Spatial Perception in Virtual Environments:

Visual Cognition Gain with Head Mounted Displays

By

Christopher Anthony Peri

B.A. (University of California, Berkeley) 1995

A thesis submitted in partial satisfaction of the
requirements for the degree of

Master of Science

in

Architecture

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor Yahuda Kalay
Professor J.P. Protzen
Professor Stephen E. Palmer

1997

Thesis Statement

The ability to navigate a space at one's own volition can substantially accelerate cognitive understanding as information about the environment is no longer abstracted into a static 2-dimensional image, so that the viewer does not have to make a cognitive leap to understand the environment that is being presented. With real-time visualization tools, like the HMD, visual cognition occurs at nearly the same rate in the virtual world as it does in the real world and integrates the user's point of view into the medium, thus allowing for a level of believability and comprehension not possible in a 2-dimensional image. We are entering an age when more and more information is becoming available in 3D format, and how we interact with that spatial information will greatly affect the completeness of the message embedded in the information.

In this thesis I intend to explore how highly immersive visualization tools help the viewer gain a cognitive understanding of a virtual environment. I will also report on an experiment designed to explore this very issue and the many other related issues that must be considered when dealing with real-time visualization tools.

"The computer, with its ability to manage enormous amounts of data and to simulate reality, provides a new window on (the) view of nature. We may begin to see reality differently simply because the computer produces knowledge differently from the traditional analytic instruments. It provides a different angle on reality." (Pagels, 1988) Similar to how perspective drawing changed our view of reality (the world) in the Renaissance, and how the microscope changed our view of the natural world, 3D visualization will affect us in ways yet unseen. By making information available at real-time, we can now reexamine how we interact with that information.

INTRODUCTION

When I first began research on this thesis, I expected to find a wealth of information about spatial cognition using a variety of display devices. Instead I found that some of the most basic issues regarding human-computer interaction in virtual environments have not been approached. I was also quite surprised at the misconception of what a Head Mounted Display (HMD)[1] was and how it could fit into the everyday use of visualization.

This thesis is to be considered an opening salvo in the investigation of real-time visualization tools for professional as well as pedagogical uses in this college. I intend this thesis to inform myself and others of the measurable advantage of using a HMD over a computer monitor to view virtual environments in real-time, as well as issues involved in taking that measure.

I have found that those using a HMD to perceive and represent a simple environment using a Head Mounted Display can represent that environment with 30% to 40% more accuracy than using a monitor to perceive the same environment.

Prior Research

Although there is not much information available in basic research, I did find one document that approaches the question of monitor vs. HMD performance.

Quantifying Immersion in Virtual Reality is a paper from the University of Virginia, written by Randy Pausch, Dennis Proffitt, and George Williams.

---

[1] The head-mounted display (HMD) is the basic output device of the VR system. Two color displays located inside the HMD provide the user with the only visual signal received, effectively isolating the user from the external world and immersing him or her in the displayed imagery. The HMD houses a pair of stereo earphones and a position/orientation sensor is located atop the HMD. The sensor, connected to the tracking device , provides information about the user's position and orientation in space. *http://www-pablo.cs.uiuc.edu/Projects/VR/vr.html* WebPage by Eric Shaffer. 1996

In this paper, we show that users with a VR interface complete a search task faster than users with a stationary monitor and a hand-based input device. We placed users in the center of the virtual room ... and told them to look for camouflaged targets. VR users did not do significantly better than desktop users. However, when asked to search the room and conclude if a target existed, VR users were substantially better at determining when they had searched the entire room. Desktop users took 41% more time, re-examining areas they had already searched.

Quantifying Immersion in Virtual Reality (currently submitted to ACM SIGGRAPH 1997)

This test demonstrates the advantage of an internal point of reference vs. an external point of reference for navigation. In the test, the desktop users did not actually use a desktop monitor but instead, viewed the environment with the HMD with the Head tracking turned off. This was done to reduce the variables of Field of View, Resolution, and so on. This exercise is similar to the pointing exercise, which I will describe later in the thesis. Like the pointing exercise, this does not include lateral navigation, which is a major difficulty in virtual environments.. The advantage in performance comes not in the search but in the formation of a cognitive map which can be seen in the desktop users larger re-examining time.

Spatial Communication

Spatial information can come in a variety of forms. The form that it takes depends on what information the author wants, or hopes that the viewer will gain. Since the spatial information needs to be as complete as possible, the information is sometimes broken up into palatable sizes and forms. The sizes and forms are chosen depending on who receives the information.

Since spatial designs are expressed to people using representation[2], the main objective is to try and communicate cognitive understanding by immersing the viewer in the design. Creating that design is only a small part of the designer's task; the largest part is effectively communicating that idea to other

[2] Representation can be defined as a process of abstraction and communication. Through some symbolic language, characteristics of a real or hypothetical object or experience are conveyed by one person to another. **Steinfeld and Kalay**

participants involved in the design process (clients, engineers, construction managers, regulatory authorities, financial institutions, etc.).

Even though architectural design is a social act (Mitchell, 1994), in the past, architects have expressed their ideas through limited conceptual media which do not efficiently communicate a cognitive understanding of the design to the viewer. Since the Renaissance, designers have been using the same representational tools to communicate spatial ideas: 2-dimensional drawings and images, (See Figure 1) and three-dimensional models. Many viewers, not initiated into the designers' systems of representation, find it difficult to infer cognitive understanding from the visual stimuli presented. The necessary abstraction inherent in these media obstructs the viewer's "visual cognition" [3], thus limiting spatial understanding.

Types of information media

Since spatial designs are communicated to people using some form of representation, designers have been searching for a medium that will enhance the cognitive content of the communication without distorting the message.

McLuhan made the distinction between television (video) and cinema as 'cool' and 'hot' media, respectively. The difference between the two is the level of immersion that the viewer can achieve. A cool medium, like television, has a display area that is below the total vision capability of the observer. A hot medium like large screen motion pictures is equal to or above the viewer's visual capacity. Both of these media employ dynamic image display, displaying images at a rate that achieves the illusion of smooth motion. The cool medium however suffers from either scale or display area. When presenting large objects, like building or ships, the image must either scale down or only be partially seen. A hot medium like cinema can display images at a much larger scale, even larger than the viewer, hence allowing more information about an environment to be displayed.

[3]"Visual Cognition": the ability of humans to infer meaning from visual stimuli rather than other forms of information. **(Bronowski 1978)**

Recently, design communicators have moved away from passive/static means, (See Figure 2) like drawings and scale models, and moved towards passive/dynamic means, such as video and cinema, where the observer follows a pre-scripted trajectory/story. What is common between the two means is that the scene or movement in the environments is predetermined, making the viewer a voyeur of that world.

Passive/Static Drawings Photos	***Passive/Dynamic*** Video Cinema
Active/Static Scale Models Holograms	***Active/Dynamic*** Real-Time Imaging VE/VR

Figure 1: Medium mode chart

Real-time imaging is the change of the medium from passive/dynamic to active/dynamic. 'Active' refers to the viewer's ability to physically alter the display of the medium. The message itself is now open to the viewer's conduction and thus the message will be influenced by the viewer's volition, not by that of a preset script. How the images are presented is still a concern. Though better than passive/dynamic, if the real-time images are presented in a small medium, like a computer monitor, it would still be 'cooler' than if presented on a large screen(s) as seen with the CAVE™ or a movable screen like the HMD.

So, if as McLuhan says, television is a *cool* medium and cinema is a *hot* medium, then interactive real-time viewing must be a *scalding* medium.

Orientation and Cognitive Mapping

The ability to navigate an environment at one's own volition can accelerate cognitive understanding of that space exponentially, and thus affect how the message is received. "People are not sponges, and do not learn by absorbing information. Rather, they construct knowledge intentionally and actively. Simply put, one learns by experience: exploring and doing." (J. Lansdown, 1994)(1.4) In reality, we construct our view of the world as it is relative to our place in it. Our medium for viewing the world is not only our eyes but also our heads, the direction that we are facing, the overlap of seeing our hands on, over, or under an object. The individual's senses provide direct sources of information and are more effective in cognitive mapping[4] formation than indirect sources. (Downs, R.M. and Stea, D.) Our bodies are constant sources of orientation to everything that is around us. One can imagine knowing that a light switch is just to the right of one's shoulder when sitting in a chair. We do not need to see the switch, because the cognitive map location of that switch is inferred from our body's orientation in space. When we view spatial images, even animated walkthroughs, we will always have a limited cognitive map of what we see because the movement through that environment is not based on our physical orientation. We gain understanding of our environment through our own movement within it. "Everyday perception occurs in a context of nested motions. Eyes move within heads, heads move on bodies, and bodies move in surroundings... Not only is motion a necessary condition for perception, but it is also a sufficient condition for the perception of a variety of environmental properties." (Proffitt & Kaiser, 1993) Thus the information produced by locomotion is fundamental to an individual's spatial orientation. (Billinghurst and Weghorst, 1995) We see the world based on how we know ourselves in the world.

This is an important point to realize when we try to communicate the qualities of a spatial design. We want the receiver of the message to gain a cognitive map of the spatial design that is as complete as possible. Since the task of communicating the cognitive map to another person (to confirm the map's accuracy) requires a new abstraction within a new representation, the receiver needs as much initial cognitive understanding as possible.

Communication and Representation

"Representation is a process with two main components: *abstraction* and *communication*. Through this process, information and meaning are transformed from one form into another... The changes in form require the party who receives the information to interpret it, in order to recover most, if not all, the conveyed meaning." (Steinfeld and Kalay.) In addition, there has to be some sort of agreement as to the

[4] Cognitive mapping is a process composed of a series of psychological transformations by which an individual acquires, codes, stores, recalls, and decodes information about the relative locations and attributes of phenomena in his everyday spatial environment. **(Downs & Stea)**

intent of the abstraction as well as the meaning of the symbols used for the communication. If there is not a prearranged agreement on the meaning of each symbol, then communication is impaired.

Tools for Real-Time Interaction

Accelerated Spatial Communication

Currently we are interacting with our computers using only a small percentage of our visual and cognitive abilities. In most applications this would be sufficient, but over the last few years, computer power and the amount of information computers can deliver have exploded exponentially. In addition to more information being available to the user, there are also more forms of that information. We can now hear sounds, look at images, watch animations, view live video from remote locations, interact with 3D models in real time, and more. Each of these new forms of information requires its own mode of transmission from the computer to the user to avoid information loss or confusion.

The advance of computer technology has made many new tools available to increase the speed and precision of spatial information storage, editing, and communication. One of these technologies is the Head Mounted Display (HMD), a common tool in the world of Virtual Reality (VR). The HMD allows the viewer to 'see' an environment similar to reality, while maintaining a high degree of navigational freedom. The HMD also involves the viewer in the scene itself, thus making the viewer's body a point of reference for navigating the simulated environment. "The key benefit of head-mounted systems is the freedom from the display screen bottleneck. Computer technology has an astonishing processing capacity but the output is limited to a 2-dimensional display generally no larger then a sheet of 36" x 42" paper." (Lansdown, 1994)

The HMD allows the viewer access to an environment in real-time. "A major distinction between virtual reality (VR) and other forms of computer simulation lies in what Meredith Bricken has called 'inclusiveness': the ability of the participant to interact with the computer-generated environment as though he or she were actually inside of a wholly contained world." (Ann Lasko-Harvill)

By involving the viewer in the orchestration of environments unfolding, information about the environment is no longer abstracted to a single 2-dimensional image. The viewer does not have to make a cognitive leap to understand the environment that is being presented. Visual cognition occurs at the real-time rate natural cognition occurs, thus adding a level of believability and comprehension that is not possible in a passive/static or passive/dynamic image.

What HMDs can't do.

Despite the leap of virtual spatial exploration and cognition that the HMD allows, one must keep in mind that the HMD is still only displaying two 2-dimensional images; HMDs are not 3-dimensional viewers as many people, including many VR advocates, seem to believe. Actual human stereo perception requires more than the simple 6.5 centimeters (average distance between two eyes) convergence adaptation. When viewing a stereo image, similar to what could be seen with the stereoscope of the early 1900's, our eyes cannot use convergence to focus on an object in the image because the convergence is preset. We cannot set the HMD LCDs closer to the eye because of the strain caused from accommodation. The following is a better and more complete explanation made available at the Edinburgh Virtual Environment Laboratory, Department of Psychology, University of Edinburgh.

> *In 1965 Ivan Sutherland [1] presented a research report discussing "the ultimate display". The prototype of this used two small CRT screens viewed through prisms and magnifying optics. It is worth noting that although this HMD did have the potential for stereoscopic presentation Sutherland [2] placed considerable emphasis on the "kinetic depth" cues the display could provide rather than disparity cues. This prototype, however, set the format for a host of HMDs that followed and most of the currently available/used models conform to a very similar design. Here-in lies a problem that arises out of a simplifying assumption that "if we place 2-dimensional images on the observer's retinas, we can create the illusion that he is seeing a three-dimensional object" (Sutherland [2]). This statement appears correct, but we would modify it by stating that "an illusion of a 3D object can be produced, but 3D space can not be rendered with integrity from dual 2-D images". The implications of this re-statement arise when we replace the static stereogram, which provides an illusion of surfaces at defined depth increments, with a representation of a 3D world with objects*

at a full range of disparity increments, which the observer then attempts to selectively sample.

1. Sutherland I.E. (1965) The Ultimate Display. Proc. IFIPS)

2. 506-508 2. Sutherland I.E. (1968) A head-mounted three dimensional display. Fall Joint Computer Conference, 33, 757-764)

Although this seems like an unresolvable problem, the ability to selectively sample objects is not critical, as visual cognition can occur even within the most limited of viewing conditions. Additionally, since stereo perception is only one of many cues of spatial perception, (e.g. binocular parallax, motion parallax, accommodation[5], convergence, size, texture, brightness, and air-perspective contrast), the computer power required, and thus the frame rate sacrifice, does not compensate for the limited benefits of stereoscope perception. What is important to note is that we are concerned with improving visual cognition in a virtual environment, not duplicating real-world perception.

Simulator Sickness

A limitation of the HMD, and what I consider to be its greatest limitation as the technology stands today, is simulator sickness.

VOR

Despite the commonality of simulator sickness, little is known about it. The most common source of simulator sickness is Vestibulo-ocular Reflex (VOR)[6]; The vestibular apparatus is a small structure that exists in the bony labyrinth of the inner ear, whose function is to sense and signal movements of the head. This function is extremely important because it contributes to the coordination of motor responses for the

---

[5] The iris, similar to a zoom lens on a camera, can focus in and out to a 1 meter field-of-view by using muscles attached to the lens thus changing its power

[6] The VOR is only one of five major types of eye movements. The other four (optokinetic, saccade, smooth pursuit, and vergence) interact dynamically with the VOR to 1) bring visual targets onto the fovea, and 2) keep them there. Of these eye-movements types, optokinetic and smooth pursuit can be considered tracking eye movements, as they both use input signals for directing the eye that are derived from the retina itself, and their function is essentially one of image stabilization, not image capture. (Oman, 1990)

body. VOR is the fundamental eye-movement reflex that functions to keep images stabilized on the retina during movement of the head, thus helping to perform a very basic but important function, to allow sight during movement. The disagreement between what is seen by the body and what the vestibular apparatus senses can cause confusion, providing disagreeing signals to the brain and thus inducing sickness. VOR can adapt its signal to the brain when the environment demands it, but not everyone has the same VOR adaptation ability. This inability to resolve vestibular and ocular discrepancy is a major contribution to simulator sickness. The greater lag time between screen position updates and vestibular sensation can accelerate simulator sickness even more.

FOV and Parallax

Humans can see about 180 deg Field of View (FOV). To display that same amount of information, we would need a media display that is 180 deg wide. Since most display media cannot meet this criterion, a certain amount of adaptation occurs.

As with any perspective image, the amount of parallax warping is variable. The smaller the display real estate, the more warping occurs to display the same amount of information. To reduce the warping, less information must be displayed. To compensate, a certain amount of parallax occurs to try to display even a fraction of that information. To display a scene within a small display medium like a piece of paper, a computer monitor or even the HMD, a certain amount of parallax warping takes place. "In real environments, Alfano and Michel (1) have shown that reduction of peripheral vision impairs perception and visuomotor performance, both of which are essential for cognitive mapping ability." (Alfano, in Billinghurst and Weghorst, 1990) Similar to wearing a new prescription pair of glasses, the unfamiliar warping of the world's edges can be quite disorienting resulting in 'spectacle sickness'. (Oman)

Virtual Motion

A contributor to simulator sickness is the tendency of a viewer to lean forward when he/she is about to navigate forward in the virtual environment. In reaction to leaning forward, the viewer will then lean back, thus inducing a rocking motion not unlike what one experiences on a boat in a choppy sea.

Currently, research is being performed at the University of Washington to try and ease this problem. Called a Virtual Motion Controller,[7] this device will provide Sufficient-Motion[8] sensitivity to allow for kinetic input of forward motion from the body. This will reduce the leaning problem currently experienced with the use of the HMD alone, and thus reduce simulator sickness.

---

[7] Maxwell J. Wells, Barry N. Peterson, Jason Aten
Human Interface Technology Laboratory
Seattle, Washington

[8] Sufficient Motion, is when the simulator affords just enough movement in the real world to create a sense of reality in the virtual world

Methodology

Objective

> ***VRR****: What do you think is the greatest obstacle facing the VR industry and why?*
>
> ***Latta****: Lack of basic research. The issues of having the most intimate form of human computer interface in Virtual Reality necessitate a thorough understanding of human perceptual, muscle and psychological systems. Yet that research foundation does not exist..*
>
> Dr. John Latta interviewed in ***Virtual Reality Report, 2 (7) p.4***

Decisions about hardware, software, and environmental and behavioral models determine how a Virtual Environment (VE) looks and behaves which, in turn, influences human performance in the environment. (**Eggleston and Aldrich, 1996)**

It is quite surprising to me just how little is known about the relationship between spatial perception in computer environments as it relates to the user. Although a large part of my thesis is concerned with the effectiveness of the HMD to communicate spatial information to a viewer, I have dedicated a significant amount of space to problems of VR use and testing. The elements involved in the execution of spatial communication deserve as much attention as does the theory behind the communication increase.

Although initially I explored the performance relationship between viewing spatial environments in the HMD vs. the monitor, I found a surprising number of variables and unexpected difficulties. I do believe that the following experiments and resulting data are valid as a performance indicator and thus warrant further testing; however, I ask the reader to refrain from viewing this information as conclusive.

Summary of Test Procedure

The test objective in this thesis is a simple one: Does visual cognition improve with the use of a HMD vs. a computer monitor? Notice that I am not concerned with spatial perception of the two devices as they compare to real world performance; this issue will have to wait for another paper.

To compare spatial cognition gain via the HMD vs. the Monitor, a test environment was created of varying levels of difficulty, and test “subjects” (students) were asked to experience them. Two subject subgroups were created; a HMD group and a Monitor group. Each subject experienced virtual environments using the HMD and the monitor. Each subject was asked to create a sketch map of the virtual environments and note measurements on the sketch map. This data was gathered and averaged in a spreadsheet to calculate mean performances of each model based on subgroup. A second exercise tested the subjects’ abilities to point to objects with their eyes closed after seeing the objects in a virtual environment. Again the error information was gathered and averaged to show performance of one subgroup to another.

Testing Parameters

Hardware

The first step in my experiment was to select a HMD. Purchase funds were limited as well as the operating system platform. I chose to use the IBM class personal computer over UNIX based on the availability of machines, software availability and HMD compatibility. The limit of funding also restricted my choice of HMD. Very few HMDs can match the Field of View (FOV) and resolution sensitivity of human vision. The variables involved in visual quality depend on the FOV, resolution, image composition, and motion lag time just to name a few.[9] The VFX1 by Forte seemed to offer the best compromise between FOV, resolution and motion feed back. Under $800 as of Jan ’97, this would make the system affordable to most users. (See Appendix C for more information.) The VFX1 also comes with a handheld navigational device, the Forte Cyberpuck. Six degree navigation was attained by tilting the puck right, left, back, and forth to

[9] Motion lag time is a variable that takes the HMDs rate of position update and the computers systems frame rate generation time.

move the view accordingly. In addition, the user could also look up, down, right, and left by turning his/her head.

The second means of viewing the virtual models was via a common 17” computer monitor. Six degree navigation was attained by using the mouse. Moving the mouse right, left, back, and forth would move the view accordingly. By holding down a mouse button, the viewer could look up or down by moving the mouse forward or back. For most people the mouse was intuitive to use but they required some time to adapt to the cyberpuck.

The platforms used were generic PC class machines. The first half of the test was on a P100 with 32 megs of ram and a graphics card with 2 megs of ram. The second half was on a P200 with 64 megs of ram and a graphics card with limited 3D graphics acceleration with 4 megs of ram.

Software

The first choice for the experiments was to use one of the commercially available VRML[10] display programs. After much investigation, I was unable to find a program that could either generate sufficient frame rate above 320 x 200 resolution[11] or was compatible with the VFX1 HMD. I looked at a first person perspective game called 'Quake' by ID Software. I found that not only was the frame rate and image quality superior to most VRML programs, but that Quake had the best support for the VFX1 HMD.

Generating the models was accomplished using a third party program called 'WorldCraft' by Ben Morris. WorldCraft is a CAD modeling program specifically designed to create models for use in Quake. Although I was restricted in the choices of texture mapping, I felt this was of little consequence given the ease and speed that I could create models. Even so, of the texture maps available, I was able to find three that served my purposes. There are ways to create models that will display any texture desired, but time restrictions kept me from fully implementing the technique.

Models

Seven models in all were created. Six models were of varying complexity and the seventh was of a single room with seven items in the room used for the Pointing Exercise.

The six models were different from each other in layout complexity and number of walls, but in most other respects they were similar. Since the task was to discern the volumetrics of the models, visual cues such as furniture, changing textures on walls and floors, were avoided. This forced the subject to obtain volumetrics from motion parallax and convergence. The lack of objects also made it harder for the subjects to maintain orientation, thus forcing them to rely solely on their cognitive maps instead of augmenting their orientation

---

[10] Virtual Reality Modeling Language. This is a text formatted file that contained information to create 3D environments. VRML was initial designed to allow 3D communication across the Internet.

[11] Resolution is computed by the number of pixels per inch.

with *imageability*.[12] I will discuss this issue in more detail later in the thesis.

There were three sets of models, each containing two models that were very similar in design complexity and number of spaces. The first set was simple both in layout and number of spaces. The second set was of two models, more complex in layout and size, but of similar number of spaces. The third set was complex in layout and quite larger in the number of spaces. These were done to allow for the range in abilities of the subjects. (See Appendix A for floor plans of all models.)

A subject would view each model using either the HMD or the monitor. A model of similar complexity and number of spaces would then be viewed using the device not used previously. This experimental design compensates for the range of abilities of different subjects. If one was to overestimate the volumetrics of one model, then they would be most likely to overestimate the volumetrics of another.

The height of the point of view was set at 5'10". Ceiling heights were set at three varying levels to assist in motion parallax, and the GFOV[13] was set at 90 deg for the first half of the test and 60 deg for the remaining.

Test Subjects

Subjects were undergraduate architecture students from the University of California, Berkeley with one or two years of drafting experience in school. Most of the subjects were recruited by posting flyers around the College of Environmental Design. I was interested in non-computer users and made efforts to try and recruit those who were unfamiliar with the 3D computer environment. This was done to try and keep a balance between computer users and non-computer users. (See Appendix E for flyer.)

Information collected about each subject was his/her drafting experience and first-person-gaming

---

[12] "that quality in a physical object which gives it a high probability of evoking a strong image in any given absolver. It is that shape, color, or arrangement which facilitates the making of vividly identified, powerfully structured, highy useful mental images of the environment." (Lynch)

[13] The GFOV is the visual angle defined by the viewing frustum used for image generation (Danas, 1995) Danas, E. (1995). Mapping auditory space onto visual space, unpublished masters thesis, University of Washington, Seattle,WA.

experience. Subjects were also asked if they had a propensity to car- or air- sickness. These people were excused from the test as well as any subjects who felt lightheaded or dizzy during the tests.

Subjects would view one model using the HMD and the next model using the monitor. For each subject that started the battery of test using one device, the following subject would go through the same battery of tests using the devices in an inverse order. For example, subject 1 would view model A using the HMD, model B using the monitor. Subject 2 would then view model A with the monitor, model B with the HMD. This compensated for the variance in difficulty between the model pairs, and for practice effects. The subjects were allowed up to five minutes to explore the models.

Subjects were asked to view each of the six models through the appropriate device and make a floor plan sketch of what they remembered. Next to each wall on the plan, I asked the subject to write down the length of that wall.

Pointing Exercise

In this last stage another test was brought in to use as an indicator of the validity of the sketching exercise. This should not be considered a validity check by any means, but merely an indicator.

In this experiment, the subjects were asked to view a model of a 20’ by 20’ room with seven items in the room: a chair, a table, a five foot tall pyramid, a bookcase, wall shelves, a box with colorful wrapping, and a can on the ground with fire coming out of it. The subjects were given up to one minute to view the model, and in both cases, HMD or monitor, the subject sat in a chair. Those wearing the HMD simply turned their heads and rotated in the chair to view the model. Those using the monitor used the left and right keys on the keyboard. In both cases, the subjects first saw the model facing a wall. They also faced a real wall in the testing room and used the real wall to get orientated to the virtual wall.

After a minute of viewing, the subjects were asked to close their eyes and envision the items’ locations. As I called off a list of the items in the room, the subjects pointed to these items. Having a map of the items’ locations, I estimated the amount of error that each subject experienced. The estimations, however, are problematical. Even with a map in hand, it was difficult to judge with any precision just how much the subject erred when pointing to an object.

I believe that to increase the precision of the pointing experiment, markers would have to be placed around the test area that correspond to the location of the objects in the model. Again because of time limitation, this was not a viable alternative.

Qualitative Comparison Test

A Gestalt-style evaluation was conducted to examine the sketches based not on the measure that was written but instead on the *wholeness* [14] of the sketches. All of the sketches were scanned and loaded into a photo editing program, and then all marking of measurement was removed. These sketches were then shown in random order to four volunteers, two designers and two non-designers. The examiners were asked to

[14] Wholeness here refers to the accuracy of the sketch in proportion as well as presence of all wall elements.

compare the sketch they saw on the screen to a print of the correct floor plan placed next to them. They would then write down a number rating from 1 to 10 (1 being poor, 10 being exact) on how the sketch compared. The examiner would perform this task twice to insure that the judgments of the models would not be affected by the examiner's adaptation to the task.

Gathering and Sorting of Data

Procedures and Problems

Preliminary testing showed that most people could handle the first set quite easily within the five minutes. Half of the preliminary test group handled the second set well, and most people had trouble finishing the last set in only five minutes. This was the basic for the formal tests. This chapter discusses the unexpected performance of the main test group from the preliminary group. The preliminary group composed of five graduate students in the Department of Architecture who worked in the CAD research lab with me.

First Iteration

Most of the subjects were able to remember and sketch the first model set. The second set was more difficult to sketch and included measurements. By the time the subjects got to the third model set, they were becoming discouraged with their abilities and were more prone to give up on the task altogether. This became a major concern and thus triggered a series of attempts to redress the problem.

Second Iteration

To deal with the problems in iteration one, I simplified the models in set 3 and shrank the size of model set 2.

I no longer asked the subjects to give measurements for the model sets 2 and 3, hoping instead that not having to concentrate on measurements, the subjects could expend more energy trying to gain a gestalt cognitive map of the model. Although this eliminated the problem of discouragement, the results were still so prone to error that no useful data could be gleaned.

Part of the advantage of navigation with the HMD is the kinetic feedback. Not only can one turn in the model by turning one's head, but physical orientation of the subject's position, (i.e., the direction he/she is facing) can be used for navigation also. Few subjects were doing this because the cord from the computer to the HMD was very short, thus intimidating the subjects from twisting around. Instead, I found that people were tilting the cyberpuck to the right or left to turn. To encourage them to use their heads and bodies to navigate, I disabled the roll function of the cyberpuck. Although this did encourage subjects to turn their bodies more, many still used their heads for left/right navigation instead of their bodies. The two clear solutions that I could see were not attainable at the time of testing; a longer cord that could be hooked above the subjects' heads, giving them freedom to turn, or a wireless link system.

Third Iteration (Early test stage)

Since the improved results on model sets two and three were still not able to produce useful data, the third iteration saw a change to the procedure and the testing models, to create a set of floor plans of the second model set for a multiple choice task. After viewing each model in the set, the subject was asked to choose which of the floor plans was most accurate, and the subject's choice was written down. The third model set was simplified even more, but the task of sketching a floor plan still remained too difficult for most people.

The results from the multiple choice was inconclusive. The first model of the model was correctly guessed every time, leaving nothing to compare the second model of the model set to compare to.

Fourth Iteration (Late test stage)

This was the final change to the testing procedure, not because many of the difficulties were sorted out, but because of lack of time. By this time it was clear that many unpredictable and confounding variables were not eliminated, and thus could have greatly affected the results of the testing. The only test that remained constant was the sketch task for the first model set. Normally this would be the first set attempted by the subjects. Towards the end of the testing period, I had them perform this last, and skip the third model set altogether. I still had the subjects perform the sketching task on model set two to see if there was any correlation between this task and performance on the first model set, but at the end of testing, I felt that I did not have enough samples to warrant a presentation in this paper. Indicators were that there was not any correlation.

Simulator Sickness

I found that a surprising number of those not used to the computer environment experienced *simulator sickness* [15] from the test equipment. Not expecting such a large number of people unable to perform the tests, I did not track those who became ill using the system. After the testing however, I must conclude that having experience in computer environments is a very important variable to consider when performing VE tests. Even the five minutes given for adaptation time were not enough to help those unfamiliar with the system. Subject attitude toward the experiment also had a significant effect on performance. Most subjects who were excited to see and use the system fared better than those that were more hesitant. Care must be given by the test administrator that those unfamiliar and tentative with using the HMD should be given extra attention to insure that they are not overwhelmed by the experience.

Since simulator sickness occurred using both media, there must have been an element common between the two that could be manipulated to reduce the effect. No one complained of simulator sickness when viewing the static model set used in the pointing exercise, which I believe is attributable to the lack of lateral acceleration. Since this is a variable that cannot be readily eliminated, other causes were sought.

Quake is set to 90 deg GFOV as a default. I changed the GFOV to 60 deg in hopes of a reduction in simulator sickness complaints. There was a sacrifice in the visual field, but this did not seem to adversely affect performance. Although there was a small reduction in complaints, the lateral disorientation still remained an obstacle.

---

[15] Simulator sickness is often used to refer to sickness symptoms that result from an incorrect presentation of a simulation, not sickness caused by a correct simulation of a nauseating experience. (Pausch, Crea & Conway, 1992) Motion sickness is caused by correct presentation stimulus causing nausea and other symptoms.

Sorting the Data

Using Excel, all of the reported volumetric measurements were entered into a database, herein referred to as Vdata, (Volumetric Data) and divided into two categories: Vdata gathered via the HMD (HMD_Vdata) and Vdata gathered via the monitor (Monitor_Vdata). The correct values for the models will be known as the Key values (Key_Vdata).

Gestalt Closure Tolerance

Subjects were asked to sketch floor plans of the models they experienced. By the end of the testing period, only data from the first test model set was still traced with measurements. As a filter, only sketches that correctly represented both of the models in the first model set were considered. Since each subject demonstrated a different level of ability to judge distance, those subjects that tended to overestimate one model most likely overestimated on the second model, thus balancing the average. Those who could only represent one of the two model sets were not included to avoid an imbalance.

The *Gestalt Closure Tolerance* is defined here as the limit of topographical information that can be missing from a sketch yet still be considered a complete image. Sketches that indicate lines for walls, will not have a measure for the width of the wall. An example is shown in Figure 7. This is an exception that would pass the tolerance criteria. If a sketch is missing a wall at any point, then the sketch falls below the tolerance. The Gestalt Closure Tolerance is set high because of the algorithms used for the quantitative analysis.

Of the 32 test sketches, only 16 were able to complete the first model set within the tolerance. Although model difficulty was a contributing factor in the 50% success rate, other factors influenced the success rate which will be covered later in the Assumptions, Discoveries, and Conclusions section.

Average Comparison of Vdata

In the above graphics, each wall in the models is assigned a number. The measure that the subject places for the wall is entered into the database as Vdata for wall #. A value for that Vdata wall # was averaged with all the Vdata wall #'s from the subgroup. Vdata for wall # from all the subjects that used the HMD were averaged together. This result was then compared to the average result of all Vdata wall # entries from the Monitor group. These two Ave_Vdata # were compared to the Key value for wall #. There are 12 walls in model A and the graph (Appendix B) shows the average performance of each group for each wall.

For Example, a subject created the sketch in Figure 15. The Vdata for wall 1 would be 30'. Vdata for wall 2 would be 5', and so on. In Figure 16, another sketch, created by a different subject, shows Vdata for wall 1 as 10' and Vdata for wall 2 as 10' also. Assuming that they were both from the Monitor subgroup, the Ave_Vdata for wall 1 would be (30+10)/2 = 20'. The Ave_Vdata for wall 2 would be (5+10)/2 = 7.5'.

A second data set was created from this information called *Average Vdata Error*. Each wall's error value was calculated from the absolute value of the difference between the key value and the subgroup value.

i.e. Error = ABS (HMD_Vdata - Key_Vdata)

Average Ratio Comparison of Vdata.[16]

Average Ratio Comparison is used for each of the subjects' Vdata, thus giving the ratio of walls based on the scale of the measure reported by the subject and not the measure of the Key. For example, if two different subjects experienced the same space, one may have measured the dimensions to be 10' x 20', and the other measured the dimensions to be 20' x 40'. If we were to judge only on the Vdata given, then the second subject will have an incorrect answer even though the proportion was the same as the first subject's. These ratios were then averaged together according to the test subjects' subgroup similar to the averaging in the liner Vdata Average.

To create this value, the wall values of a local perimeter were added together. Values for opening into other areas were not considered because the value for those opening could correspond to walls that would be used in the calculation for another local perimeter. A ratio of a wall's value to the value of its local perimeter is calculated for each wall. An example of this formula would be:

(a+b+c+d) / a = ratio_Vdata a

(a+b+c+d) / b = ratio_Vdata b

…etc.

The calculation for Average Ratio is the same as with Average Vdata. This holds true for Average Ratio Error to Average Vdata Error.

---

[16] Not to be confused with the 'ratio estimate technique' used in Sadalla's Room Size experiment (Sadalla 1984, Allen 1978).

Assumptions, Discoveries, and Conclusions

Imageability, Occupancy, and Empty Rooms

One of the most difficult choices I had to make was whether I would allow objects in the models or not. Placing objects in an environment brings different levels of imageability depending upon the viewer. There are two ways that people organize information about their environments: linearly and spatially. (Passini, 1984) Linear wayfinding is the sequential following of environmental cues, while spatial wayfinding is more of a Gestalt understanding of the environment topography. To insure that the spatial abilities were tested, objects that could have been used as linear cues were omitted.

The presence of objects can cause an effect I will call '*occupancy*.' Imagine your house or apartment before you moved in the furniture. The space looked smaller, (or sometimes larger), than it did after all of your furniture was in place. How much so depended on what furniture you had and where you put it. Occupancy in the test models could have tricked the subjects into thinking the space was smaller (or larger) than it really was depending on the imageability of the objects. Many environmental designers learn how to avoid this illusion by ignoring furniture when perceiving the room. Since this logical, *concrete operation* is easily swayed by appearance unless practiced, I believed that an empty room would level the test playing field.

Another consideration was follow-up tests. If I allowed objects in the virtual room, the effects of occupancy would be different based on which objects I chose and where those objects were placed. If someone else wished to duplicate this test, they would be forced to use the same objects in the same configuration that I used to eliminate the effects of occupancy. At present, I do not know just how much influence occupancy will have on spatial perception, but like any other unknown variable, it is best to try and limit its influence when possible.

The lack of objects in a model can become disorienting. So much so that a majority of subjects would try to find the "beginning" of the model, or at least where they started. This can be attributed to the lack of any change in wall texture and the lack of landmarks. Subjects would ask, "Where is the beginning?" or "I can't remember where I came in." One person wrote 'start' on her sketch where she entered the model. It is clear that from these questions and other comments that spatial reasoning occurs in a hierarchical fashion (Wayfinding) because of the need to break the task of learning the model into pieces. To deny the subjects of landmarks[17] to orient themselves put a severe limitation on their ability to gain a firm cognitive map of the model. Since some people are better then others at dealing with this lack of information, then somewhere a balance needs to be found for those who have different *wayfinding styles*.[18]

The dilemma of wanting landmarks for orientation purposes, and yet the desire to not have objects in the virtual room because of occupancy effects, seems to be a difficult one. One possible solution would be to have each room mapped[19] with a different texture map, thus helping to subjects remember which room was visited while still avoiding the effects of occupancy.

Sketch Validity

One of many definitions for *Cognitive maps* are "mental models of the relative locations and attributes of phenomena in spatial environments." (Billinghurst and Weghorst, 1995) Since my main interest is to compare the ability of the test subjects to gain a cognitive understanding of the virtual environment, the use of sketching that environment was employed. Mark Billinghurst and Suzanne Weghorst have conducted a test to check for the validity of using sketch maps.

---

[17] Objects become landmarks for two reasons; their distinctiveness and personal meaning (Lynch, 1960).

[18] The assessments of the two groups can be seen to reflect the difficulty each group experienced in finding the information it was looking for. It can therefore be assumed that some people rely on linearly organized information as it is presented by directional signing, while other people tend to rely on spatially organized information that permits an understanding of the setting as a spatial ensemble. The tendency to rely on one type of information more than another marks two distinct wayfinding styles. The first we may call linear, the second spatial.

[19] Each object, in this case a wall, is created by a assemblage of flat triangle elements called faces. These elements combined create the object that is viewed in the model. Each face can have a image assigned to it, whether the image is brick, plaster, or whatever, that will be seen as the faces material. This image is called a *texture map*.

Results show a high positive correlation between subjective ratings of orientation, world knowledge and sketch map accuracy, supporting our hypothesis that sketch maps provide a valid measure of internal cognitive maps of virtual environments.

In Billinghurst and Weghorst's work, they described three ways that their sketch maps were analyzed: Map Goodness, Object Classes, and Relative Object Positioning.

'Map Goodness' is a subjective judgment as to the usefulness of the sketch map as a navigational tool. 'Object Classes' refers to the number of objects in a scene and 'Object Positioning' refers to the object's position relative to other objects in the world.

In the paper, the authors concluded that these sketch evaluation techniques are good for complicated environments but that another technique would better serve testing sparse environments like the ones used in this thesis.

Of course there are pitfalls to the use of sketch maps as explained by Daniel Henry in his thesis.

To begin with, they [sketches] also are a measure of a person's memory of a space. The sketches are drawn after the space has been visited. This means that the accuracy of the sketch is also a function of the variation in people's ability to remember spaces. The memory of a space is not of primary importance in this study. What is of interest is the measure of peoples perception of virtual spaces while they are in them. Furthermore, sketches also measure people's ability to draw. A participant who has a good sense of what a certain place is like, but who has difficulties drawing plan views to scale, will appear to have an inaccurate cognitive map.

In this thesis, memory was an important element. My concern is more the usefulness of spatial information after the fact of communication, but the warnings are still valid.

Interpretation of the test data

I employed two techniques to evaluate sketches within this thesis; qualitative analysis, and quantitative analysis using measurement and ratio. Although measurement has an inherent limitation as discussed earlier, the use of wall to area ratio may serve as a calculation algorithm for evaluating scale. [20]

The basis of this thesis is to examine the potential benefits of examining 3D environments with a HMD vs. the traditional computer monitor. To test for this, I compared the ability of one subgroup to represent their perception of the model using sketches and measurement to another subgroup. I am also interested in exploring the possibilities of quantifying this information so that decisions can be made about the benefits of introducing the HMD into pedagogical and production environments.

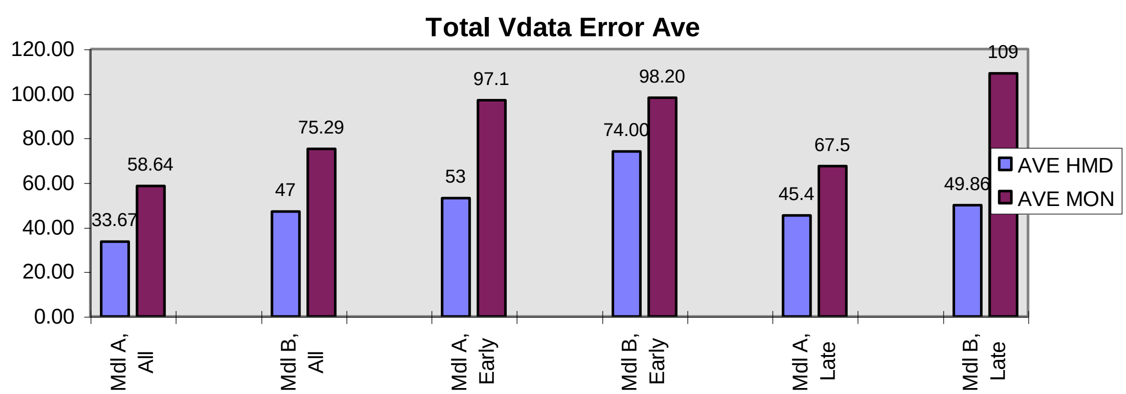

Figure 2 HMD Group

The information on the above graph shows the Average Vdata Error. Graphs for earlier and later iterations are also shown. (See Appendix B)

The amount of Vdata Error generated by the HMD subgroup for model A (33.67) is app. 43% lower than for the Monitor group A (58.64). (*The lower the bar, the lower the amount of average error*) However, when looking at the Total Average Ratio Error, (*below*) the monitor subgroup performed 25% better than the HMD counterpart for model A. At this point in time, I do not have an explanation for this;

[20] I do not know how this compares with the Multidimensional scaling analysis computer program called *Congru*. It was not readily available to me at this time. (Allen 1978).

however, I hope to perform follow-up tests in the future to see if this contradiction repeats itself.

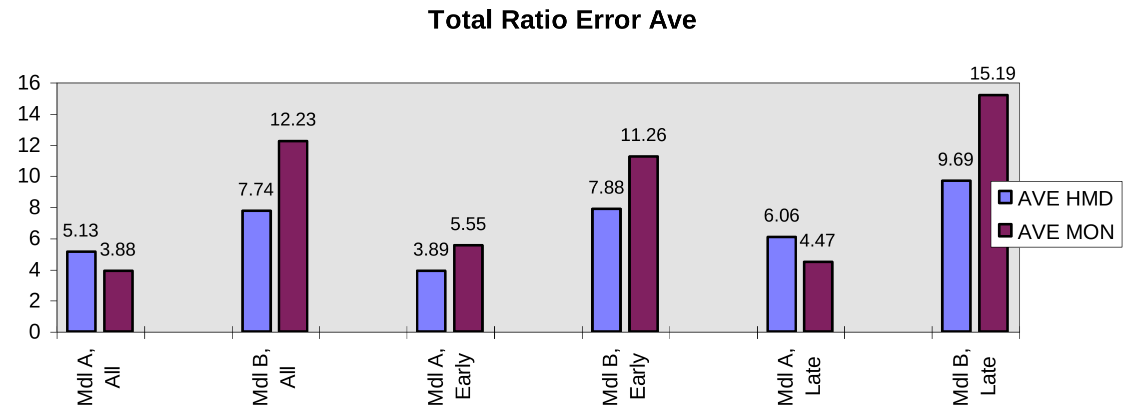

Figure 3 Monitor Group

When looking at the ratio graphs for model A in the early stage vs. the late stage, the difference is significant. The best I can explain for this event phenomenon is that there are only two sketches for the model A late group, thus inducing the greater possibility for localization error.

All other total averages seem to indicate a notable improvement in performance of the HMD subgroup to the Monitor subgroup, from 30% to 40%.

These findings support the basic thesis that spatial cognition, as measured here, improves using the HMD over a monitor.

The Pointing Exercise

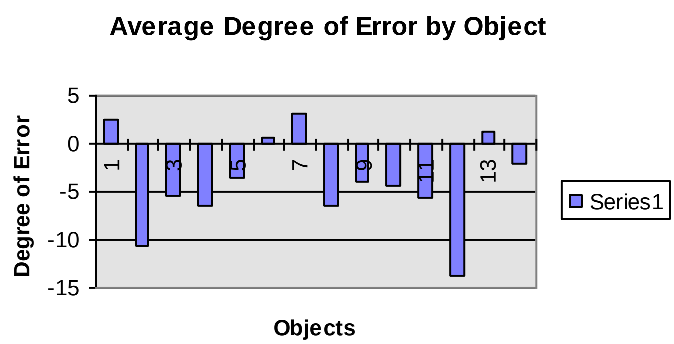

Figure 4 Ave degree of Error for the Pointing Exercise

The pointing exercise enjoyed far better and more consistent results than the sketching task. The total average degree of error for the HMD subgroup was 6.81, and the average degree of error was 11.05 for the Monitor group, resulting in an approx. 38% improvement, correlating with most of the findings from the sketching tests. This finding correlates with Pausch's finding of target location improvements with the HMD at 41%. The Graph (Figure 20) shows the Difference of the Ave Degree of Error between the two subgroups. Bar graphs that are plotted above the zero line indicate that the HMD subgroup's Average Error was larger then the Monitor. Bar graphs that are below the zero line, show the HMD subgroups Average Error as lower then the Monitor group. However, the ability to point to an object is only the ability to remember and locate an object that is relative to the observer. This does not say anything about the quality of the cognitive map the subject possesses. I included this test to see if the percentage of performance improvement would be similar to the sketching exercise, and in this case, the results seem to support the thesis.

Qualitative evaluation

A qualitative (subjective) evaluation was performed on the Gestalt goodness of all the sketches. The results showed little variance (about .5), between the rating (a number from 1 to 10) of the HMD subgroup to the Monitor subgroup. The same result was found when comparing ratings only for those sketches that complied with the Gestalt Closure Tolerance. It seems that this information supports Billinghurst and Weghorst's conclusion that qualitative evaluation for sparse environments may not be the best alternative.

Unforeseen variables in the testing procedure

There are many more important influences in performing cognitive tasks such as these than what I have already covered. I do not wish the following to be considered an exhaustive list by any means, but instead concerns that should be noted and addressed when performing similar perception tests.

Mazes and Subjects' Disorientation

A very common complaint among subjects is the feeling of getting lost in the models. This feeling, even if overcome, affected the subjects' abilities to concentrate on the task at hand. Models should be designed to avoid the adverse effects of occupation without becoming labyrinthine. One recommendation already mentioned is to change the map texture for each room.

'Wow' Factor

Many subjects were quite impressed with the HMD. So much so that they would spend time exploring the abilities of the headset instead of trying to gain an understanding of the virtual environment. I gave the test subjects five minutes to explore a space with the HMD, so they could adapt to navigation and VOR effects. Although this was enough time for most subjects, it is clear that for many novice computer users more time would be needed. A proper amount of adaptation time should be derived according to the needs of the least experienced user.

Subjects' Preconception of What They Are About to See

A common distraction, similar to the 'Wow Factor', is the subjects' preconceived ideas of what they were about to see. Many thought that they would experience fully textured rooms filled with objects that could be interacted with as seen in many Hollywood films. The disappointment of the actual experience tended to reduce the enthusiasm of the subjects to explore the spaces. I would recommend that subjects should be forewarned not to expect a VR experience similar to what they may have seen in films.

Attention to the Subjects

A serious error in test administration was the amount of attention I paid to the subjects as they were taking the tests. After a review of as many of the test subjects that I could contact, I found that there was an even distribution of attention and lack of attention paid across the subgroups, which was purely accidental. In future tests, the amount of interaction with the subject before and during the test must be closely monitored. It is a well documented fact that paying attention to someone performing a task usually improves performance on the task..

Sound

Although spatial sound localization is important in creating a convincing VR experience, one must also remember that sound outside of the VE can serve as a disorienting element when trying to immerse the subject in the Virtual Environment. Radio, other subjects, equipment and any other sound, can be localized by the subject, thus creating conflicting orientation cues.

Suggestions for Further Study

Estimation of Room Size

Although not discussed in the main body of the thesis, a clear trend can be seen in both models A and B, and within both subgroups. For some reason, the averages show that most people under-estimated the room dimensions, particularly with the monitor group. There also seemed to be a trend in the ratio graphs. What is interesting to note, is that the ratio graphs seem to try and smooth out the harshness of the Key ratio graph. This seems to be an effect of symmetry bias in perception(Tversky, 1984).

Another suggestion for further study is to quantify the effects of occupation. This can be done by creating various rooms with no, some, and many objects. Subjects will then be asked to estimate the dimensions of the room. From this data, comparisons of error will show the effect of occupation in virtual environmental perception.

Real World Testing

The experiment done in this thesis looked at visual cognition performance of the HMD vs. the Monitor. To put these results in real-world terms, a third group would need to be created that goes through real rooms. Of course, suitable environments will have to be found or constructed that can be replicated using a computer model.

Impact on Design

Since real-time graphics can offer immediate feedback when designing, a studio can be set up with students working on designs using the HMD as a viewing device. Samples of the design process can be taken and examined to see if the HMD impacts the direction of the design and the impact of the designer's decisions. A second group of designs, working on the same problem will be used to compare again. Evaluations will be purely qualitative of course.

BIBLIOGRAPHY

Alfano, P.L. and Mitchell, G.F. (1990). Restricting the field of view: Perceptual and performance effects. Perceptual and Motor Skills, 72. Pp. 35-45.

Billinghurst and Weghorst. (1995). The use of sketch maps to measure cognitive maps of virtual environments. Institute of Electrical and Electronics Engineers.

Jacob Bronowski. (1978). The origins of Knowledge and Imagination. Yale University Press. New Haven.

Card, S.K. (1985). Windows: Why they were invented, how they help. The Office, March, pp. 52-54.

Downs, Roger M. and Stea, David. (1973). Image and Environment: Cognitive Mapping and Spatial Behavior. Chicago: Aldine Pub. Co.

Eggleston, Robert G., Janson, William P., and Aldrich, Kenneth A. (1996). Virtual reality system effects on size-distance judgments in a virtual environment. Institute of Electrical and Electronics Engineers.

Erickson, Thomas. (1993). Artificial Realities as Data Visualization Environments. In Alan Wexelblat (Ed.), Virtual Reality: Applications & Explorations. Boston: Academic Press Professional.

Fisher, I. (1990). Virtual interface environments. In Brenda Lauren (Ed.), The Art of Human-Computer Interface Design. Reading, Mass: Addison-Wesley Publishing Company.

Ginsburg, Herbert P., and Opper, Sylvia. (1969). Piaget's Theory of Intellectual Development. New Jersey: Prentice Hall.

Held, R. and Durlach, N. (1991). Telepresence, time delay and adaptation. In Stephen R. Ellis (Ed.), Pictorial Communication In Virtual and Real Environments, (2nd ed.). Bristol, PA: Taylor & Francis.

Hughes, Kevin. (1995). From webspace to cyberspace. Enterprise Integration Technologies. Menlo Park, CA.

Lansdown, John. (1994). Visualizing design ideas. In Lindsay MacDonald and John Vince (Eds.), Interacting With Virtual Environments. New York: Chichester.

Lasko-Harvill, Ann. (1993). Interface devices. In Teresa Middleton (Ed.), Virtual Worlds : Real Challenges. Papers from SRI's 1991 Conference on Virtual Reality. Westport, CT: Meckler.

Laurel, Brenda (Ed). (1990). The Art of Human-Computer Interface Design.
Reading, Mass: Addison-Wesley Publishing Company.

Lynch, Kevin. (1960). The image of the city. Series title: Publications of the Joint Center for Urban Studies. Cambridge, Mass: Technology Press.

MacDonald, Lindsay, and Vince, John (Eds). (1994). Interacting with Virtual Environments. New York: J. Wiley & Sons.

Malone, T.W. (1983). How do people organize their desks? Implications for the design of office information system. ACM Transactions on Office Information Systems, 1(1), 99-112.

Mitchell, N. J. (1977). Computer-aided Architectural Design. Van Nostrand Reinhold.

Nasar, Jack L. (winter 1990). The evaluative image of the city. Journal of the American Planning Association, 56, no. 1.

Neisser, Ulric. (1976). Cognition and Reality : Principles and Implications of Cognitive Psychology. San Francisco: W. H. Freeman.

Omen, Charles M. (1990). Sensory conflict in motion sickness: an observer theory approach. In Brenda Lauren (Ed.), The Art of Human-Computer Interface Design. Reading, Mass: Addison-Wesley Publishing Company.

Pagels, Hienz R. (1988). The Dreams of Reason: The Computer and the Rise of the Sciences of Complexity. New York: Simon and Schuster.

Passini, Romedi. (1984). Wayfinding In Architecture. New York: Van Nostrand Reinhold Company.

Pausch, Randy, M. Shackelford, Anne, and Proffit, Dennis. (1993). A user study comparing head-mounted and stationary displays. Proceedings of the IEEE 1993 Symposium on Research Frontiers in Virtual Reality, Pp. 41 - 45.

Pausch, Crea, and Conway. (1992). A literature survey for virtual environments: military flight simulator visual systems and simulator sickness. Presence 1(3) Pp. 344-363.

Proffit, Dennis R., and Kaiser, Mary K. (1993). Perceiving environmental properties from motion, information: minimal conditions. In Stephen R. Ellis (Ed.), Pictorial Communication In Virtual and Real Environments, (2nd ed.). Bristol, PA: Taylor & Francis.

Rapoport, Amos. (1990). The Meaning of the Built Environment: A Nonverbal Communication Approach. Beverly Hills: Sage Publications.

Rheingold, Howard. (1990). What's the big deal about cyberspace? In Brenda Lauren (Ed.), The Art of Human-Computer Interface Design. Reading, Mass: Addison-Wesley Publishing Company.

Snyder, Harry L. (1988). Image quality. In Martin Helander (Ed.), Handbook of Human-Computer Interaction. New York, NY: Elsevier Science Pub. Co.