether or not the perceived tubes are connected.

For the spatial region question, the answers of all three groups are analyzed using Receiver Operator Curve (ROC) statistics. ROC is a graphical plot of the *sensitivity* versus *1-specificity* for a binary classifier system. The ROC can also be represented by plotting the fraction of *true positive* (TP) versus the fraction of *false positive* (FP). In the context of this user study, there are four possible combinations of whether or not a sub-region is passed through by the longest curved tube in a trial, on the one hand, and a subject's answer based on his/her own judgment, on the other hand (see Table 6).

A *true positive* situation is that the longest curved tube in a trial passes through a sub-region, and a subject does recognize this fact correctly. A *false positive* situation is that the longest curved tube in a trial does not actually passes through one sub-region, but a subject thinks it does by mistake. The other two situations can be described similarly. From these situations, we can derive several statistics to describe the user performance. The fraction of *false negative* P_{FN} (also known as the chance of missing P_M) is defined as $P_{FN} = FN / NP$, while the fraction of *false positive* P_{FP} (also known as the fraction of false alarm P_F) is defined as $P_{FP} = FP/NE$. Both probabilities can be analyzed as a function of the experimental conditions using Chi-Squared statistics.

The analyses for the false negatives are summarized in Table 7. The fish tank system performs significantly different from the other two systems. In absolute term, the fish tank group is more accurate than the other two groups, with P_{FN} equal to 0.39, 0.31 and 0.39 for the HMD, fish tank, and haptic groups, respectively. This is also reflected in Figure 12.

The analyses for the false positives are summarized in Table 8. In this case, only the differ-

Figure 11. Mean P_e for counting during the connectivity task divided according to the number of curved tubes

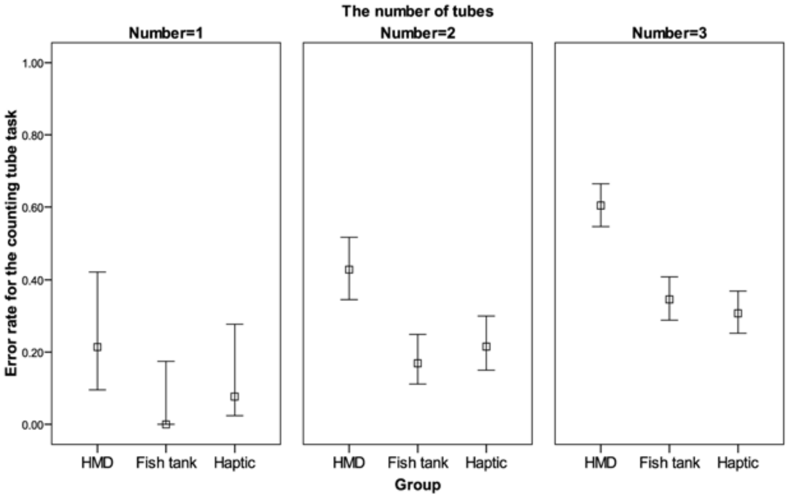

Table 6. Four situations in judging whether the longest tube passes through a sub-region

	pass (real situation)	**no pass (real situation)**
pass (subjects' answer)	*true positive* (TP)	*false positive* (FP)
no pass (subjects' answer)	*false negative* (FN)	*true negative* (TN)
sum	*number_of_pass* (NP=TP+FN)	*number_of_empty* (NE=FP+TN)

Table 7. Results of Chi-Squared analyses for P_{FN}

	HMD	**fish tank**
fish tank	**$\chi^2 = 10.896$; $df = 1$; $p = 0.001 < 0.05$**	-
haptic	$\chi^2 = 0.021$; $df = 1$; $p = 0.884 > 0.05$	**$\chi^2 = 11.474$; $df = 1$; $p = 0.001 < 0.05$**

ence between the HMD condition and the fish tank condition without haptics is shown to be statistically significant. In absolute terms, the HMD group is less accurate then the two other groups, with P_{FP} equal to 0.20, 0.17 and 0.17 for the HMD, fish tank, and haptic groups, respectively. This is reflected in Figure 12.

The overall performance can be quantified by attributing costs to both the false negatives P_{FN} (misses) and the false positive P_{FP} (false alarm). In summary, the fish tank (without haptics) group has the best performance. The HMD group is least accurate in both finding the sub-regions the longest curved tube passes through and ignoring the regions that the longest curved tube does not pass through. The haptic group has intermediate performance, with almost the same frequency of false negatives as the HMD group but slightly lower frequency of false positives.

Figure 12. Mean P_{FN} and P_{FP} values for the different experimental conditions, all results are divided by display system (HMD, fish tank, fish tank with haptics), error bars represent 95% confidence intervals: P_{FN} and P_{FP} values for locating the longest curved tube during the connectivity task.

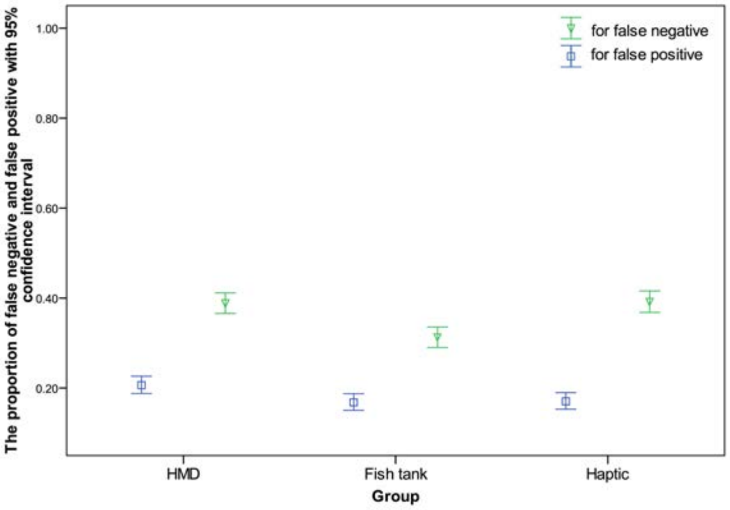

Table 8. Results of Chi-Squared analyses for P_{FP}

	HMD	**fish tank**
fish tank	**$\chi^2 = 4.123$; $df = 1$; $p = 0.042 < 0.05$**	-
haptic	$\chi^2 = 3.604$; $df = 1$; $p = 0.058 > 0.05$	$\chi^2 = 0.017$; $df = 1$; $p = 0.896 > 0.05$

Interpretation of Results

The time needed to complete tasks in the HMD system was significantly longer, compared to both the fish tank and the fish tank with haptics systems. One explanation is that the HMD system requires participants to walk around within the tracking space, which takes more time to explore compared to moving hands and head in the fish tank systems. Another critical issue was the reported inability of HMD participants to remember where they had previously seen target items within the volume because of the high density of data sets. They would often have to re-search the volume for objects they had previously located, but had "lost" as they walked into a different region. Finally, participants may simply be more familiar with a standard desktop system.

The fish tank group was also significantly faster than the haptic group. When touch was available, participants often spent more time "feeling" inside the volume to confirm their decisions, even when a correct answer could be derived from visual evidence alone.

A curve of the time spent on each trial indicates a similar learning effect for all three groups. Time decreases as the participants complete more tasks (Figure 13). The first five trials show the strongest learning tendency. After the first five trials, time spent on each trial still varies mainly due to the different difficulty of each trial, which is caused mainly by the density of the data set for that trial. The learning effect did not affect our ability to draw conclusions because the three groups shared the same learning pattern. For users who have become familiar with the task and equipment,

Figure 13. A time curve for each VR system, in the order that participants completed the trials

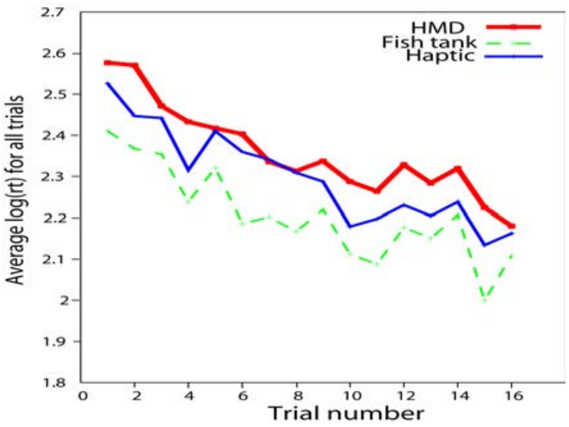

we believe performance will stabilize to times similar to those seen during the later trials.

The HMD group was significantly less accurate than the fish tank and haptic groups in the shape task. Error results showed that participants from all three groups found it relatively easy to identify the sphere object. The HMD group made more mistakes in identifying cylinders than the other two groups. It is difficult to judge the number of shapes for the HMD group when all four shapes exist. The accuracy of the haptic group consistently increases as the number of shapes increases, and gets the best performance among three groups when all four shapes of objects are present. It indicates that touch does help the participants for the shape tasks. Finally, participants from all three groups sometimes misjudged the curved tube as a cylinder. This was also mentioned in the post-experiment feedback from the participants.

Although there were no significant differences in accuracy during the size task, absolute performance was poor across all three groups. Not surprisingly, when there was only one size of sphere, the responses were quite accurate (P_e = 0.43, 0.15, and 0.0 for HMD, fish tank, and haptic, respectively). When two sizes of sphere with a large difference in radii were presented, the participants from all groups also did well. The haptic group performed best in the cases of one or two sizes. However, when the radius difference between the two spheres was small, or when there were three or four different sizes of sphere, all participants had difficulty determining how many different sizes there were (P_e = 0.76, 0.78, and 0.96 with three sizes of sphere, and P_e = 0.93, 1.0, and 1.0 with four sizes of sphere for the HMD, fish tank, and haptic groups, respectively). The average accuracy of each group is even lower than the probability of guessing in the three or four sizes situation. The participants with the HMD system produced the best performance when there are more than two sizes. This suggests that: (1) some of the radii differences were too small to be distinguished reliably; (2) touch did not help much in distinguishing such small differences.

For the density task, the HMD participants were significantly less accurate than the fish tank and haptic participants. None of the three groups had high accuracies, however. The reason might be both the characteristics of high-density and slight-density differences between adjacent sub-regions. There was no significant difference between the fish tank and haptic groups, implying that haptic feedback did, again, not assist in identifying spatial regions with different densities of objects, particularly in high-density situations. The errors for all three groups were spread out across the trials, and showed no learning effects. This suggests that identifying regions of varying density (especially small differences) within a 3D volume is a difficult task, that none of these three display systems fully supports.

In the connectivity task, participants are asked to count the number of curved tubes in the volume, then locate the longest curved tube and identify which sub-regions of the volume it passed through. For the numerosity question, the HMD participants were significantly less accurate than the fish tank and haptic participants. Identifying that only one tube exists does not present a challenge for any of the participants. When more than one tube is present, the average accuracies for all three groups decrease as the number of tubes increases. However, the ability to "feel" along the inside of the tubes helped the haptic group provide slightly more accurate counts of the number of unique tubes when there are three or more. It indicates that touch can be useful in such a situation. The lack of overview of the volume for the HMD group and the absence of clear complete view of the path of every tube for all three groups create the major difficulty for all participants to judge whether different segments they saw belonged to the same tube or not.

For the spatial region question in the connectivity task, the HMD group was significantly less accurate than the fish tank and haptic groups. The fish tank group was slightly more accurate than the haptic group. Task analysis for this question revealed that participants first had to identify which tube was the longest one by comparing the lengths of all tubes, before continuing with determining which sub-regions this target tube passed through. If the wrong tube was identified as longest, the answer on the second question was obviously also wrong (apart from the cases where the wrong tube crossed the same sub-regions as the longest one). In the HMD system, participants often misjudged a tube to be the longest one. For the fish tank and haptic systems, when the length differences among the tubes were large, haptic feedback helped participants locate the longest tube by touch. They could then correctly identify the sub-regions containing the tube. When the length differences were small, however, the haptic system provided insufficient assistance. This explains the slightly different error rates between the fish tank and the haptic systems. Our results match the findings of Ernst and Banks (2002): when visual and haptic feedbacks are present and haptic feedback can add a definite assistance for a task or judgment, it will be used. Otherwise, visual feedback is still the dominant sensory input.

In addition to statistical results, a number of interesting anecdotal findings were made, pointing to: (1) the desire for an overview display in the HMD system; (2) the desire for immersion in the fish tank VR systems; (3) fatigue in the HMD system; and (4) the preference for including touch in the haptic system.

Several HMD users spontaneously suggested adding the ability to see a high-level overview (which might be provided through a button press, Mine's head-butt zoom, or a worlds-in-miniature interface). One casual user was tall enough that he stood above the data, enabling him to get an overview in the HMD system, which he reported to be useful. This matches our later analysis as well as issues related to the effects of memory on participants' results. Some participants in the fish tank and haptic groups wanted to zoom in and see the volume from the inside (some tried to do this by moving their head near the screen). We concluded that both overview and immersion are helpful for performing our tasks. Anecdotal and formal results indicate that a system designed for the study of dense volumes should include both capabilities.

Most participants said that the HMD and haptic systems were "cool" or "neat" upon initial exposure. Several participants mentioned without being asked that they liked the HMD VR system or the haptic system. However, participants in the HMD group requested more breaks after five trials and sometimes asked "How many trials do I still have?" after around ten trials, indicating heavy workload and a dissatisfaction with the system. We believe this is due to physical or mental fatigue. The increased number of breaks requested did not happen in the fish tank or the haptic cases.

Subjective Results

Subjective measurements were obtained through analysis of the post-experiment questionnaires (see Appendix). Most questions used a standard seven point rating scale (some used a five-point rating scale). The answers indicated that overall, participants preferred the haptic and HMD VR systems due to perceived ease of use, presence, and immersion. We summarize our findings over the following categories of questions we asked.

Perception of the VR Systems

The first category of questions addressed perception properties and characteristics of VR systems, including immersion, presence, depth cues, and spatial relationships. For the question: "the extent that you felt you were within a virtual environment" the HMD system ranked significantly higher than the fish tank with haptics systems, F(2, 37) = 5.481, *p* = 0.008, with a post-hoc comparison between HMD and haptic of *p* = 0.006, and absolute rankings of 6.0, 5.4, and 4.4 for HMD, fish tank, and haptic, respectively (Figure 14a). There was also a significant difference on the question: "the extent you had a sense of acting in the virtual space, rather than operating something from outside." The HMD system ranked significantly higher than the other two systems, F(2, 37) = 15.666, *p* = 0.001, with scores of 5.9, 3.1, and 4.4 for HMD, fish tank, and haptic, respectively (Figure 14b). Further post-hoc comparison showed the fish tank with

Figure 14. Mean values for the different questions about the perception of VR systems, all results are divided by display system (HMD, fish tank, fish tank with haptics), error bars represent 95% confidence interval: (a) mean rank for the presence question; (b) mean rank for the question of acting inside VR space; (c) mean rank for the question of VR surrounding the subject; (d) mean rank for the immersion question.

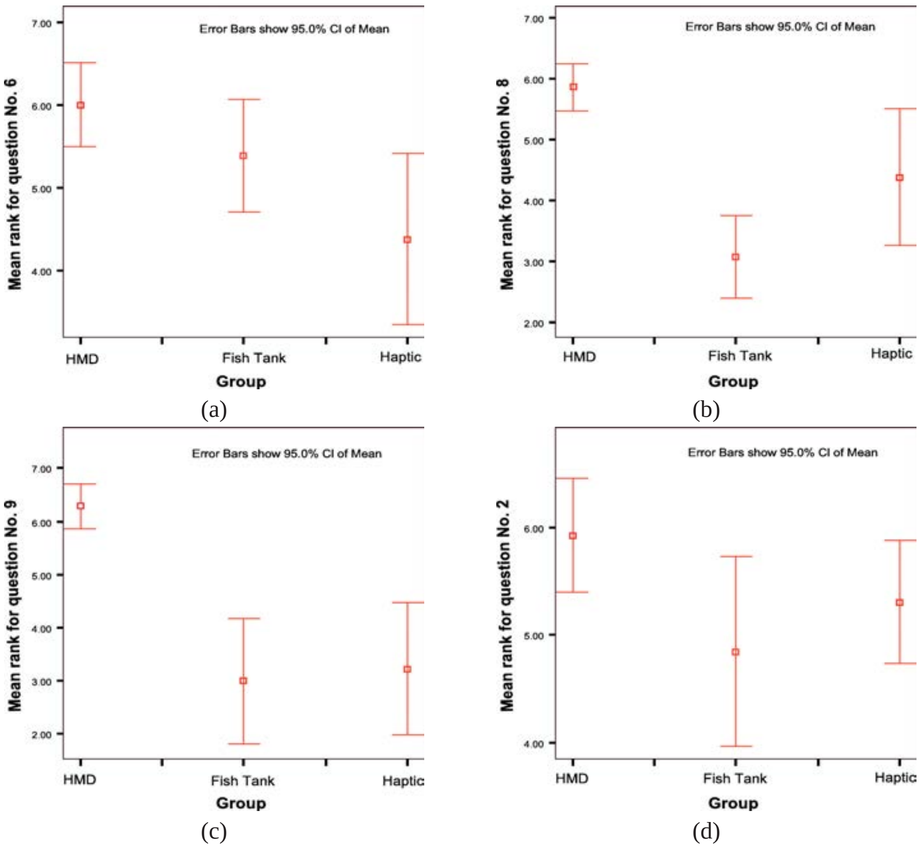

(a) (b)

(c) (d)

haptics system ranked significantly higher than fish tank alone due to the existence of touch (p = 0.03), indicating that haptic feedback does add an inside-out property to a fish tank display. For the question: "the extent you felt that the virtual environment surrounded you", the HMD group again ranked higher than the other two groups, F(2, 37) = 16.464, p = 0.001, with scores of 6.3, 3.0, and 3.2 for HMD, fish tank, and haptics, respectively (Figure 14c). This suggests that HMD participants felt more strongly that they were acting within a virtual environment. We found no notable statistics differences on the questions: "a sense of being there," "a sense of immersion," "difficulty of understanding the spatial relationships," "the quality of multiple view points," or "the quality of depth cues," although the HMD system did rank slightly higher in absolute terms in the immersion (Figure 14d), presence, multiple viewpoints and depth cues questions.

Usability of the VR Systems

The ease of learning and using a VR system is the main focus of this category of questions. Answers to the question: "how much consistency did you experience in the VR system compared with a real world experience" were similar for participants from each group, indicating the act of moving from place to place was judged to be relatively natural and easy. There were no obvious differences on the question about system delay, although HMD participants reported a slightly shorter perceived delay. No participant from any group complained about the resolution, frame rate or delay; these parameters did not seem to bother them. The haptic system ranked higher than the other two systems for identifying the shape and location of individual objects, and the shape of the global topology. Although participants from all three groups felt their system was easy to use,

Figure 15. Mean values for the different questions about usability issues of VR systems, all results are divided by display system (HMD, fish tank, fish tank with haptics), error bars represent 95% confidence interval: (a) mean rank for the level of demand on the participants' memory; (b) mean rank for the level of confidence in the answers.

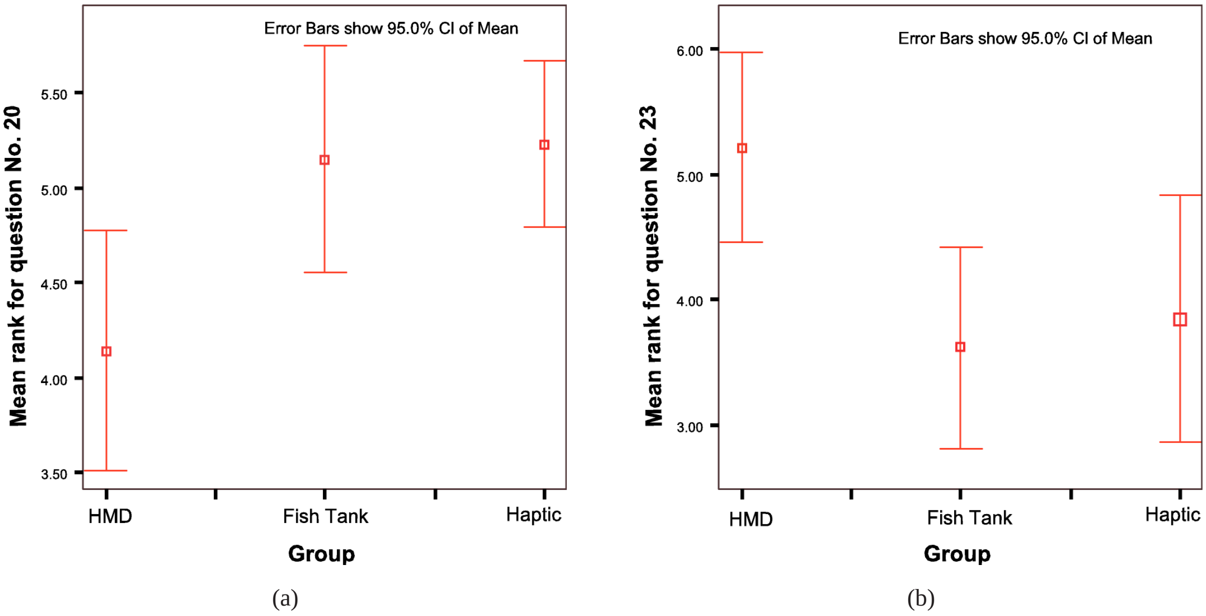

the HMD group ranked highest for the perceived difficulty in carrying out their tasks. Moreover, HMD participants reported a significantly higher demand for memorizing than the other two groups, F(2, 37) = 5.017, *p* = 0.012, with scores of 5.2, 3.6, and 3.8 for HMD, fish tank, and haptics, respectively (Figure 15a). Finally, HMD participants were less confident about the accuracy of their answers, F(2, 37) = 5.521, *p* = 0.008, with scores of 4.1, 5.2, and 5.2 for HMD, fish tank, and haptics, respectively (Figure 15b).

The "Added Value" of Haptics

The use of haptics requires participants to employ multiple sensory modalities to perform tasks. Most participants in the haptic group were excited about the additional functionality, and claimed that haptic feedback did help in some way. Participants in the haptics group were asked four questions that related to their experiences:

- Consistency of the information from multiple senses
- Ease of searching within the virtual environment through touch
- The effects of multimodal sensory on understanding the space
- The effects of multimodal sensory on understanding the structure of the data set

The first two questions used a standard seven-point rating scale. The last two questions used a standard five-point rating scale. Eighty percent of the participants from the haptic group thought the visual and haptic information was consistent, and that searching the virtual environment through touch was easy. Seventy-five percent of the participants thought touch helped them better understand the space, and 80% thought it helped understand global structure. Participants reported that haptics was especially helpful for the connectivity questions: "How many curved tubes are there?" and "Please name all the sub-regions the longest tube crosses" since the tubes are hidden behind other objects.

TANGIBLE INTERFACE FOR CLIPPING PLANE

Spatial reasoning and 3D perception skills are very important for interacting with volumetric data in 3D space. Currently, the dominant interface for 3D manipulation with volumetric data is the computer desktop with a graphical user interface that is controlled by a mouse and keyboard. As described earlier, recent developments within 3D interfaces that add tangible elements to the interface have the potential of improving the 3D interaction process for the purpose of data analysis. However, previous studies have focused on more generic 3D manipulation tasks, such as selection, positioning, and so forth (Chen, 1988; Zhai, 1995; Hinckley, 1997).

The user study in the previous section has indicated that an immersive VR environment without overview capability does not help users with most of the selected tasks (identifying data structure and properties), which was surprising to us. On the other hand, the dense nature of our data sets may explain the low task performances observed. The desktop VR environment helped users achieve better performance (than in immersive VR) in terms of accuracy and time, but absolute performance is still not really acceptable. Naturally, we ask whether there are other means that could further improve the user performance within such an environment. In this section, we investigate whether the inclusion of a tangible interface can help users performing those tasks. We focus on a more in-depth investigation into one specific interface aspect, that is, the positioning of a clipping plane within volume-rendered data. We propose the design of tangible interface prototypes for a clipping plane that employ wireless vision-based tracking (Mulder, 2002; van Liere, 2003). Such a tangible clipping plane can assist a user in exploring the inside of a dense volumetric data set through creating 3D and two-dimensional (2D) intersection images. By varying the design, these prototypes allow us to study and compare different user interface strategies for performing the clipping plane interaction task. A user evaluation is being planned with these prototypes for measuring their effectiveness.

Design Practice of Tangible Interfaces

Positioning of a clipping plane is a common but complex operation in volume visualization for data analysis. Such a plane cuts through a 3D data set in order to explore its interior structure. The common method of controlling the 5 (or 6) DOF of the virtual clipping plane, that is, its position (3 DOF) and orientation (2 DOF in case of a plane, or 3 DOF in case of a window), is by means of a 2 DOF control device such as a mouse. In order to accomplish this, the positioning task needs to be decomposed in at least three subtasks that require at most two DOF at a time. Despite the fact that such a 2D interface (in principle) allows users to perform the task, it is often difficult to obtain enough awareness of the spatial relationships to manipulate the data efficiently. This is due to the fact that the 2D interaction is unrelated to the natural interaction process in 3D space. We therefore propose alternative interface designs for the clipping plane task that make use of 3D (tangible) interaction devices. Our goal is to combine current knowledge and understanding of 3D user interface design (Bowman, 2004) with new technical possibilities to create clipping interfaces for more demanding and realistic tasks. The principles that we adhere to in the design of these interfaces are the following:

1. **Easy to use:** The interface should not distract the user from the actual clipping task.
2. **Easy to learn:** The interaction should be natural and intuitive, requiring little explanation and training.
3. **Adequate perceptual feedback:** The interface should provide (passive) tactile and visual cues that assist in the interaction.
4. **Low-cost setup:** The interface and system should be created using off-the-shelf and inexpensive technology.
5. **Real-time interaction:** The interface should work in real time while still generating volume rendering of realistic high-resolution data.

In order to achieve real-time performance, the hardware setup is organized around two *DELL* graphics workstations with different interface components.

The first workstation is mainly used for tracking the tangible interface and consists of the following components:

- One *DELL* workstation Precision 530 (Pentium IV, 2.4 GHz, 512 MB RAM with ATI *FireGL 4* graphics card coupled to an infrared emitter from StereoGraphics Inc.
- Two analog Leutron Vision LV-8500 progressive scan CCD cameras (720x576 pixels, 50Hz frame rate) with COSMICAR/PENTAX lenses with a focal length of 12mm and infrared transparent filters (that block visible light); these cameras are connected to two synchronized Leutron Vision PictPort H4D frame grabbers.
- A 15’ LCD display from *DELL*.

The second workstation includes the following components:

- One *DELL* workstation Precision 670 (Intel Xeon, Dual CPU, 3.2 GHz, 2.0 GB RAM with NVidia *Quadro FX 4500* graphics card)
- A 14’ CRT monitor from *DELL* with a vertical refresh rate up to 120Hz, so that stereoscopic images can be viewed with the help of active liquid crystal shutter glasses (CrystalEyes 3)

- A 15’ LCD display from *DELL*, used for showing intersection images

The separation of tracking and 3D rendering across two different machines enables us to achieve better performance for 3D interaction with a large data set. A wooden chassis has been constructed for integrating the different components and creating a workspace for the users. The two infrared cameras are mounted on the upper layer of the wooden chassis (Figure 16). A silver mirror mounted on a wooden slab is hung in front of the chassis at an angle of 45 degrees to reflect an image of the user’s hands with the tangible devices to the cameras. The use of the wooden cabinet with the cameras makes the system set-up stable and allows for easy transportation. In the current prototype there are no provisions for tracking the user’s head (which might be useful for also providing motion parallax feedback in the displayed image).

Due to the hardware difference in graphics cards, the volume rendering algorithm used in previous study has been modified to provide direct volume rendering. The rendering engine uses hardware-supported 3D texture mapping in OpenGL, rather than a Marching Cubes algorithm. The algorithm to render the non-polygonal isosurfaces in the data is based on the approach presented by Westermann and Ertl (1998). In a pre-processing step, the gradient vector is computed for each voxel of data set using the central differences method. The three components of the normalized gradient vector together with the original scalar value of the data set are stored as RGBA quadruplets in a 3D texture.

Figure 16. The diagram of the system setup

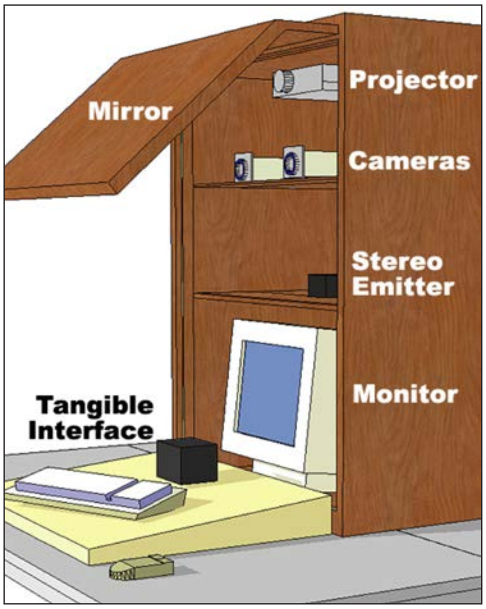

The vector components must be normalized, scaled and biased to adjust their signed range [-1; 1] to the unsigned range [0; 1] of the color components. The alpha test discards incoming fragments conditional on the outcome of a comparison of the incoming alpha value with a user-specified reference value. In our case the alpha channel contains the scalar intensity value and the alpha test is used to discard all fragments that do not belong to the isosurface specified by the reference alpha value. The setup for the OpenGL alpha test is:

```
glDisable(GL_BLEND);
// Enable Alpha Test for isosurface
glEnable(GL_ALPHA_TEST);
glAlphaFunc(GL_EQUAL, fIsoValue);
```

Within the prototype interfaces, the user is given a wooden cube that can be tracked by the system. This cube can be rotated to control the orientation of the virtual cube, with its associated volumetric data, and moved towards or away from the user’s body to control the zoom factor (see Figure 17). The image on the display, which is the result of volume rendering for a fixed camera position that is optimized according to the observer’s actual viewpoint (the size of objects observed is comfortable to users), changes in accordance with the movement of the tangible cube. The cube can also be placed on a small (physical) pedestal in case the user prefers to perform clipping operations

Figure 17. 3D manipulation of volumetric data with tangible cube

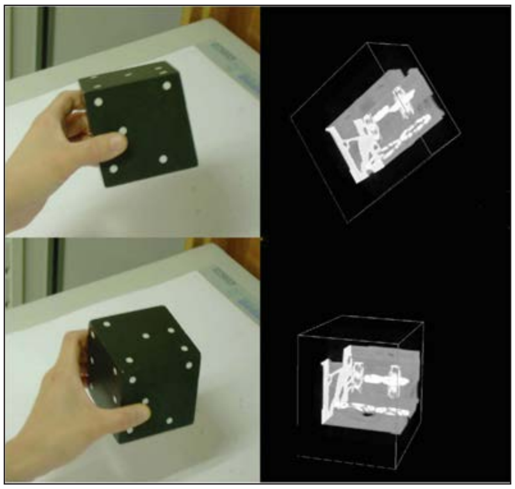

with one hand. Although the data set remains in a fixed position in such a case, its orientation can still be varied discretely in a very simple way, that is, by changing the side of the cube that is resting on the pedestal.

Together with the cube, a square-shaped metal frame is used to control the clipping plane. Five infrared-reflecting stripes on three of its sides form a unique planar pattern as shown in Figure 19. The six DOF of the frame (three DOF for position and three DOF for rotation) are monitored continuously by the vision-based tracking algorithm. The appearance of this device makes its purpose very obvious. While using this prototype, the user positions the cube with his non-dominant hand, and grasps the frame with his dominant hand on the side that has no dots on it. The physical cube intersects with the physical frame in a way that agrees one-to-one with the intersection of their virtual counterparts on the screen. As a result, even though all six DOF are enabled when moving the plane, the interface does not seem difficult to control.

The clipping plane interface can operate in two modes: the slice mode and the opaque clipping mode. In slice mode only the planar intersection image is displayed in 3D space, as shown in Figure 18a. In the opaque clipping mode, the part of the volume data that is in front of the clipping plane is made transparent, as shown in Figure 18b.

Tangible Devices for 3D Interaction

Considering ease of use, we pursue a further design by adding different handles to the frame. Three different types of handles have been designed in

Figure 18. (a) the slice mode for the 3D intersection interaction; (b) the opaque clipping mode for the 3D intersection interaction

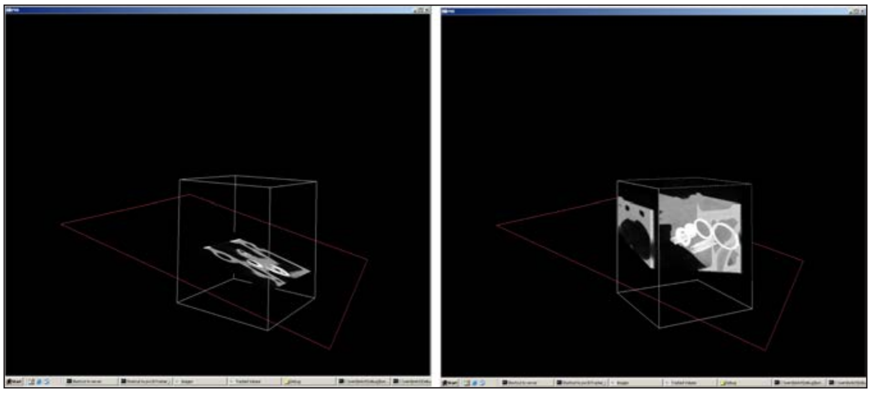

(a) (b)

order to find an effective solution. First, paper mockups of three handles are constructed to give an impression of the final look of the physical objects. Three wooden plane frames with different handles are made for further comparison as Figure 19a shows.

We have asked several colleagues from our department to try out the different prototypes and give informal feedback. The tangible clipping interfaces have received many positive feedbacks. Further survey among the users indicated that there is no difference in the preference of the shape of the handle, which means it will not be a major factor within future user experiments. The final design of the clipping plane frame is shown in Figure 19b.

With these interfaces, the next logical step is to undertake a more structured and formalized experiment with the different prototypes. We adopt the tasks performed in the previous section as a starting point since the systems in the previous experiment did not have a constantly satisfying user performance. More specifically, the experimental goal is to ask users to cut through 3D simulation data with a clipping plane in different forms and to answer the same questions regarding data properties as were asked in the previous experiment. By repeating that experiment with new interface prototypes, we plan to establish the relationship between our tangible prototypes and the absolute performance (time and accuracy) for a desktop VR environment. We can also verify whether or not the observed preference for the presence of a tangible clipping plane and a 2D intersection image holds in terms of user performance for our visualization tasks. Moreover, we can study the effect of two-handedness and of the form factor of the interaction device. In addition, we want to see how performance relates to the spatial ability of our individual subjects.

FUTURE WORK

Our planned user study will compare five kinds of setups within one baseline system: non-immersive VR, with fixed virtual clipping or a tangible clipping frame and with/without a 2D intersection image. The relative performance of these configurations will be compared for the four generic tasks that were applied in the earlier experiment. The rendering paradigms are only tested in their most common configuration: outside-in for non-immersive VR. A between-subject design is planned, with interface type as an independent factor:

Figure 19. (a) the wooden model of three handles; (b) the final design of the plane-like tangible interface for virtual clipping plane

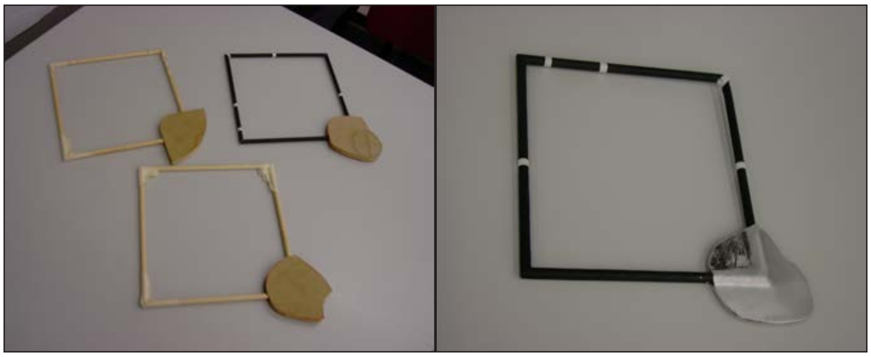

(a) (b)

- **Condition 1 (abbreviated as cube or C):** Baseline system is a non-immersive VR system with a tangible cube to orient the data set. The visual feedback is only a 3D image of the data set (Figure 20);
- **Condition 2 (abbreviated as fixed-plane or CF):** Baseline system with a fixed virtual clipping plane. A user can manipulate the cube and cut through the data with a fixed clipping plane (Figure 21a);
- **Condition 3 (abbreviated as tangible-frame or CT):** Baseline system with a tangible frame in the shape of a clipping plane. The movement of the visual clipping plane corresponds to the physical plane-shaped object. The visual feedback is a 3D representation of the data set and the virtual clipping plane on the screen (Figure 22a);
- **Condition 4 (abbreviated as fixed-intersection or CFI):** The interaction devices are the same as in condition 2. However, the visual feedback consists of both a 3D rendering result and a synchronized 2D intersection image in another window (Figure 21a and b);
- **Condition 5 (abbreviated as tangible-intersection or CTI):** The interaction devices are the same as in condition 3. However, the visual feedback consists of both a 3D rendering result and a synchronized 2D intersection image in another window (Figure 22a and b).

Participants will be randomly assigned into one of five groups. The first group will use only the baseline system, which contains a tangible wooden cube within the tracked environment to explore the volumetric data. The second group will use the same baseline system, together with a virtual plane on the screen in a fixed position. The participants in this group can manipulate the cube and intersect the virtual cube with the fixed virtual plane, so that part of the 3D data can be made transparent. The third group will use a tangible frame to position the virtual clipping plane and cut through the data. The fourth group will use the same setup as the second group plus an additional 2D intersection image for a fixed clipping plane. The fifth group will use a tangible frame to generate both a 3D rendered result and an additional 2D intersection image presented in a separate window to help participants observe the intersection result.

CONCLUSION

This chapter presents two main results. The first is an empirical evaluation comparing human performance using different VR systems for four generic volume visualization tasks. The second is the design of alternative tangible interfaces for performing a clipping plane operation. The tasks in the first experiment were derived from data

Figure 20. Condition 1 in the planned experiment

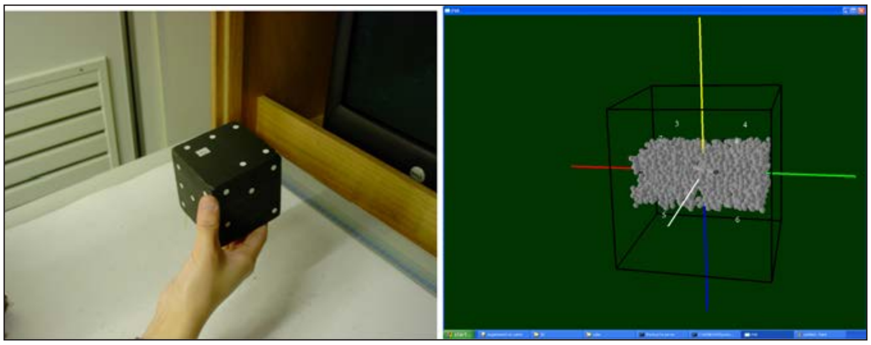
(a) (b)

Figure 21. The diagram of condition 2 and condition 4: condition 2 includes (a); condition 4 includes the intersection image (b) as well

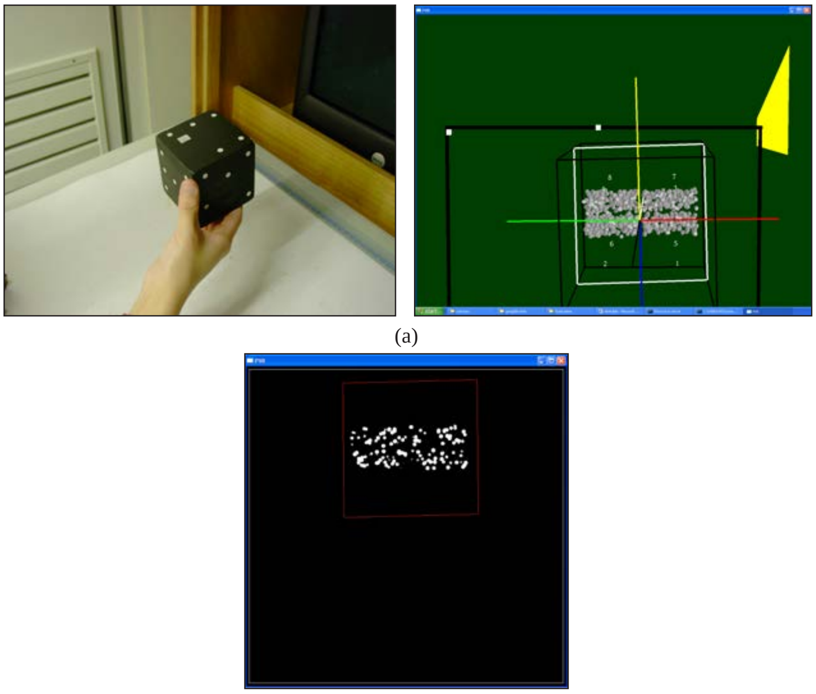

(a)

(b)

Figure 22. The diagram of condition 3 and condition 5: condition 3 includes (a); condition 5 includes the intersection image (b) as well

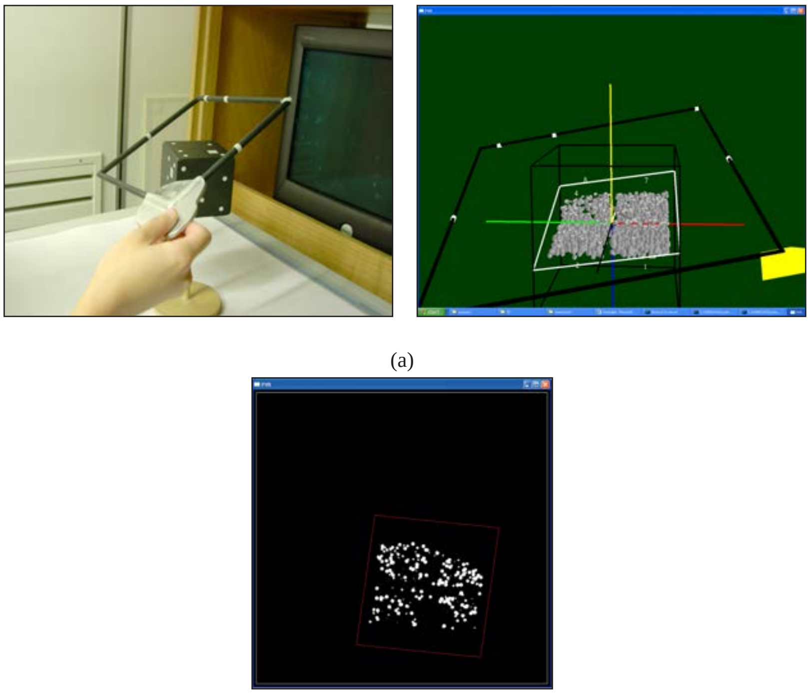

(a)

(b)

with characteristics and questions being asked by researchers studying mucuciliary clearance in CF. Results showed that the haptic system offered participants both an inside-out and an outside-in perspective on a volume, a property that was identified as important to completing our tasks. Participants using the HMD VR system were significantly slower than participants from the other two systems, and were less accurate for the shape, density, counting, and spatial tracking questions. Finally, none of the systems allowed for accurate judgment of different sizes of objects, or of which regions of a volume had the densest spatial packing.

The speed difference for the HMD system was not unexpected, but the inferior task performance was quite surprising. Participants' responses to questionnaires and anecdotal comments reveal that memory load was a significant factor. In the absence of an overview capability, participants were forced to make an internal representation of the total volume; the dense nature of the data removed visible landmarks that can normally provide such a frame of reference. It is believed that a future planned system that includes both an overview and an inside-out capability within the HMD would produce a system whose performance is at or above the level of the haptic-enabled system for some tasks. Furthermore, the poor performance of the HMD VR system for this data visualization task does not mean it is not appropriate for other tasks or applications. The lack of reference of the frame does not exist for other applications, for example, gaming or architecture.

The design of a multimodal interface asks the designer to consider how the brain combines and integrates different sources of information in order to make the interface truly helpful. Correct combination and integration of multiple sources of sensory information, for example vision and touch, is the key to create a robust perception and judgment for search tasks in a multi-modal interaction situation. Combination does not only mean the presence of two modalities, but an integration and coordination that match the user's physiological senses. We can observe from the experimental results that the haptic system has different effects on the user's performance for different tasks or different conditions of the same task. For some tasks, the presence of haptics maximizes the information received from both modalities (vision and touch). It also reduces the variance of the sensory estimation to increase its reliability. For some tasks, it does not.

Based on the experience of the VR study and its results, we proposed the design of tangible interfaces, particularly a tangible clipping plane for improving task performance. The detailed design strategy for such interfaces was described and the planned further user study will give us the opportunity to quantitatively measure the benefits brought by the tangible objects for the same tasks studied in our original experiment.

FUTURE RESEARCH DIRECTIONS

3D interaction through AR, VR or tangible interfaces continues to be an interesting field for interface experiments. Currently, 3D interaction itself is still an experience, instead of a routine. Practically, though, there still lacks enough evidence that 3D interfaces may improve the speed of interaction, or give the user a better understanding of the observed data. It is argued that current 3D interfaces are simply not right for simulated 3D environments yet. It is also difficult to control a 3D space with the interaction techniques that are currently in use, since they were designed for 2D manipulation (dragging, scrolling) or derived from 2D manipulation.

So regarding 3D interfaces and interactions, the future work should, in our view, concentrate on the following:

- Advancing technology: technology advances can solve part of the existing problems in 3D interaction. A new generation of 3D LCD displays already can be found in the market in terms of hardware development.

- Improving the ergonomics of 3D interfaces: one factor that can lead to frustration with 3D interfaces is the poor ergonomic design of 3D interfaces (devices). For example, it takes time until a user gets used to the heavy goggles, sometimes too long for acceptance.
- Further understanding the perceptual and cognitive issues behind 3D interfaces: Navigating through a 3D space can be natural and attractive in the beginning, but after a while such space and way of interaction may become an obstacle for a user. Although bad ergonomic design may be one of the reasons, further study on perceptual and cognitive issues may uncover additional facts.
- Proposing better evaluation methods and collecting experimental evidences: The novelty and the limitless possibilities of 3D user interfaces and interaction research have resulted in a practice where researchers mostly focus on developing new devices, interaction techniques, and user interface metaphors. Evaluations have not been used much to actively influence the design process. At the same time, although there are several evaluation methods in HCI, customized evaluation methods for 3D interaction are still in their infancy. Proposing customized evaluation methods and carrying out more systematic evaluations should be important future work so that more experimental evidence can be collected to guide the design of 3D interaction.

So with the future development of technology and more understanding of the principles behind 3D interaction, 3D interface will undoubtedly become more useful for data analysis by scientists and professionals.

REFERENCES

Arthur, K., Booth, K., & Ware, C. (1993). Evaluating 3D task performance for fish tank virtual worlds. *ACM Trans. Inf. Syst, 11*(3), 239-265.

Bowman, D., Kruijff, E., LaViola, J., Jr., & Poupyrev, I. (2001). An introduction to 3D user interface design. *Presence, 10*(1), 96-108.

Bowman, D., Datey, A., Ryu, Y., Farooq, U., & Vasnaik, O. (2002). Empirical comparison of human behavior and performance with different display devices for virtual environments. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting,* (pp. 2134-2138).

Brederson, J., Ikits, M., Johnson, C., & Hansen, C. (2000). The visual haptic workbench. *Proceedings of PHANToM Users Group Workshop,* (pp. 46-49).

Brooks, F. (1999). What's real about virtual reality?. *IEEE Computer Graphics and Applications, 9,* 16-27.

Brooks, F., Ouh-Young, M., Batter, J., & Kilpatrick, P. (1990). Project GROPE- haptic displays for scientific visualization. *Computer Graphic*s, *24*(4), 177-185.

Chen, M., Mountford, S., & Sellen, A. (1988). A study in interactive 3D rotation using 2D control devices. *Computer Graphics, 22*(4), 121-129.

Cruz-Neira, C., Sandin, D., & DeFanti, T. (1993). Surround-screen projection-based virtual reality: The design and implementation of the CAVE. *ACM Computer Graphics, 27*(2), 135-142.

Demiralp, C., Laidlaw, D., Jackson, C., Keefe, D., & Zhang, S. (2003). Subjective usefulness of CAVE and fish tank VR display systems for a scientific visualization application. *Proceedings of IEEE Visualization*. Seattle, WA.

De Guzman, E., Ho-Ching, F., Matthews, T., Rattenbury, T., Back, M., & Harrison, S. (2003).

Eewww!: Tangible instruments for navigating into the human body. *Extended Abstracts of CHI 2003,* 806-807.

Ernst, M., & Banks, M. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429-433.

Froehlich, B., & Plate, J. (2000). The cubic mouse: A new device for three-dimensional input. *Proceedings of the SIGCHI conference on Human factors in computing systems,* (pp. 526-531).

Hansen, C., & Johnson, C. (2004). *The visualization handbook*. MA: Elsevier Butterworth Heinemann.

Hinckley, K., Tullio, J., Pausch, R., Profftt, D., & Kassell, N. (1997). Usability analysis of 3D rotation techniques. *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology,* (pp. 1-10).

Hinckley, K., Pausch, R., Kassell, N., & Goble, J. (1994). A three-dimensional user interface for neurosurgical visualization. *The SPIE Conf. on Medical Imaging,* (pp. 126-136).

Ishii, H., Ben-Joseph, E., Underkoer, J., Yeung, L., Chak, D., Kanji, Z., & Piper, B. (2002). Augmented urban planning workbench: Overlaying drawings, physical models and digital simulation. *ISMAR '02: Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR'02),* (pp. 203). Washington, D.C.: IEEE Computer Society.

Kosara, R., Healey, C., Interrante, V., Laidlaw, D., & Ware, C. (2003). User studies: Why, how, and when?. *IEEE Computer Graphics and Applications*, (pp. 20-25).

Kreuger, W., Bohn, C., Froehlich, B., Schueth, Strauss, H., & Wesche, G. (1995). The responsive workbench: A virtual work environment. *IEEE Computer, 28*(7), 42-48.

Levoy, M. (1998). Display of surfaces from volume data. *IEEE Computer Graphics and Applications, 8*(5), 29-37.

Martens, J., van Liere, R., & Kok, A. (2007). Widget manipulation revisited: A case study in modeling interactions between experimental conditions. *IPT-EGVE 2007.* The Eurographics Association.

Meehan, M., Whitton, M., & Brooks, F., Jr. (2002). Physiological measures of presence in stressful virtual environments. *Proceedings of ACM SIGGRAPH2002, 21,* (pp. 645-652).

Mulder, J., & van Liere, R. (2002). The personal space station: Bringing interaction within reach. *Proceedings of VRIC 2002 Conference* (pp. 73-81).

Qi, W., Taylor, R., Healey, C., & Martens, J. (2006). A comparison of immersive HMD, fish tank VR and fish tank with haptics displays for volume visualization. *Proceedings of the 3rd symposium on Applied perception in graphics and visualization 2006,* (pp. 51-58).

Razzaque, S., & Whitton, M. (2001). Redirected walking. *Proceedings of Eurographics 2001,* (pp. 289-294).

Schroeder, W., Avila, L., & Hoffman, W. (2000). Visualizing with VTK: A tutorial. *IEEE Comput. Graph. App., 20*(5), 20-27.

Schulze, J., Zeleznik, J., Forsberg, A., & Laidlaw, D. (2005). Characterizing the effect of level of immersion on a 3D marking task. *Proceedings of HCI International,* (pp. 447-452). HCI International.

Sutherland, I. (1968). A head-mounted three-dimensional display. *Proceeding of the Fall Joint Computer Conference,* (pp. 757-764).

Taylor, R., Chen, J., Okimoto, S., Llopis-Artime, N., Chi, V., Brooks, F., et al. (1997). Pearls found on the way to the ideal interface for scanned-probe microscopes. *Proceedings of Proceedings of IEEE Visualization*, (pp. 467-470).

Taylor, R., Hudson, T., Seeger, A., Weber, H., Juliano, J., & Helser, A. (2001). VRPN: A device-independent, network-transparent VR peripheral system. *VRST'01: Proceedings of the ACM symposium on Virtual reality software and technology,* (pp. 55-61).

Ullmer, B., Ishii, H., & Glas, D. (1998). Mediablocks: Physical containers, transports, and controls for online media. *Proceedings of SIGGRAPH '98,* (pp. 379-386).

Ullmer, B., & Ishii, H. (2001). Emerging frameworks for tangible user interfaces. In: J. Carroll (Ed.), *Human-computer interaction in the new millennium*, (pp. 579-601). Addison-Wesley.

van Liere, R., & Mulder, J. (2003). Optical tracking using projective invariant marker pattern properties. *Proceedings of the IEEE Virtual Reality Conference 2003,* (pp. 191-198).

Ware, C., Arthur, K., & Booth, K. (1993). Fish tank virtual reality. *Proceedings of CHI 93,* (pp. 37-42).

Ware, C., & Franck, G. (1996). Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Trans. Graph., 15*(2), 121-140.

Westermann, R., & Ertl, T. (1998). Using graphics hardware in volume rendering applications. *Proceedings of SIGGRAPH'98,* (pp. 169-178).

Zhai, S. (1995). *Human performance in six degrees of freedom input control*. Doctoral Thesis, University of Toronto, Toronto.

ADDITIONAL READINGS

Card, S., Mackinlay, J., & Robertson, G. (1990). The design space of input devices. *Proceedings of the 1990 ACM Conference on Human Factors in Computing Systems (CHI'90)* (pp. 117-124). ACM Press.

Card, S., Moran, T. & Newell, A. (1986). The model human processor. In: K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance* (vol. 1, pp. 451-455). John Wiley & Sons.

Card, S., Moran, T., & Newell, A. (1983). *The psychology of human-computer interaction*. Lawrence Erlbaum Associates.

Fitts, P. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology, 47,* 381-391.

Gabbard, J., Hix, D., & Swan, J. (1999). User-centered design and evaluation of virtual environments. *IEEE Computer Graphics & Applications, 19*(6), 51-59.

Hix, D., & Hartson, H. (1993). *Developing user interfaces: Ensuring usability through product & process*. John Wiley and Sons.

Jacob, R. (1996). Input devices and techniques. In: A. Tucker (Ed.), *The computer science and engineering handbook,* (pp. 1494-1511). CRC Press.

LaViola, J. (2000a). A discussion of cybersickness in virtual environments. *SIGCHI Bulletin, 32*(1), 47-56.

Mine, M., Brooks, F., & Sequin, C. (1997). Moving objects in space: Exploiting proprioception in virtual environment interaction. *Proceedings of SIGGRAPH'97,* (pp. 19-26). ACM Press.

Milgram, P., & Kishino, F. (1994). A taxonomy of mixed reality visual displays. *IECE Transactions on Information and System, E77-D*(12), 1321-1329.

Preece, J., Rogers, Y., & Sharp, H. (2002). *Interaction design: Beyond human-computer interaction*. John Wiley and Sons.

Shneiderman, B. (1998). *Designing the user interface: Strategies for effective human-computer interaction*, (3rd ed.). Addison Wesley.

State, A., Livingston, M., Hirota, G., Garrett, W., Whitton, M., Fuchs, H., & Pisano, E. (1996). Technologies for augmented reality systems: Realizing ultrasound-guided needle biopsies. *Proceedings of SIGGRAPH'96,* (pp. 439-446). ACM Press.

Stanney, K., Mourant, R., & Kennedy, R. (1998). Human factors issues in virtual environments: A review of the literature. *Presence: Teleoperators and Virtual Environments, 7*(4), 327-351.

Srinivasan, M., & Chen, J. (1993). Human performance in controlling normal forces of contact with rigid objects. *Advances in Robotics, Mechatronics, annd Haptic Interfaces, ASME,* (pp. 119-125).

Simon, A., & Fröhlich, B. (2003). The yoyo: A handheld device combining elastic and istonic input. *Proceedings of INTERACT 2003,* (pp. 303-310). Zurich: IOS Press.

Usoh, M., Catena, E., Arman, S., & Slater, M. (2000). Using presence questionnaires in reality. *Presence: Teleoperators and Virtual Environments, 9*(5), 497-503.

Usoh, M., Arthur, K., Whitton, M., Bastos, R., Steed, A., Slater, M., & Brooks, F., Jr. (1999). Walking > Walking-in-Place > Flying in Virtual Environments. *Proceedings of SIGGRAPH '99,* (pp. 359-364). ACM Press.

van Dam, A., Forsberg, A., Laidlaw, D., LaViola, J., & Simpson, R. (2000). Immersive VR for scientific visualization: A progress report. *IEEE Computer Graphics & Applications, 20*(6), 26-52.

van Dam, A. (1997). Post-WIMP user interfaces: The human connection. *Communication of the ACM, 40*(2), 63-67.

Welch, G., Bishop, G., Vicci, L., Brumback, S., Keller, K., & Colucci, D. (1999). The hiball tracker: High-performance wide-area tracking for virtual and augmented environments. *Proceedings of the 1999 ACM Symposium on Virtual Reality Software and Technology (VRST'99),* (pp. 1-10). ACM Press.

Welch, R., & Warren, D. (1986). Intersensory interactions. In: K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance*, (vol. 2, pp. 251- 256). John Wiley & Sons.

Ware, C. (2000). *Information visualisation: Perception for design*. San Francisco, CA: Morgan Kaufman.

APPENDIX: POST-EXPERIMENT QUESTIONNAIRE

Participant ID:_____________________________

The following questions relate to the Virtual Reality (VR) system you have experienced during the experiment. Please select the correct one:

- ○ **Immersive HMD VR**
- ○ **Fish tank VR**
- ○ **Fish tank with Force-feedback (haptic)**

1. Please rate your sense of being in the virtual environment that has the simulation data on the following scale from 1 to 7, where 7 represents your normal experience of being in a place.
   I had a sense of being in the virtual environment containing the simulation data:

   (Not at all) 1 2 3 4 5 6 7 (Very much)

2. Please rate any sense of immersion you experienced when looking into the dataset.
   The sense of immersion I experienced was...

   (Not at all) 1 2 3 4 5 6 7 (Very much)

3. How difficult or straightforward was it for you to understand spatial relationships between objects in the virtual environment while working with the system?
   The spatial relation was...

   (Very difficult) 1 2 3 4 5 6 7 (Very straightforward)

4. Did you find it relatively simple or relatively complicated to move through the virtual environment and the simulation data?
   To move through the virtual environment was...

   (Very complicated) 1 2 3 4 5 6 7 (Very simple)

5. The act of moving from place to place in the virtual environment can seem to be relatively natural or relatively unnatural. Please rate your experience of this.
   The act of moving from place to place seemed to be...

   (Very unnatural) 1 2 3 4 5 6 7 (Very natural)

6. How often did you to feel you were in virtual environment when observing the simulation data and searching for the required structure?
   I felt I was in virtual environment...

   (Very few) 1 2 3 4 5 6 7 (Very often)

7. To what extent were there times during the experience when you felt dissatisfied with the interface?

continued on following page

There were times during the experience I felt dissatisfied...

(At no time) 1 2 3 4 5 6 7 (Almost all of the time)

8. To what extent did you have a sense of acting in the virtual space, rather than operating something from outside?

(Not at all) 1 2 3 4 5 6 7 (Very much)

9. To what extent did you feel that the virtual environment surrounded you?

(Not at all) 1 2 3 4 5 6 7 (Very much)

10. To what extent did you feel like you just perceived pictures?

(Not at all) 1 2 3 4 5 6 7 (Very much)

11. How much did your experience in the virtual environment seem consistent with your real-world experience?

(Not at all) 1 2 3 4 5 6 7 (Very much)

12. To what extent do you think the virtual reality system you experienced helped you identify the structure within the volumetric simulation data?
I thought the virtual environment helps me...

(Not at all) 1 2 3 4 5 6 7 (Very much)

13. How easy to use was this virtual reality system?

(Hard to use) 1 2 3 4 5 6 7 (Very easy)

14. How effective do you feel you were when working with the virtual reality system compared with a traditional desktop system?

(No difference) 1 2 3 4 5 6 7 (Very effective)

15. How helpful were the depth cues in this virtual reality system compared to the traditional desktop system?

(No difference) 1 2 3 4 5 6 7 (Very much)

16. Rate the degree of difficulty in carrying out the task, for the virtual reality system you experienced:

(Not difficult at all) 1 2 3 4 5 6 7 (Very difficult indeed)

17. How well could you examine objects from multiple viewpoints?

(Very difficult) 1 2 3 4 5 6 7 (Very easy)

continued on following page

18. How much delay did you experience between your actions and expected outcomes?

(None at all) 1 2 3 4 5 6 7 (Very much)

19. When exploring the virtual space, did the objects appear too compressed or too magnified?

(Not at all) 1 2 3 4 5 6 7 (Very compressed)

(Not at all) 1 2 3 4 5 6 7 (Very magnified)

20. During the experiment, your general level of confidence in the accuracy of your answers was:

(Just guessing) 1 2 3 4 5 6 7 (Very sure)

21. During the task, identifying the individual shape and location of an object within the environment was:

(Difficult) 1 2 3 4 5 6 7 (Easy)

22. During the task, identifying the global topology of the simulation data was through the VR system:

(Difficult) 1 2 3 4 5 6 7 (Easy)

23. During the task, what was the level of demand on your memory?

(Small) 1 2 3 4 5 6 7 (Large)

24. How effective was the sense of perspective (further objects appeared the correct size compared to nearer objects)?

(Ineffective) 1 2 3 4 5 6 7 (Effective)

25. Did you think that the VR system you experienced changes the way you observe and analyze the data, in comparison to traditional media?

○ **Didn't change at all**
○ **Changed just a little bit**
○ **Changed slightly**
○ **Changed quite some**
○ **Changed radically**

Answer the following questions if you experienced fish tank VR with haptic device

1. How consistent or inconsistent was the information coming from your various senses (visual and haptic feedback)?

(Inconsistent) 1 2 3 4 5 6 7 (Very Consistent)

continued on following page

2. How well could you actively survey or search within the virtual environment using touch? Searching within the virtual environment through touch was

(Very difficult) 1 2 3 4 5 6 7 (Very easy)

3. Do you think the addition of multimodal sensory stimulation (haptics) would help you?

A) Create a better understanding of space:
- ○ **Not at all**
- ○ **Small chance**
- ○ **Possibly**
- ○ **Likely**
- ○ **Very likely**

B) Create better understanding the structure and topology of the data:
- ○ **Not at all**
- ○ **Small chance**
- ○ **Possibly**
- ○ **Likely**
- ○ **Very likely**