Mississippi State University

Scholars Junction

12-9-2011

Peripheral visual cues and their effect on the perception of egocentric depth in virtual and augmented environments

James Adam Jones

Follow this and additional works at: https://scholarsjunction.msstate.edu/td

Recommended Citation

Jones, James Adam, "Peripheral visual cues and their effect on the perception of egocentric depth in virtual and augmented environments" (2011). *Theses and Dissertations*. 3513.
https://scholarsjunction.msstate.edu/td/3513

PERIPHERAL VISUAL CUES AND THEIR EFFECT ON THE

PERCEPTION OF EGOCENTRIC DEPTH IN VIRTUAL

AND AUGMENTED ENVIRONMENTS

By

James Adam Jones

A Dissertation
Submitted to the Faculty of
Mississippi State University
in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy
in Computer Science
in the Department of Computer Science and Engineering

Mississippi State, Mississippi

December 2011

PERIPHERAL VISUAL CUES AND THEIR EFFECT ON THE

PERCEPTION OF EGOCENTRIC DEPTH IN VIRTUAL

AND AUGMENTED ENVIRONMENTS

By

James Adam Jones

Approved:

_____
J. Edward Swan II
Associate Professor of
Computer Science & Engineering
(Major Professor)

_____
T.J. Jankun-Kelly
Associate Professor of
Computer Science & Engineering
(Committee Member)

_____
Carrick C. Williams
Associate Professor of Psychology
(Committee Member)

_____
Philip Amburn
Research Associate Professor
Geosystems Research Institute
(Committee Member)

_____
Edward B. Allen
Associate Professor of
Computer Science & Engineering,
and Graduate Coordinator

_____
Sarah A. Rajala
Dean of the James Worth Bagley
College of Engineering

Name: James Adam Jones

Date of Degree: December 9, 2011

Institution: Mississippi State University

Major Field: Computer Science

Major Professor: Dr. J. Edward Swan II

Title of Study: PERIPHERAL VISUAL CUES AND THEIR EFFECT ON THE PERCEPTION OF EGOCENTRIC DEPTH IN VIRTUAL AND AUGMENTED ENVIRONMENTS

Pages in Study: 99

Candidate for Degree of Doctor of Philosophy

The underestimation of depth in virtual environments at medium-field distances is a well studied phenomenon. However, the degree by which underestimation occurs varies widely from one study to the next, with some studies reporting as much as 68% underestimation in distance and others with as little as 6% (Thompson et al. [38] and Jones et al. [14]). In particular, the study detailed in Jones et al. [14] found a surprisingly small underestimation effect in a virtual environment (VE) and no effect in an augmented environment (AE). These are highly unusual results when compared to the large body of existing work in virtual and augmented distance judgments [16, 31, 36–38, 40–43]. The series of experiments described in this document attempted to determine the cause of these unusual results. Specifically, Experiment I aimed to determine if the experimental design was a factor and also to determine if participants were improving their performance throughout the course of the experiment. Experiment II analyzed two possible sources of implicit feedback in the experimental procedures and identified visual information available in the

lower periphery as a key source of feedback. Experiment III analyzed distance estimation when all peripheral visual information was eliminated. Experiment IV then illustrated that optical flow in a participant's periphery is a key factor in facilitating improved depth judgments in both virtual and augmented environments. Experiment V attempted to further reduce cues in the periphery by removing a strongly contrasting white surveyor's tape from the center of the hallway, and found that participants continued to significantly adapt even when given very sparse peripheral cues. The final experiment, Experiment VI, found that when participants' views are restricted to the field-of-view of the screen area on the return walk, adaptation still occurs in both virtual and augmented environments.

Key words: dissertation, virtual environments, virtual reality, augmented environments, augmented reality, mixed reality, perception, vision, locomotion

DEDICATION

To my parents, James & Jane Jones, and to all those without whom this would not have been possible.

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF FIGURES

ABBREVIATIONS

AR Augmented Reality

AE Augmented Environment

MR Mixed Reality

VR Virtual Reality

VE Virtual Environment

HMD Head-Mounted Display

OST Optical See-Through

VST Video See-Through

NASA National Aeronautics and Space Administration

FOV Field of View

SPAAM Single Point Active Alignment Method

CRT Cathode Ray Tube

LCD Liquid Crystal Display

IS-1200 InterSense IS-1200 VisTracker Motion Tracking System

CHAPTER 1

INTRODUCTION

Augmented Reality, in its most general sense, is the ability to enhance or modify individuals' perception of their surroundings. Though these modifications can include visual, olfactory, tactile, or auditory augmentation, for the purposes of this document, Augmented Reality (AR) will refer specifically to visual stimulation presented through an optical see-through head mounted display. In AR, observers retain their view of their surroundings while virtual, computer generated elements are added to this view. The degree of augmentation can vary from simply adding textual information to a scene to completely replacing an observer's view of their surroundings. The latter case is more commonly known as Virtual Reality.

Though Virtual Reality (VR) and AR are similar technologies, VR is somewhat more thoroughly studied, as it has historically been less technically complicated to implement. This is largely due to the complexities of accurately matching and merging the views of the virtual and real worlds in AR. VR has also been seen as an interesting research tool, enabling researchers to present observers with stimuli which would otherwise be impractical or unsafe in reality. This raises the question, are stimuli presented in Virtual Reality actually comparable to those presented in the real-world? In an attempt to answer

this question, much research has focused on understanding how accurately observers can make depth and layout judgments in VR as compared to the real-world.

It is also important to note that both AR and VR have numerous other applications. These include AR/VR assisted surgery, post traumatic stress syndrome exposure therapy, guided navigation, combat training, air traffic awareness, entertainment, and many more. For tasks such as these to be usefully applied, a deeper understanding is needed of how observers perceive augmented and virtual imagery.

Depth judgments in VR have been widely studied, and observers have historically reported that the perceived position of computer generated imagery is not congruent to that of co-located real-world objects. Specifically, observers tend to view virtual environments as compressed relative to their actual geometric size. This compression leads to an underestimation of the size of virtual spaces. Studies have reported as much as 68% compression of depth judgments in virtual environments [38]. Though the exact sources of these misperceptions are largely unknown, it is likely to be a combination of several factors, running the gamut from inadequate calibration to limitations in display technologies.

CHAPTER 2

PREVIOUS WORK

2.1 Master Thesis

The work described in this document is an extension of that discussed in the author's Master Thesis, entitled "Egocentric Depth Perception in Optical See-Through Augmented Reality" [13]. The goal of this work was to determine if the well-studied underestimation of distances found in virtual environments also existed in augmented environments. This thesis was the first to perform a direct comparison of depth judgments between congruent augmented, virtual, and real-world environments. Additionally, this work introduced a novel optical see-through calibration method that required no prior training. This method was referred to as the 3D Compass. The major findings from the two experiments discussed in the thesis were detailed in Jones et al. [14] and Swan et al. [36].

The first experiment, described in Swan et al. [36], was the first experiment of its kind to utilize blind walking, a common distance judgment protocol when studying virtual environments, in combination with an augmented environment. This experiment also collected distance judgments via verbal report. It was found that verbal report lacked consistency between observers, but judgments made using blind walking were much more stable across the group of participants. Four viewing conditions were used for this experiment: real-world, real-world viewed through an HMD, real-world with superimposed

augmentation, and a purely virtual stimulus. All four environments were viewed from a stationary, non-tracked optical see-through head-mounted display. The stimulus used in all conditions was a wireframe pyramid measuring 23.5cm along the base and 23.5cm in height.

The results of the first experiment demonstrated significant underestimation in all environments with all conditions being significantly more underestimated than the real-world stimulus. One of the more interesting findings of this experiment was that a real-world stimulus viewed through an HMD did not significantly differ from either of the augmented conditions. This result began to hint that the cause of the underestimation may not lay in the stimulus but in the manner in which it is viewed. Viewing of the real-world stimulus was the only condition where observers' head movements were not restricted by looking through the rigidly mounted HMD. Suspecting that these results may be an effect of motion based cues not available during fixed viewing, a second experiment was formulated.

The second experiment, described in Jones et al. [14], aimed to determine whether or not the addition or restriction of motion parallax as a depth cue would influence the trends observed in the first experiment. To test this, observers saw two parallax conditions: Still and Motion. In the Still condition, observers were instructed to observe the stimulus while standing as still as possible. In the Motion condition, observers were asked to observe the stimulus while swaying from side to side. Since the verbal reports in the first experiment lacked consistency, only blind walking judgments were collected in the second experiment. This experiment also incorporated a purely virtual environment as one of its viewing conditions.

The results of this experiment revealed no consistent effect of motion parallax. Contradictory to the first experiment, the augmented environment showed no significant difference from the real-world viewing condition. The distance judgments in the purely virtual environment did, however, exhibit a significant underestimation compared to distance judgments in the real-world. A surprising result was that the underestimation in the virtual environment was markedly less than has been previously reported in a wide range of related studies [1, 4, 11, 16, 25, 29, 37, 38, 40, 42, 43]. The work described in this document aimed to determine the cause of these unusual and conflicting results.

2.2   Related Work

The appeal of virtual and augmented environments is that they can provide observers with views and information that may not be possible in the real-world. This has been of particular interest in the field of psychology, where observers can be placed in tightly controlled artificial environments for various experimental or therapeutic purposes. Virtual environments, in particular, have been successfully used for exposure based therapies to treat conditions such as acrophobia and post-traumatic stress syndrome in a safe and monitored fashion [33].

However, using virtual and augmented environments in this manner relies on the assumption that these artificial environments are, in fact, analogous to the real-world. There is some behavioral evidence to indicate that observers do, to some degree, perceive real and artificial environments similarly. There is even some neurological evidence indicating that mice utilize the same areas of their brains for processing spatial relationships in both

real-world and computer generated mazes [9]. Studies such as this seem to indicate that there is some basic comparability.

Though virtual environments and the real-world may be seen as roughly similar, numerous studies have found that they are far from equivalent. In particular, a large body of work exists that has thoroughly studied how observers perceive egocentric depth judgments in immersive virtual environments. Though the results of these studies vary, they almost uniformly describe perceived distances in virtual environments as being compressed relative to the real-world [1, 4, 11, 16, 25, 29, 37, 38, 40, 42, 43]

Work has been done that has shown that observers can adapt to a virtual environment by interaction and navigation to the point that they accurately judge spatial relationships [30]. However, this seems to be an issue of the observers calibrating their movements to suit the compressed environment. A related study has shown that by adjusting the scale of the virtual environment through magnification that observers were able to more accurately judge distances [35]. Unfortunately, this approach seems somewhat insufficient as the virtual environment is no longer geometrically congruent to the real-world. This is an incredibly important factor when dealing with augmented environments where virtual elements are added to an otherwise unaltered view of the real-world environment. Though both of these studies provide insight into means of compensating for the depth underestimations in strictly virtual environments, they are only addressing symptoms of some underlying cause.

Foley et al. [6] discuss the perception of location and extent as a means of describing the geometry of visually perceived space. Foley describes a nonuniform transformation

between perceived space and physical space which corrects for inconsistencies between a strictly Euclidean correspondence between perceived and physical space. This model accounts for the apparent tendency for observers to overestimate extents while underestimating locations. This model centers around the concept that observers tend to perceive their effective visual angle as being greater than their physical visual angle.

It is also important to note that Foley et al. [6] make the distinction between location and extent estimations and that underestimation is found in one while overestimation is found in another. This is extremely similar to the results reported in publications by Lappe et al. [17]. However, Lappe et al. make no distinction between location and extent and consider them both to be conflicting measurements (not differing perceptual phenomenon) of the same general distance judgment. Foley's model of perceptual space could answer questions posed by their research [6]. It might be a possibility that Lappe et al. were measuring two different perceptual phenomenon that are related to distance but exhibit different biases. This has been personally suggested to the authors.

Though studies such as these can provide a theoretical model of how these mispercep-tions behave, they provide little insight into their causes or means of mitigation. Cutting, however, places strong emphasis on the idea that the combinations of and fidelity by which depth cues are presented may be a significant factor in these misperceptions [3]. Hu et al. [10] found that observers increase their accuracy in estimating the position of virtual surfaces as shadows and interreflections were added. Phillips and Interrante [29] found that by removing cue fidelity from an otherwise photorealistic environment, observers performed distance judgments with greater underestimation. However, a similar study by

Thompson et al. [38] found conflicting results, indicating that photorealistic fidelity had no effect on distance judgments in an immersive virtual environment. Restriction of visual information even in the real-world has been shown to cause underestimation similar to that seen in virtual environments. Wu, Ooi, and He demonstrated that by restricting an observer's field-of-view that they could modulate the degree by which distances would be underestimated [44].

Substantially less work has been done to determine if the compression seen in virtual environments exists in augmented environments as well. The work that has studied augmented environments has been limited and somewhat inconsistent. For instance, Livingston et al. [19] found a tendency to overestimate distances in an outdoor augmented environment. One study by Swan et al. [37] found that observers exhibited no significant underestimation when performing a perceptual matching task in an indoor augmented environment. Another study by Swan et al. [36] found up to 21% underestimation of distances when performing visually directed walking in an augmented environment [36]. However, in a follow-up study to Swan et al., Jones et al. [14] found no significant difference between visually directed walking judgments in a congruent real-world and augmented environment. These conflicting results leave many questions still to be answered about how distance judgments work in augmented environments.

CHAPTER 3

EQUIPMENT & EXPERIMENTAL SETUP

3.1 Head-Mounted Display

Possibly the most crucial piece of equipment required for the research discussed in
this document is the head-mounted display (HMD). HMDs are a common display device
used to present users with immersive virtual environments. A distinct advantage of HMDs
over other immersive VR display devices, such as Cave Automatic Virtual Environments
(CAVEs), is that they can offer users a wide range of movement. When using most HMDs,
users' movement is typically only limited by the length of the data and power cables that
drive the device and the effective range of any associated motion tracking devices. This
flexibility has made HMDs very popular for research that requires immersive virtual en-
vironments. However, for presenting augmented environments where both real-world and
virtual elements are combined, there are typically two classifications of HMDs: Video
See-Through and Optical See-Through.

3.1.1 Head-Mounted Display Technologies

Video See-Through HMDs (VST-HMD), depicted in Figure 3.1, work by overlap-
ping computer generated images with camera video-feeds that are approximately aligned
with the observer's eyes. Though this method is somewhat easy to implement and can be
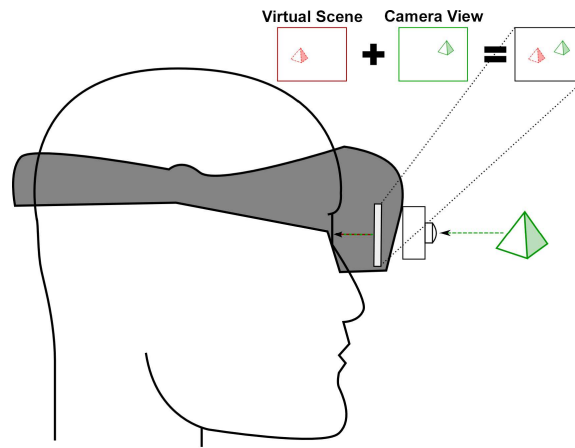
Figure 3.1

A typical video see-through HMD

achieved by attaching cameras to any HMD capable of displaying a virtual environment, real-world visual information presented to the observer is restricted by the technical limitations of the display elements and cameras. For instance, the presentation of the real-world scene is optically limited to the focal distance of the HMD's display elements and the combined fields-of-view of the camera and the display elements. Additionally, matching the cameras' positions and alignment to approximately that of the observer's eyes can be quite difficult.

Optical See-Through HMDs (OST-HMD), on the other hand, preserve a direct optical path from the observer's eyes to the real-world. By preserving this view, the real-world scene is presented to the observer in full resolution and with little loss of visual information. The virtual components of the scene, however, are still generated in the OST-HMD's display elements and suffer the same limitations imposed by display elements' design and internal optical path. The method by which an OST-HMD combines views of the real
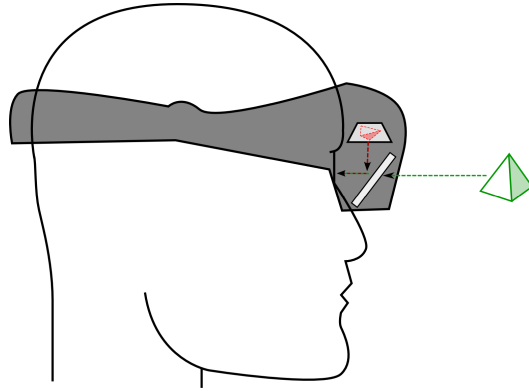
Figure 3.2

A typical optical see-through HMD

and virtual elements is referred to as optical combination. Optical combination preserves a direct optical path from the observer's eyes to the surrounding real-world environment and the OST-HMD's display elements using an optical combiner. The optical combiner is typically either a half-silvered mirror or two conjoined resin-gap prisms. A common arrangement is to have the display elements located above the observer's eyes, facing downward to an optical combiner positioned at 45° to the observer's eyes. This optical combiner allows light from the real-world to directly pass through while partially reflecting light from the display elements toward the observer's eyes. Figure 3.2 depicts this basic arrangement.

3.1.2 NVIS nVisor ST60 Optical See-Through HMD

The research discussed in this document focuses exclusively on Optical See-Through HMDs and, for brevity's sake, will simply refer to them as HMDs. The HMD used for the following experiments is an NVIS nVisor ST60 Optical See-Through Head-mounted
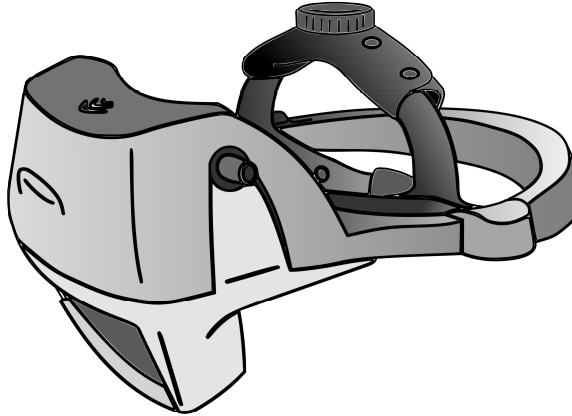
Figure 3.3

NVIS nVisor ST60

Display, see Figure 3.3. This HMD boasts 100% overlap of the real-world and display area with a diagonal field-of-view (FOV) of 60°, horizontal FOV of 48°, and a vertical FOV of 40°. The ocular separation, or interpupilary distance (IPD), is adjustable and spans the range of 53mm to 73mm.

3.2   InterSense IS-1200 Motion Tracking System

The head tracking system used for the experiments described in this document was an InterSense IS-1200 VisTracker, as depicted inFigure 3.4. The IS-1200 is a 6-degree-of-freedom tracker that utilizes both optical and inertial information to provide real-time positional information. In these experiments, an IS-1200 was attached to the head-mounted display in order to properly model the observer's forward view direction. The tracking values returned by the IS-1200 consists of three translational (X, Y, and Z positions) and three rotational components (roll, pitch, and yaw).
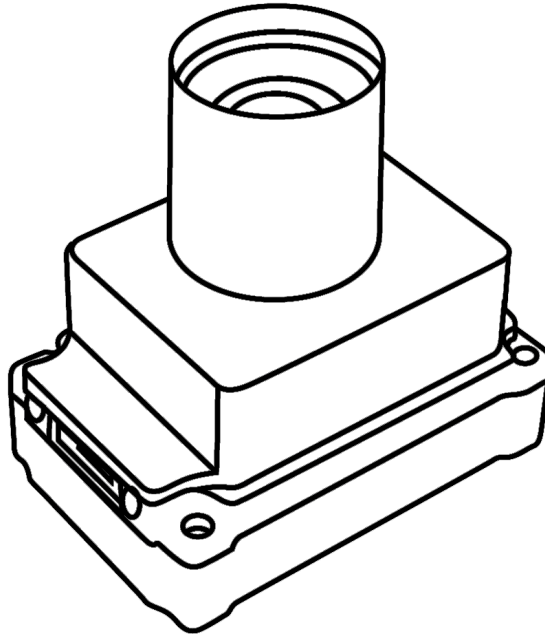
Figure 3.4

InterSense IS-1200 VisTracker [12]

The IS-1200 uses a series of preprogrammed circular markers, referred to as fiducials, that define the area over which the system can provide data. The arrangement of the fiducials is known as a constellation. The constellation used in this experiment was attached to a ridigly mounted board that was suspended from the ceiling and hung directly behind the observer's head. Figure 3.5 shows this arrangement. This allowed the constellation to be easily moved and stored when experiments were not being performed.

3.3 Nonius Apparatus

When deprived of all visual stimulation, an individual's eyes typically default to a resting state. There has been some evidence to indicate that this resting state may affect
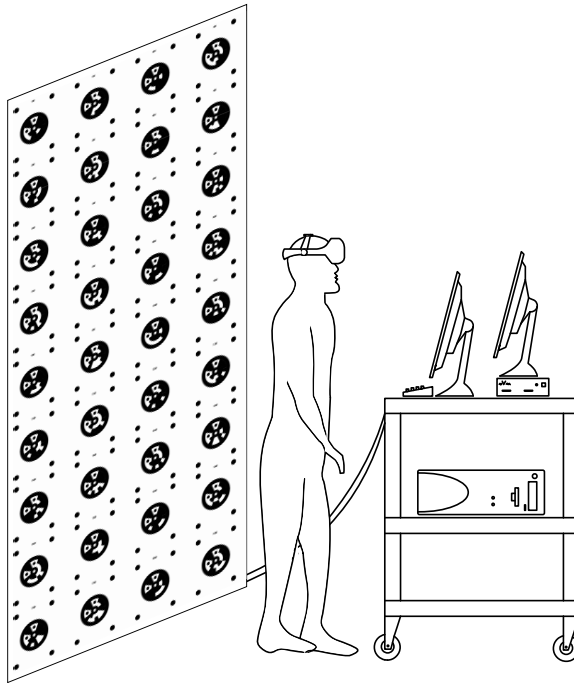
Figure 3.5

Tracking constellation configuration
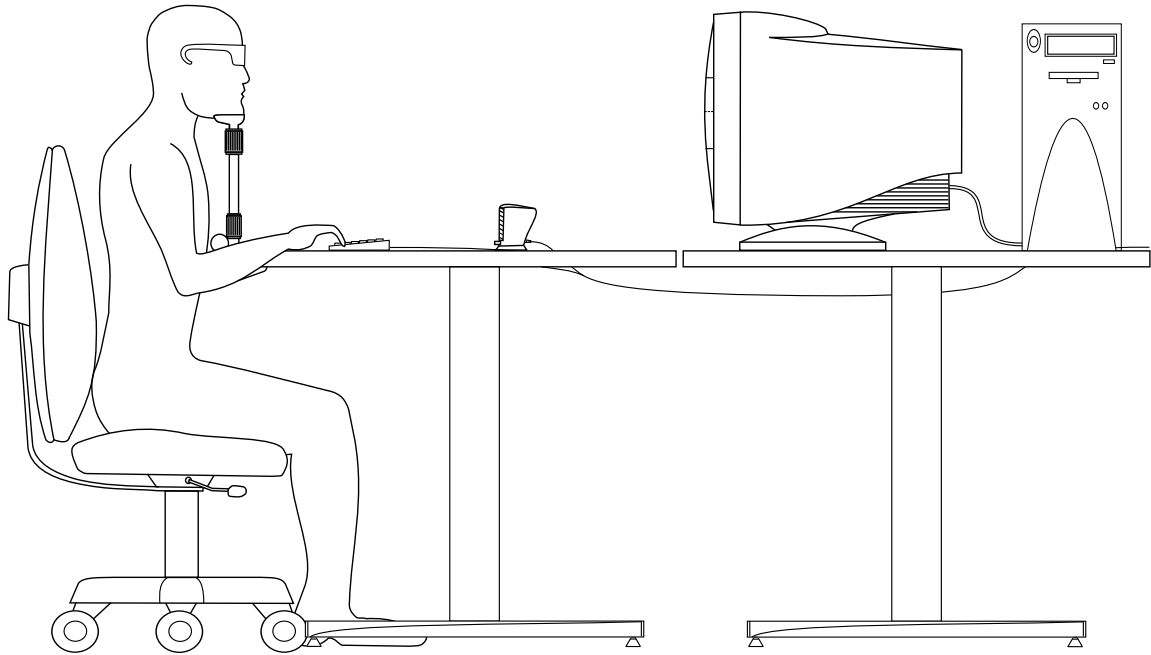
Figure 3.6

Nonius apparatus

visual perception [8, 22–24, 27]. The angle at which the eyes point while resting is referred to as dark vergence. This section describes the design of an experimental apparatus built to measure dark vergence for the purposes of investigating a possible connection with medium field distance judgments in augmented and virtual environments.

A well established method of measuring dark vergence is by performing an Nonius alignment task in the absence of all other visual stimulation. Capturing this measurement proved to be a nontrivial task. Previous work done by Miller [21–23] and Owens and Leibowitz [27] laid out a basic framework for constructing a general purpose Nonius alignment apparatus. The following section will describe the Nonius alignment apparatus

constructed for the current study. The design of this apparatus is loosely based on that described by Miller [21].

The Nonius alignment apparatus measures a participant's vergence by dichoptically presenting flashing stimulus lines to each of the participant's eyes, with one eye's stimulus above the other. In the case of our device, the right eye's stimulus was presented above the left eye's stimulus. The top line is kept at a constant position while the participant had the ability to reposition the bottom line. It is important to note that the lines must be presented in flashes no longer than 350ms to avoid activating the accommodative reflex and there by altering the participant's convergence through the accommodative-convergence reflex. The participant is given the task of adjusting the position of the bottom line such that it appears to be perfectly aligned with the top line.

The stimuli were presented on a 19 inch ViewSonic G790 CRT monitor at 1280x1024 resolution and 70Hz vertical refresh. In order to enable dichoptic presentation of the stimuli, linear polarizing filters were applied to the upper and lower halves of the screen's surface with the upper and lower filters differing in polarizing angle by 90°. A pair of cardboard frame, polarizing filter glasses were constructed in order to enable dichoptic viewing of the stimuli on the monitor. The polarizing filters in the glasses were oriented such that the right and left eyes corresponded to the upper and lower portions of the screen, respectively. The brightness of the monitor was adjusted to a very low level to prevent the glow associated with the ambient phosphor excitation of the dark pixels. This ambient glow could act as a cue to the location of the monitor relative to the participants, which could bias the dark vergence measurement. However, even with the brightness at a min-

imal level, the edges where the screen met the monitor housing were very apparent. To mitigate this effect and prevent unnecessary ambient illumination, aluminum foil strips were applied to the screen edges in order to completely block the visibility of the edges. Layers of black masking tape were applied to the edges of the aluminum foil that were internal to the screen area. The layers of tape provided a diffusing effect which caused the screen's ambient glow to fade as it approached the completely black edge of the aluminum foil. This allowed only the necessary portions of the screen to be visible while still ambiguating the monitor's position.

It is important to note the reasons for using this particular CRT monitor, as opposed to a more common LCD monitor. Typical desktop LCD monitors have several shortcomings which prevented them from being used for the Nonius apparatus. Unlike CRT monitors, which excite individual sets of RGB phosphors to illuminate a pixel, LCD monitors use a white back-light which is directed through adjustable liquid crystal filters that modulate the amount of red, green, and blue light that is transmitted through the LCD panel. The typical LCD monitor's ability to block the back-light when displaying black pixels is somewhat limited, causing them to have much brighter black levels than CRT monitors. The brighter a monitor's black level, the easier it is to localize the monitor's surface in the dark.

In order for the measurement of dark vergence to be successful, certain other issues had to be addressed in order to help ensure accurate measurement. Firstly, participants need to be deprived of visual stimulation in near absolute darkness in order for their eyes to remain in a resting state. The room used for this task was an internal, windowless room, but it still required further conditioning in order to be sufficiently dark. Even though the

room had no windows, a substantial amount of light could enter from around the room's only door, even when closed. Enough light could enter the room from these seams that one could easily navigate the room after only a few minutes of adaptation to the darkness. This was unacceptable for the purposes of this experiment. To solve this problem the hinged side of the door was completely sealed with an aluminum foil seam which was taped to the door and to the connecting wall. An aluminum foil lip was taped to the opposite side and top of the door. When the door was closed magnets were used to hold the seam against the door's metal frame. The bottom of the door was sealed using an L-shaped aluminum bar which slid underneath the door and spanned its width. Aluminum foil was an excellent material for sealing the room since it was inexpensive, readily available, and has no light conducting properties.

Once light seepage from the doorway had been resolved, it became apparent that there was light seepage from small gaps around the ceiling tiles and light fixtures in the drop-ceiling of the laboratory. The source of the light were attic lamps and ambient light from adjacent offices which seeped upward from the gaps between their ceiling tiles and were reflected back downward from the upper ceiling. The seepage was not sufficient to allow one to navigate the room, but was sufficient to enable one to distinguish dark and light surfaces. This was resolved by taping aluminum foil strips to the largest ceiling gaps and disabling the attic lamps.

There was one final obstacle in light proofing the room, which was to lightproof the technology required for the experiment. In much current technology, there exists an abundance of LEDs and other light sources that are used to communicate various system infor-

mation or to simply for aesthetic purposes. The computer used for this experiment was a Dell XPS 730, which was bristling with LEDs. Though its LEDs could be disabled by a software application, they were automatically re-enabled after each reboot. For simplicity, the LEDs were physically removed from the computer. However, many other devices in the lab, including monitors, printers, and telephones, had light sources which were not as easily removed. To block these light sources, aluminum foil was taped to all light emitting surfaces. Once these tasks had been completed, the room was dark enough that after as much as 20 minutes of dark adaption one could not visually detect the presence of their hand in front of their face.

The participants interacted with the Nonius apparatus via keyboard. Only three actions were necessary for the participants to interact with the system: left adjust, right adjust, submit response. These functions were mapped to three keys on the keyboard. Since this portion of the experiment took place in the absence of light, unique physical textures were adhered to each of the three keys so they could be easily identified without vision.

Figure 3.7 shows the stimuli presented on to the participants. The upper stimulus, presented only to the right eye, remained stationary in the horizontal center of the screen while the participants adjusted the position of the lower stimulus, presented only to the left eye. The participants were then tasked with adjusting the lower stimulus until it appeared to be aligned with the top stimulus, as depicted in Figure 3.8. Though the participants controlled the directional movement of the lower stimulus, they did not directly control the distance increment that it moved in either direction. The movement increments were adjusted based on a distance bracketing procedure that analyzed the adjustment patterns
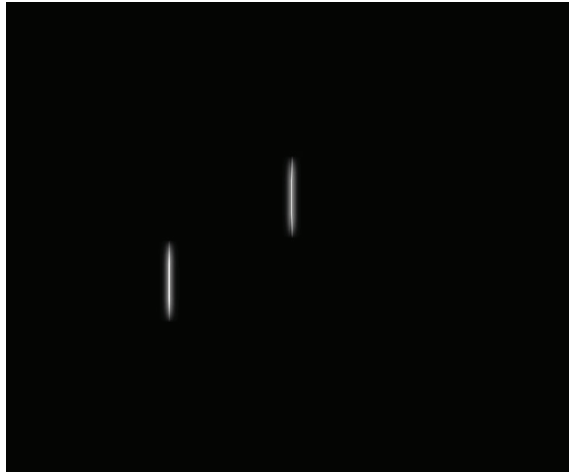
Figure 3.7

The stimuli presented by the nonius device
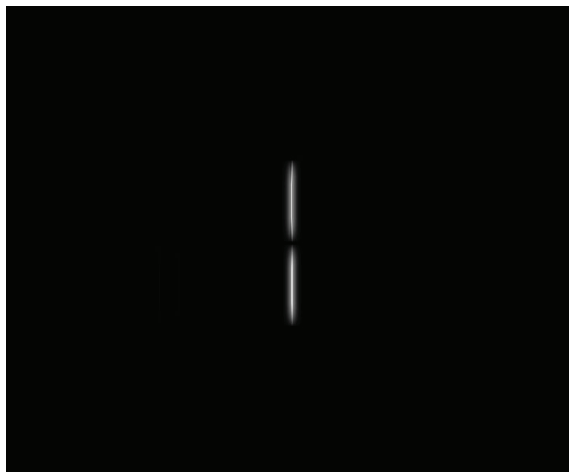
Figure 3.8

Aligned stimuli on the nonius device

of the participants. The software used to present the stimuli attempted to detect whether the participants were making adjustments that were either honing in on the subjective target position or correcting for the previous adjustment. It did so by looking for alternating direction changes in sequential movements. For instance, if a participant adjusted the stimulus left and then right, this indicated that the participant was honing in on the subjective target position. However, two consecutive adjustments in the same direction indicated that a participant was correcting for a previous adjustment. In order to avoid trapping the participants' movements in local minima, inescapable brackets, the adjustment increments were altered based on whether the participant appeared to be correcting or honing. The increment size decreased by half when honing movements were detected. However, the increment size doubled when corrective movements were detected, allowing the participant to move beyond the previous positional bracket. This method allowed the participants to rapidly adjust the stimulus to the subjective target position.

CHAPTER 4

EXPERIMENTAL PROCEDURES

4.1 Environment Calibration

One of the most important prerequisites for performing meaningful perceptual exper-
iments in virtual and augmented environments is the proper modeling of the relationships
between the observers' eyes and the display plane on which the graphics are displayed.
The calibration techniques vary based on the display technology used for a given environ-
ment. The research detailed in this document focuses exclusively on head-mounted virtual
and augmented environments, thus the calibration procedures discussed in this chapter will
be exclusive to those applicable to HMD-based environments. The ultimate goal of a cali-
bration procedure is to, as closely as possible, estimate the parameters of the optical system
used to produce graphical signals (the head-mounted display) and the parameters of the
optical system used to receive these graphical signals (the observers' eyes). Often, HMD
manufacturers provide detailed specifications of the optical characteristics of the displays
they sell. Though these specifications are typically close to those of the actual hardware,
they apply more generally to the display model as a group with individual displays varying
somewhat from the provided values. For applications that require only a loose correspon-
dence between the real-world and graphical augmentation, these values may suffice. For
other applications where exact correspondences are required, the variations between in-

dividual displays within a model series render the general specifications less useful. In some instances, manufacturers will provide "build reports" with exact measurements of the display's parameters as measured at its assembly. These values are very useful and, from the experiences of the author, can be sufficient for sub-centimeter level accuracy. Unfortunately, the display's parameters often change with time as a result of repeated use and repair procedures. This drifting makes calibration of the display parameters an on-going process throughout the life of the HMD. These parameters, however, only describe half of the optical system involved in displaying virtual and augmented environments. The yin to the HMD's yang is the eye of the observer. Humans are very biologically diverse creatures and can exhibit large variations in eye-height and interpupillary distance from one person to the next. In circumstances where a display device may have multiple users, the device may need to have a unique calibration for each person. The bimodal nature of the optical systems involved in viewing a virtual or augmented environment presents an interesting problem for display calibration: how can one easily measure the parameters of these independent optical systems? One answer to this question is to consider the display and the observer as a single optical system where the end result is a proper projection on the observers' retina. This approach is referred to as a single phase calibration. Another approach is to measure and model the optical systems separately. This is referred to as a two phase calibration. Though each of these methods are valid approaches to the calibration problem, each has their limitations. Much work has been done with regard to making the calibration procedure as easy as possible for both researchers and end-users, but a

comprehensive method for all potential uses is illusive. Generally speaking, calibration techniques can be broken down into two groups: single phase and two phase methods.

Single phase methods involve having the observer perform a series of alignments between virtual and real-world markers. By motion tracking the position of the observers' gaze direction during the alignment tasks, the parameters necessary to display an accurately projected virtual environment can be estimated. A single phase method takes the elegantly simplistic approach of combining the measurements of the ocular and HMD parameters into one unified method. This typically involves having the observer, while wearing the HMD, perform a series of bore-sighting alignments between real-world and virtual markers. By performing these alignments, the eye-to-virtual environment projection can be estimated, providing an increasingly accurate projection as more alignments are performed. A very common single phase approach used for HMD calibration is known as the Single Point Active Alignment Method, or SPAAM for short [39]. SPAAM-like methods generally provide robust results, but this method requires the observer to have training in the calibration procedure prior to using the system. Additionally, observers are required to perform many alignments to build up sufficient data to estimate the optical parameters.

Two phase calibration methods take a somewhat more complicated approach by measuring the HMD and observer as two separate optical systems. This is done by treating the HMD as a static optical system that does not change over time. Though drift in the HMD's optical parameters does occur with use, changes to the optical parameters are generally very small unless the optical elements undergo maintenance or suffer significant abuse. This changes the nature of the HMD portion of the calibration process from a

per-use procedure to an occasional procedure that prevents the optical parameters from drifting significantly from those modeled in the virtual environment. The only parameters necessary to be measured on a per-use basis then becomes the parameters related to the observer's eyes. These are specifically the interpupilary distance (IPD), sometimes referred to as ocular separation, and principle ray, the vector representing the ray that passes through the optical center of the eye and the center of the HMD's display elements.

Though the two phase approach is somewhat more complicated than the single phase approach, it has two major advantages. Firstly, the observer typically needs no training in the calibration procedure. Secondly, the observer spends very little time seeing virtual or augmented elements prior to the actual experiment. This is especially important for experiments, such as those discussed later in this document, that aim to measure the effect of exposure to virtual or augmented environments. It is for these reasons that a two phase calibration method was developed for this research.

4.1.1   Phase I: HMD Calibration

The first phase of the calibration measures the optical properties of the HMD itself while the second phase calibrates for properties that change on a per-observer basis, similar to that described by Owen et al. [26]. During the first phase, we measure the following properties: 1) field of view, 2) principle ray, 3) optical distortion. Firstly, the HMD was rigidly mounted in a scaffolding on an optical workbench in such a way as to allow for small rotational movements in roll and pitch. The mounting scaffolding used for the experiments described in this document can be seen in Figure 4.1. Yaw adjustments were not
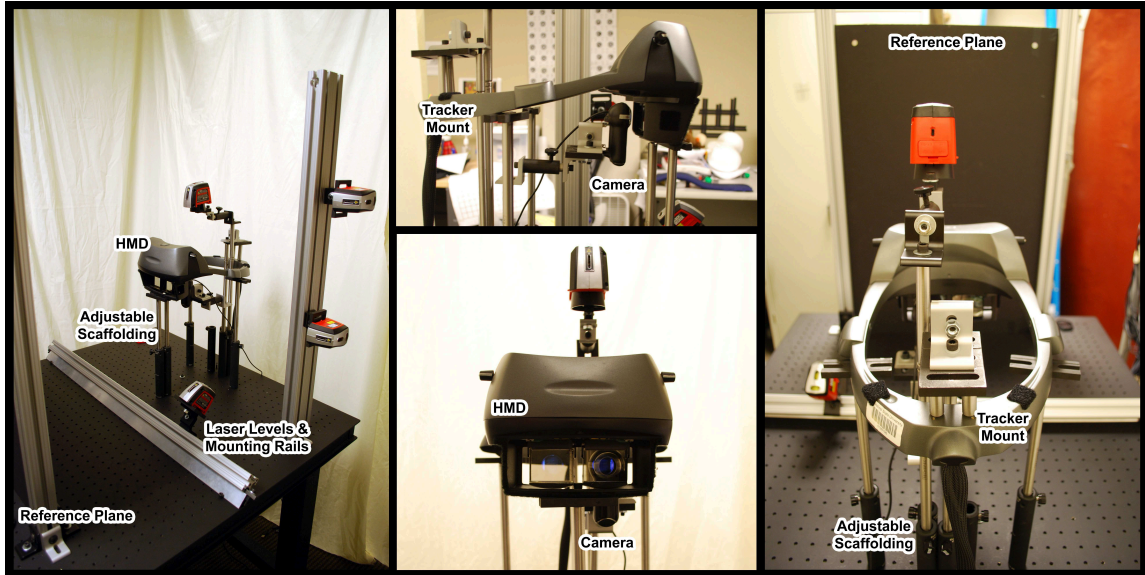
Figure 4.1

Scaffolding constructed for HMD calibration

necessary as the precision grid of threaded ports were used to tightly control the HMD's yaw and placement of all real-world references with millimeter or better accuracy. The optical workbench, on which this phase of the calibration was performed, was equipped with locking, leveling, vibration resistant casters. Using a bubble level, the casters were adjusted such that the workbench surface was leveled perpendicular to the direction of gravity. The casters were then locked into position to prevent movement and ensure a stable surface.

The tracker mount was chosen as the reference position from which rotational adjustments for leveling the HMD were measured. According to the nVisor ST60 schematics provided in the operator's manual, the tracker mount is parallel along all axes to the forward view through the HMD. Using a bubble level, the roll and pitch of the HMD were

26

adjusted such that the forward view would, according to the schematics, be parallel to the surface of the workbench.

Once the HMD has been positioned and secured, a small digital camera was rigidly mounted behind one of the HMD's oculi. The camera was then leveled using the same bubble level used for leveling the HMD. Using a Leica Total Station TPS800-Power (Figure 4.2), the exact height from the workbench surface, distance from the reference plane, and central offset from the HMD of the camera's lens was measured. Using this information, a real-world crosshair was projected on the reference plane using two laser levels. This crosshair defined the optical center of the camera's forward view. A virtual crosshair was position on the video feed from the camera such that it was centered on the pixel center of the feed. If the physical position of the camera was correctly measured, the real and virtual crosshairs would perfectly overlap. This enabled small misalignments to easily be seen and corrected by making very small rotational adjustments in the camera's position.

At this point, it was necessary to make sure that the camera is centered in the oculus' exit pupil. A method similar to that described in Rolland et al. [34] was used to ensure centering. A series of concentric circles were displayed in the HMD's graphics and the mechanical IPD control was adjusted until the view of the concentric circles appeared horizontally centered in the camera's video feed. The camera's height was then adjusted until the circles were vertically centered in the video feed. Due to the collimated nature of the vNisor ST60's optical system, no adjustments were needed at this step, but they were still checked as to not simply assume that no error existed.

27
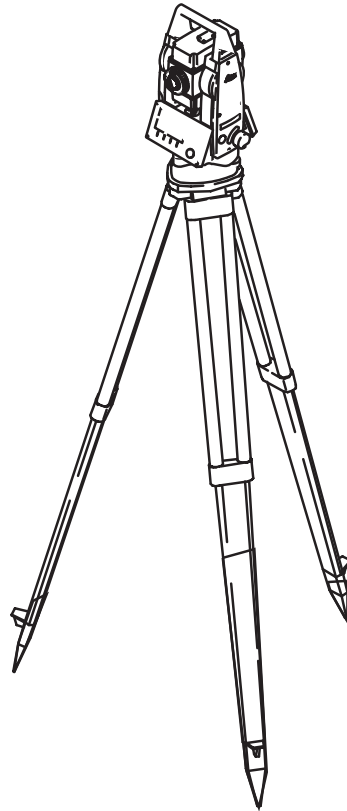
Figure 4.2

Leica Total Station TPS800-Power [18]

4.1.1.1 Field-of-View

Once both the camera and the HMD were properly positioned, the field-of-view for the oculus could be measured. A three pixel, red border is displayed in the HMD's graphics. Through the camera's video feed, this border appears as though it was projected on the real-world reference plane. Using a laser level, the vertical center of the camera's video feed was marked with a horizontal line which was aligned with the cross hair in the feed. Another laser level was then used to mark the video feed's horizontal center with a vertical line. These two reference lines establish a camera's center of view and should exactly overlap with the crosshair in the video feed. Using another laser level, a line was projected that exactly overlapped with the rightmost edge of the graphical border as seen through the video feed. The distance between the camera's center of view and the border was then measured using the TotalStation. This distance is referred to as the horizontal distance to the right edge of the display area, or simply as $hDist_r$. The laser reference line was then moved to exactly overlap with the leftmost edge of the graphical border as seen through the video feed. The distance between the camera's center of view and the left border was then measured using the TotalStation. This distance is referred to as the horizontal distance to the left edge of the display area, or simply as $hDist_l$. The distance from the camera to the reference plane was also measured and will be referred to as $camDist$. Using these three measurements the total horizontal field-of-view, $hFOV$, was calculated as the sum of two half-fields using the following equation:

$$hFOV = aTan(camDist/hDist_r) + aTan(camDist/hDist_l) \tag{4.1}$$

Similarly, a laser level was then used to project a line that exactly overlapped with the topmost edge of the graphical border as seen through the video feed. The distance between the camera's center of view and the border was then measured using the TotalStation. This distance is referred to as the horizontal distance to the top edge of the display area, or simply as $vDist_t$. The laser reference line is then moved to exactly overlap with the bottommost edge of the graphical border as seen through the video feed. The distance between the camera's center of view and the bottom border was then measured using the TotalStation. This distance is referred to as the horizontal distance to the bottom edge of the display area, or simply as $vDist_b$. Using these measurements the total vertical field-of-view, $vFOV$, was calculated as the sum of two half-fields using the following equation:

$$vFOV = aTan(camDist/vDist_t) + aTan(camDist/vDist_b) \qquad (4.2)$$

4.1.1.2 Principle Ray

Once the field of view had been measured, the principle ray, or the direction which the center of HMD's display is pointing, was measured. This was done by measuring the rotational offsets from a direct forward view through the display elements of the head-mounted display. This step is characterized by measuring the differences between a virtual crosshair displayed in the HMD's graphics and a reference crosshair displayed in the camera's video feed. Both cross hairs are placed such that they perfectly bisect the horizontal and vertical portions of their respective screen areas, intersecting in the center. The differences between these crosshairs are then measured and modeled in the virtual environment.

This ensures that a forward view from the observers eye will match the projected forward view of the graphics displayed in the HMD.
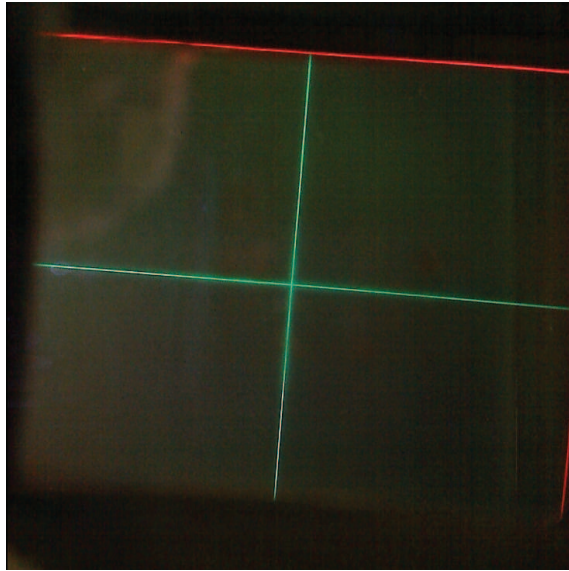
Figure 4.3

Calibration graphics as seen through the HMD

Firstly, a crosshair was displayed in the HMD graphics. Figure 4.3 depicts the graphics a seen through the HMD. Another crosshair was then overlaid on the video feed, marking the center of the forward view through the camera. The camera was then adjusted in yaw until the center of the video feed intersects the center of the HMD graphics, Figure 4.4. A series of laser lines were then projected onto the real-world reference plane such that they transcribed a triangle formed by the intersection of the HMD and video crosshairs and any differences in their roll. This arrangement is depicted in Figure 4.5. The Leica
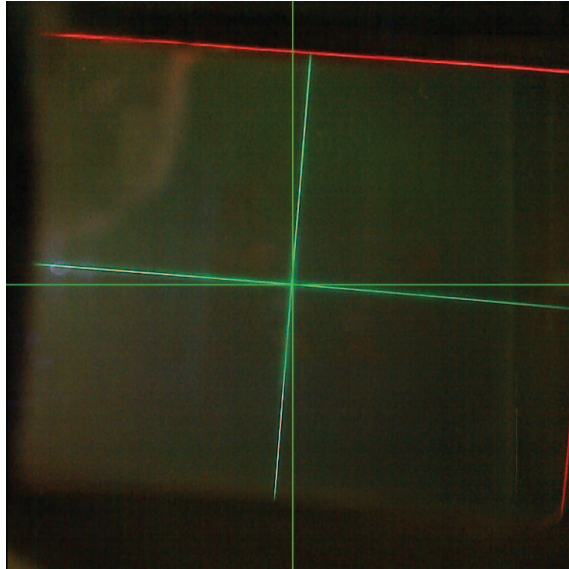
Figure 4.4

HMD graphics with video feed crosshair overlaid

TotalStation was then used to measure the sides of this triangle. These measurements were then used determine the difference in roll between the video feed and HMD graphics.

The HMD's roll was then adjusted such that the vertical lines of both the video and HMD crosshairs are aligned. Once this has been done, any difference in the pitch between the camera and HMD graphics should be visible, as seen in Figure 4.6. At this point, laser lines were projected onto the real-world reference plane and the pitch differences were measured with the TotalStation.

It is at this point that a vertical laser line presenting the corrected forward view of the camera is projected on the real-world reference plane, see Figure 4.7. The camera's yaw is then adjusted until the vertical portion of the video crosshair completely overlaps the projected laser line. Another laser line is then projected onto the reference place such
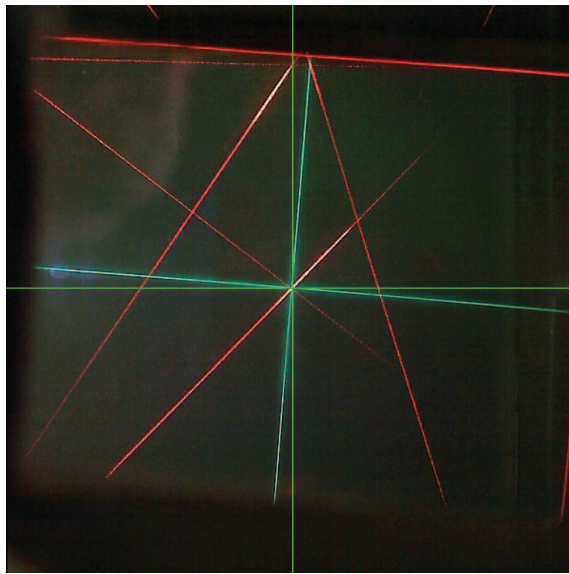
Figure 4.5

HMD graphics with video feed crosshair overlaid and real-world reference lines
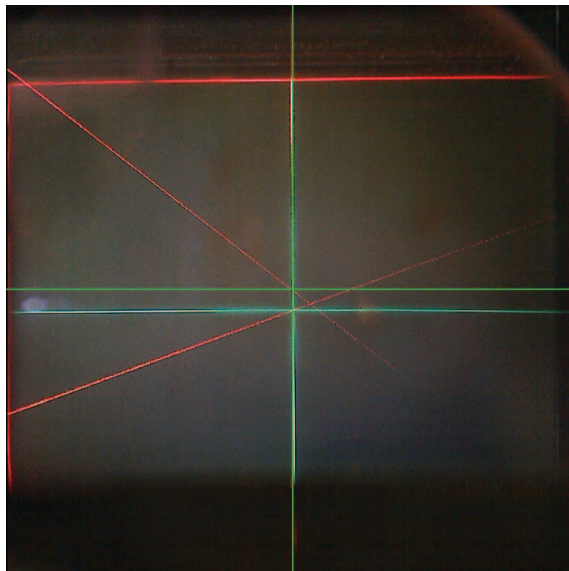
Figure 4.6

Measurement of pitch difference between the HMD and camera graphics

that it completely overlaps with the vertical portion of the graphics crosshair displayed in the HMD's graphics. A third laser line is then projected on the reference plane such that it completely overlaps with the horizontal portion of the crosshair in the video feed, see Figure 4.8. The TotalStation was then used to measure the offset between the intersection points of the two vertical laser lines and the horizontal laser line. These measurements provided the vergence of the HMD graphics relative to the forward view of the camera.
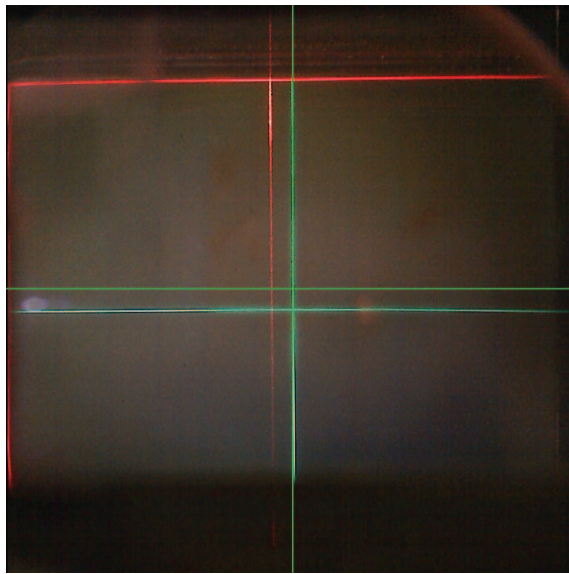
Figure 4.7

Crosshairs adjusted to measure the HMD's vergence

4.1.1.3 Optical Distrotion

These measurements enabled the principle ray of the given eye of the HMD to be modeled as three rotational offsets. The final step in the procedure was to measure the optical
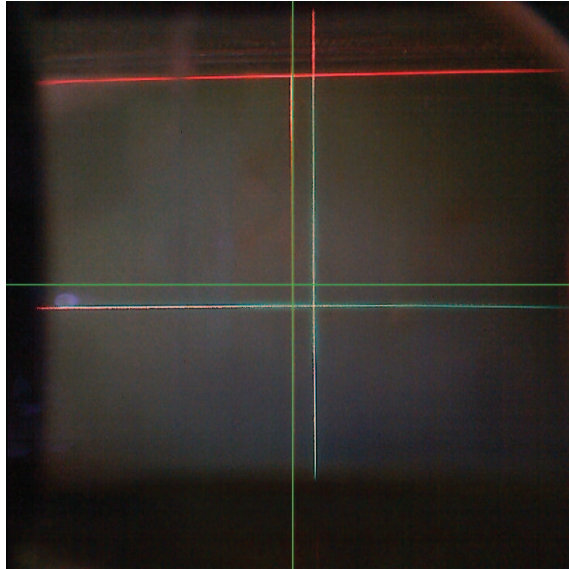
Figure 4.8

Reference lines and crosshairs used to measure the HMD's vergence

distortion introduced by the HMD's lens system. This was done using a method similar to that described in Owen et al. [26]. With the camera still centered with its view perpendicular to the real-world reference plane's surface, the HMD's roll, pitch, and yaw were adjusted until the HMD crosshair perfectly overlapped with the video feed's crosshair. This was done to ensure that the camera's view was both aligned with the reference plane and the HMD's graphical center. A uniform, rectilinear grid was displayed in the HMD's graphics. An image of this grid was then captured from the video feed, see Figure 4.9. The grid was then removed from the HMD's graphics and a real-world, laser projected grid of the same dimensions was projected on the reference plane, see Figure 4.10 for a realistic mock-up of the grid used in the calibration process. A pixel space analysis of the grid intersection locations was then conducted between the image of the virtual and real-world

grids. This analysis revealed the differences in the virtual and real-world grids to be, on average, less than 1 pixel and did not vary more than 2 pixels. This very closely matches the NVIS specifications of less than 1% pixel shift along a given axis due to optical distortion. This completed the Phase I calibration process for a single oculus. The camera then must be moved to the next oculus and this process repeated.
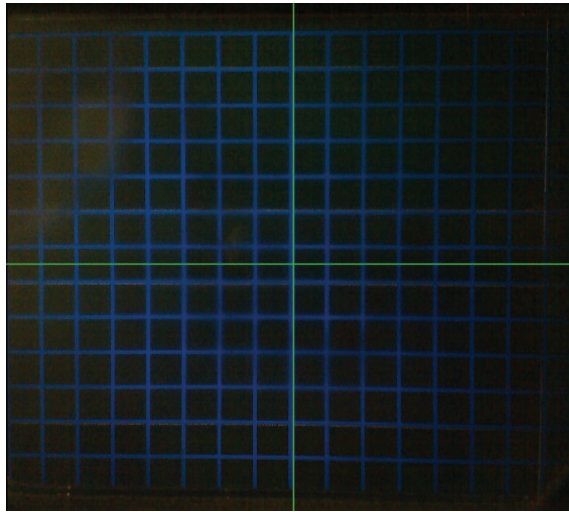
Figure 4.9

Actual virtual grid used to measure optical distortion

To achieve reliable results for all parameters discussed for this phase of the calibration, the parameters had to be measured several times and then averaged to get an set of values that approximated that of the HMD. However, even averaging over many measurements, the values often needed be slightly adjusted based on how the graphics appeared to the trained eye of an experimenter. This method was also somewhat laborious and required far more delicate handling of the hardware than is likely required by other methods.
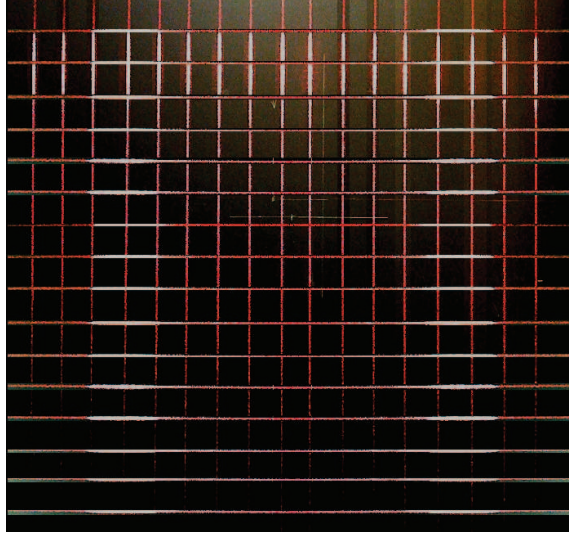
Figure 4.10

Mock of real-world grid used to measure optical distortion

4.1.2 Phase II: Observer Calibration

The calibration procedure used in this experiment consisted of three steps to correct for (1) optical alignment as well as (2) translational and (3) rotational errors reported by the head tracker.

The first step in the calibration procedure ensures that, for each eye, the observer's optical axis is aligned with the HMD's optical axis. To accomplish this, we implemented the calibration procedure presented by Rolland et al. [34], who also demonstrate that without this alignment an optical system presents optically incorrect depth cues. The observers were presented with a series of concentric circles that were centered about the optical axis of the display elements (see Figure 4.11a, top). The HMD has a knob on top of the head which raises and lowers the entire display frame relative to the observer's eyes. The observers were instructed to turn this knob until they could see an equal amount of the upper

37
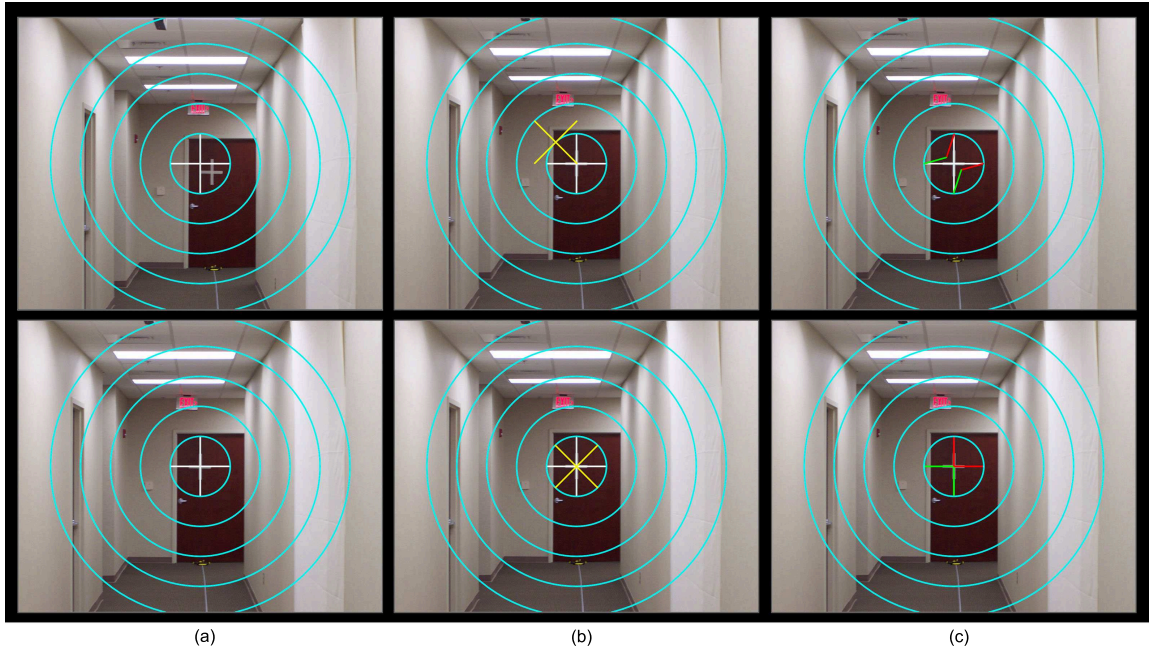
Figure 4.11

Participant's view of the calibration procedure

and lower portions of the outermost circle. The HMD also has knobs that independently shift the left and right display elements horizontally; observers were instructed to turn these knobs until an equal amount of the outermost circle could be seen on the left and right sides of each display. This procedure was performed monocularly for each eye. After these procedures, the optical axis of each of the observers' eyes was both horizontally and vertically aligned with the optical axis of each display element. In addition, each observer's interpupillary distance was measured using a pupilometer, an ophthalmic device specifically designed to perform this measurement. The graphics system used this distance when generating stereo imagery.

As part of developing the experimental apparatus, we carefully calibrated the 6 degree-of-freedom tracker for the hallway. However, because of differences in the way the HMD sits on the head, there are always noticeable translational and rotational errors, even if the display is removed and then replaced on the same observer's head. The goal of the second calibration step was to correct for tracker errors along the observers' x (horizontal) and y (vertical) axis. While similar errors also existed along the z (depth) axis, it was not necessary to correct for them, because the experimental task was always conducted at the same z location for each observer. For this calibration step, the observers were shown a virtual crosshair and a real-world cross placed at their eye height at the end of the hallway (Figure 4.11a, top). The observers were then asked to align the two crosshairs by moving their heads (Figure 4.11a, bottom). Once the observers had aligned the crosshairs, their line of sight was parallel to the floor. They were next handed a game controller and shown a virtual, yellow "X" that was translationally controlled by the head tracker (Figure 4.11c, top), which shows a typical degree of translational error). The initial position of the X represented the location where the real-world crosshair should be located according to the tracker. The observers then used the game controller to adjust the position of the X until it was aligned with both the real and virtual crosshairs (Figure 4.11d, bottom). This adjustment added a translational offset to the values reported by the head tracker, which translationally corrected for the way the HMD was sitting on their head.

The goal of the third calibration step was to correct for rotational tracker errors around the observers' pitch (up/down) and yaw (side/side) axis. The tracker also had roll (twist) errors, but these errors were not important for this task. The observers were shown the

39
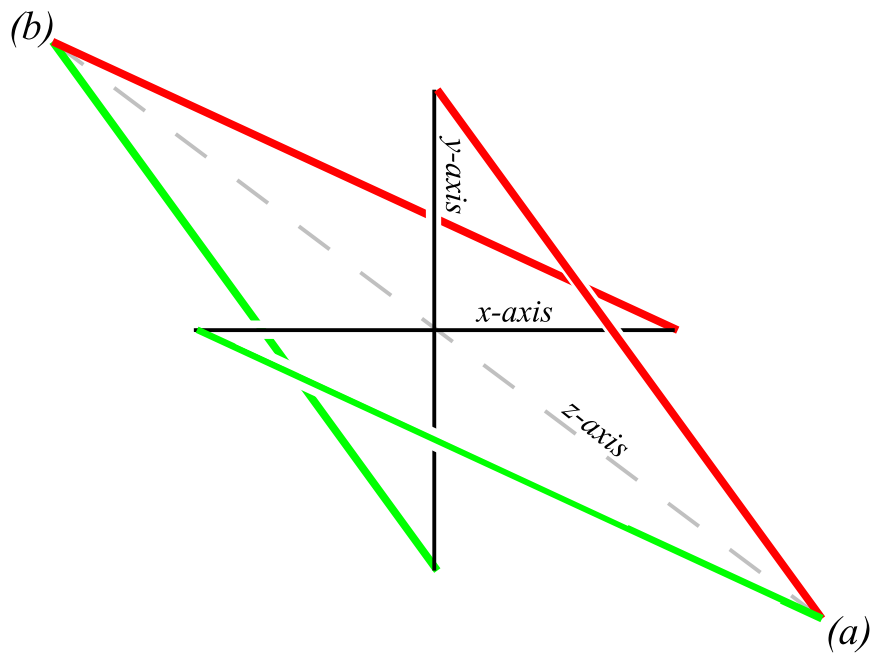
Figure 4.12

Side view of the 3D Compass

same real and virtual crosshairs as in the previous step and asked to perform the same boresighting task (Figure 4.11a). This time the observers were shown a 3-dimensional crosshair that we called the 3D Compass (see Figure 4.12). The 3D Compass is rotationally controlled by the head tracker, but it is translationally centered at the virtual crosshair. The shape of the 3D Compass is such that if there is any rotational offset when aligned with the real world crosshair, its 2D projection results in an accidental view with a star-like shape (Figure 4.11c, top, which shows a typical degree of rotational error). However, when all rotational errors have been compensated, the 2D projection results in another accidental view that looks like a plus sign (Figure 4.11c, bottom). The observers were given a game controller and asked to adjust the shape until it became a plus. This adjustment added a pitch and yaw offset to the values reported by the head tracker. Together, these calibration procedures resulted in accurate registration between the virtual and real worlds. Observers were required to perform this calibration before every block of trials in the AR and VR viewing conditions. Also, if the observers touched, moved, or otherwise jostled the HMD at any point during the trials, the calibration procedure was repeated before any further data was collected.

4.2   Motion Tracking Verification

The positional data provided by the IS-1200 was verified for accuracy prior to being used in any experiments. This was an important step as the position reported by the tracker is used to determine the participants' position and viewing direction in the virtual and
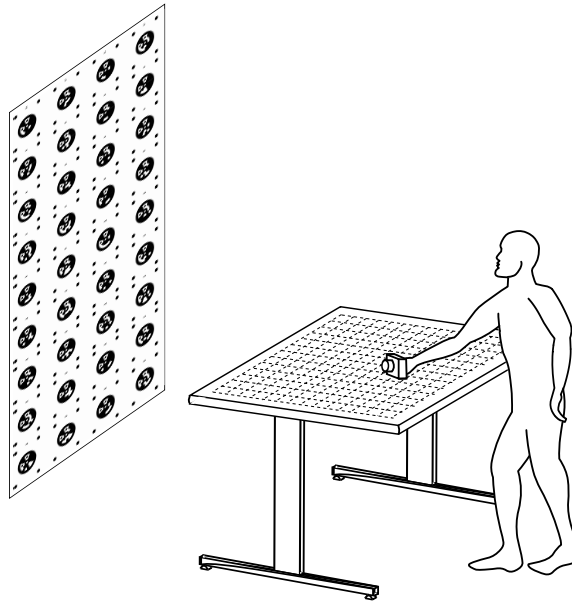
Figure 4.13

Real-world verification of tracking data

augmented environments. Errors in this data would adversely affect the viewing of nearly all computer generated imagery in these experiments.

First, the fiducial constellation to be used for the experiments was hung at a well known and measured position. A small table was then placed in front of the constellation, centered at roughly where participants would be standing during an experiment. This was done to provide a reference surface that would encompass the positions at which a participant would reasonably view the virtual or augmented environment. The table was positioned such that its backmost edge would be parallel to the surface of the hanging constellation. A uniform, rectilinear grid was then drawn on the table's surface. Given that the table and grid were in well established positions, relative to the constellation, the IS-1200 was then

systematically placed at each grid position. Its reported position was then compared to its real-world position. Figure 4.13 demonstrates this procedure.

The position reported by the IS-1200 was typically within 1cm of its actual position. However, the reported position tended to drift in a circular pattern around its actual position over time. The drift typically did not exceed ±2cm. These measurements were taken while the tracker was sitting on the table's surface. It was soon noticed that when the tracker was moved or bumped the drift in the reported position would drop to roughly ±1cm. A small amount of vibration applied to the table's surface provided the same decreased drift. This seems to be an effect of the hybrid optical/inertial nature of the IS-1200. Its positional data was most stable when both the optical sensor and the inertial sensor were receiving active stimulation. The near complete stillness of sitting the IS-1200 on the table's surface seems to have been introducing drift in the inertial sensor. Since it is very difficult and somewhat unnatural for people to hold their heads completely still, participants in the experiment would be continually introducing small movements that would be detectable by the inertial sensor. This would then aid in decreasing drift introduced by the inertia sensor.

4.3   Stereo Vision Test

Prior to participating in any of the experiments detailed in this document, all volunteers were required to pass a stereo vision test. The stereo vision test consisted of presenting a participant with a series of nine, numbered red-blue anaglyph stereo objects, see Figure 4.14. The objects were outlined diamond shapes containing four circles. There
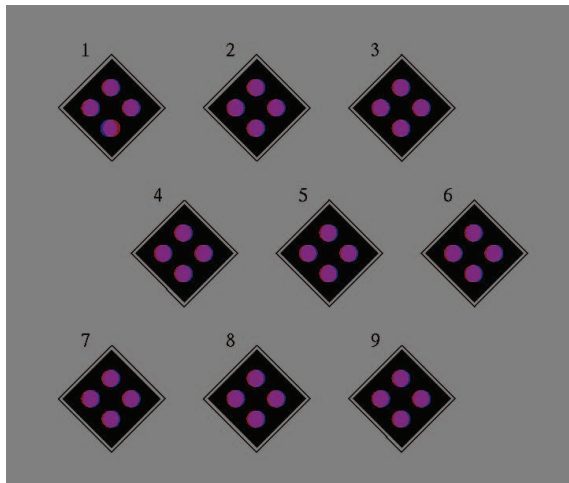
Figure 4.14

Stereo vision test
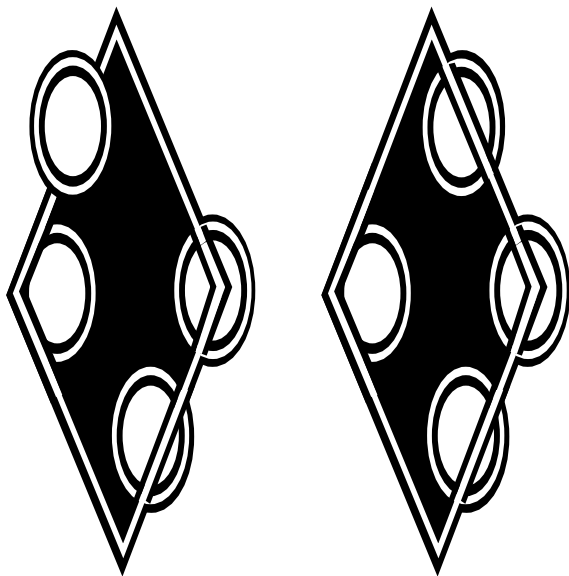
Target Object   Distracter Object

Figure 4.15

Target and distracter objects

were two classes of objects: targets and distracters. In the distracter object, when viewed stereoscopically, the four circles appear to be sunk into the diamond such that they appear to be behind the screen's surface. However, in the target object, one of the four circles would appear to be sticking out of the diamond such that it would appear to be hovering in front of the screen. These objects are depicted at an oblique angle to exaggerate their features in Figure 4.15.
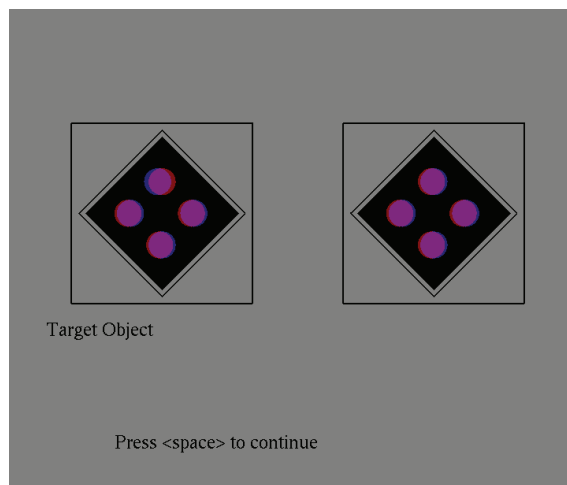
Figure 4.16

Stereo object training examples

As seen in Figure 4.18, the test was administered on a standard LCD screen at a distance of 67cm from the participant. Participants' heads were not restricted, so this distance varied somewhat on a person-to-person basis. Prior to viewing the test, participants were asked to put on a pair of red-blue anaglyph stereo glasses whose filters were specifically designed to match the hues generated by most LCD screens. The stereo vision test begins
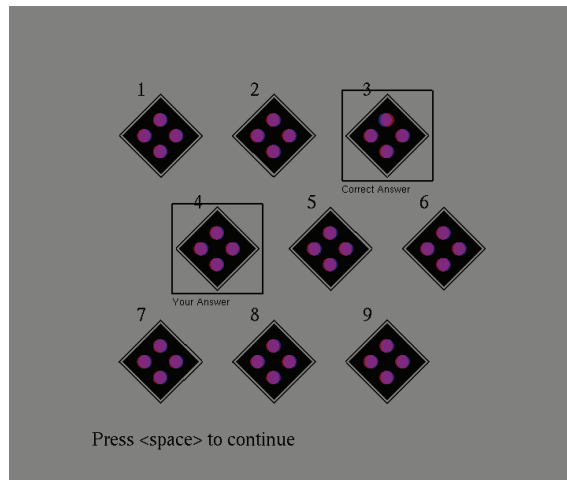
Figure 4.17

Stereo object training trial

by showing the participants exaggerated target and distracter objects side-by-side, Figure 4.16. The experimenters explained that participants are to search for a target object hidden among a field of distracter objects. At this point, the participants are presented with five practice trials where they received feedback on whether or not they have successfully identified a target object for a given trial. Participants indicated their response by pressing the number on the keyboard that matched their numbered selection on the screen. Figure 4.17 shows the feedback that participants received for incorrectly identifying a target object during the practice trials. When participants correctly identified a target object during the practice trials, they were simply presented with a blank screen with the word "Correct" before progressing to the next trial.

After completing the five practice trials, participants began the measurement portion of the stereo vision test. This part of the test worked in the same manner as the practice

Figure 4.18

Stereo vision test apparatus

trials except there was no feedback or pause between trials and it consisted of 10 trials. At the end of this portion of the test, the percent of correctly identified targets was presented on the screen. A passing score was classified as correctly identifying at least 20% of the target objects. Participants that scored below this 20% were not permitted to continue the experiment.

4.4   Judgment Techniques

Much work has been done to understand what factors effect depth judgments in VR. However, a much smaller body of work exists that analyzes depth and layout in AR. Current evidence indicates that similar distance underestimation may occur in AR but to a lesser degree [14, 36]. The vast majority of this work focuses specifically on medium-field distances (approximately 1.5m to 30m) [3]. Depth judgments in this range are typically taken with one of three implicit measurements: blind walking, triangulated walking, or imagined walking. These techniques all rely on the observer's ability to navigate based on an internalized, cognitive representation of their surroundings.

Blind walking is the most well established of these techniques. Blind walking requires an observer to view an object and then attempt to walk to its position without the aid of vision, see Figure 4.19. Loomis and Knapp [20] compiled the results of numerous studies that indicate that observers are highly accurate at performing this task in the real-world.

However, for use in virtual environments, this technique requires that a direct path exists to the viewed position in both the real and virtual worlds. For some virtual environments, such as CAVEs or display walls, this simply may not be possible. In this case,
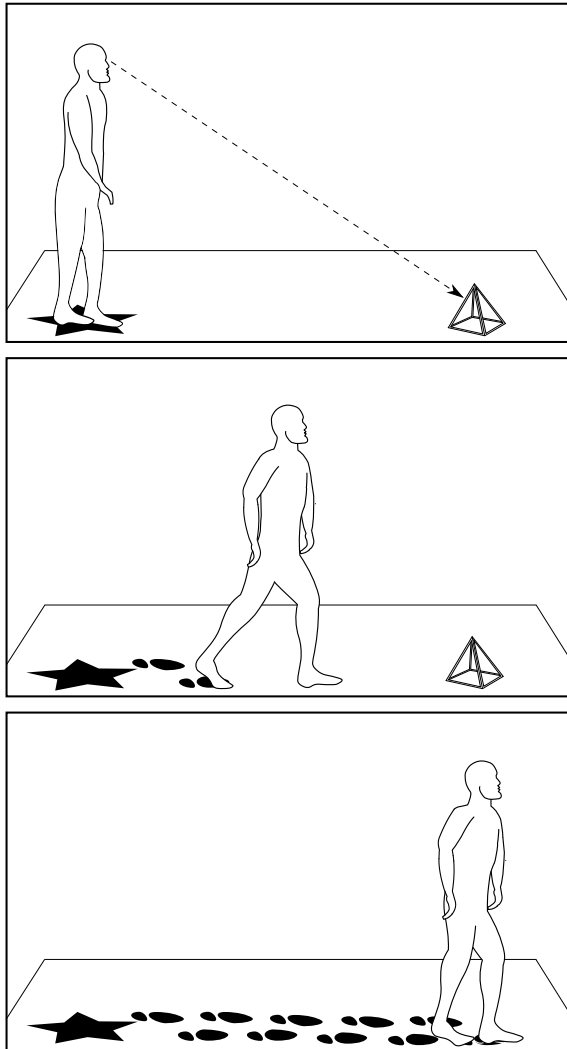
Figure 4.19

Visually directed blind walking

techniques such as triangulated walking may be more appropriate. Triangulated walking is similar to blind walking except the observers view an object, close their eyes, turn an oblique angle, walk some short distance, stop, and point at the remembered location of the object, see Figure 4.20. The starting position, stopping position, and pointing direction are then marked and measured. Using these values, the judged position of the object can be calculated. This technique does, however, have the drawback of producing results that have a somewhat higher deviation from measurement to measurement.

In some cases there is even insufficient room to perform triangulated walking. In these cases, imagined walking can be used. In imagined walking, the observers view an object, close their eyes, start a timer, and imagine walking to the object's position, see Figure 4.21. Upon arriving at the imagined position, the observers stop the timer. The elapsed time of the imagined walk is then recorded. Afterward, the observers are taken to another location where locomotion is not restricted and asked to walk a specific distance. Their total travel time during this walk is recorded and used to estimate their walking speed, thereby producing a value by which the previously recorded imaged walking times can be multiplied. The resulting value provides an estimated distance for each judgment of the object's position. As with triangulated walking, this technique can produce results comparable to those of blind walking [15].

For the experiments described in this document, blind walking was exclusively used as the means of measuring egocentric distance judgments. However, since the participants were tethered by the data and video cables associated with the HMD and motion tracker, a few modifications had to be made to the general blind walking procedure. These modi-
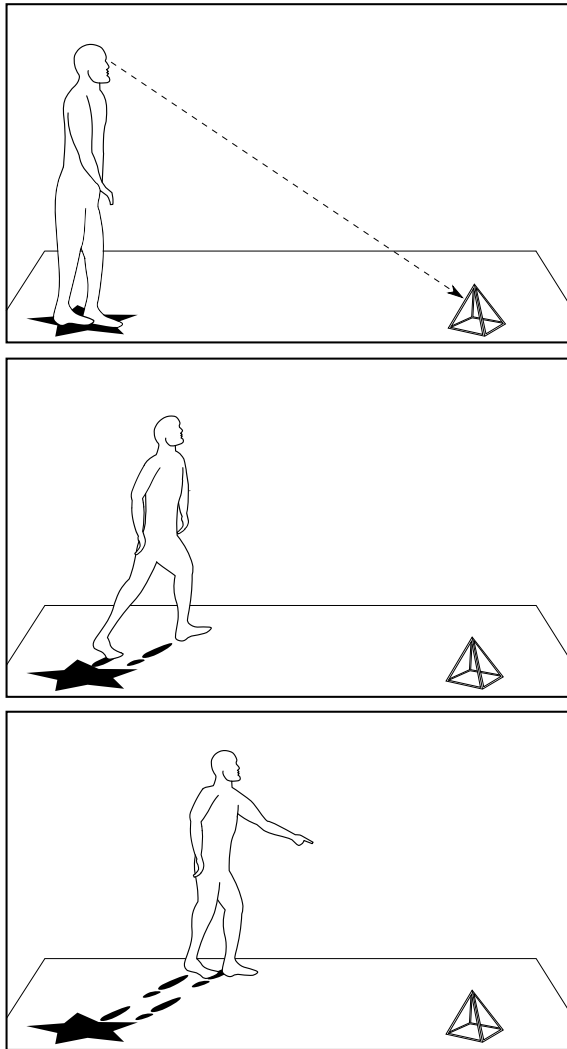
Figure 4.20

Triangulated walking
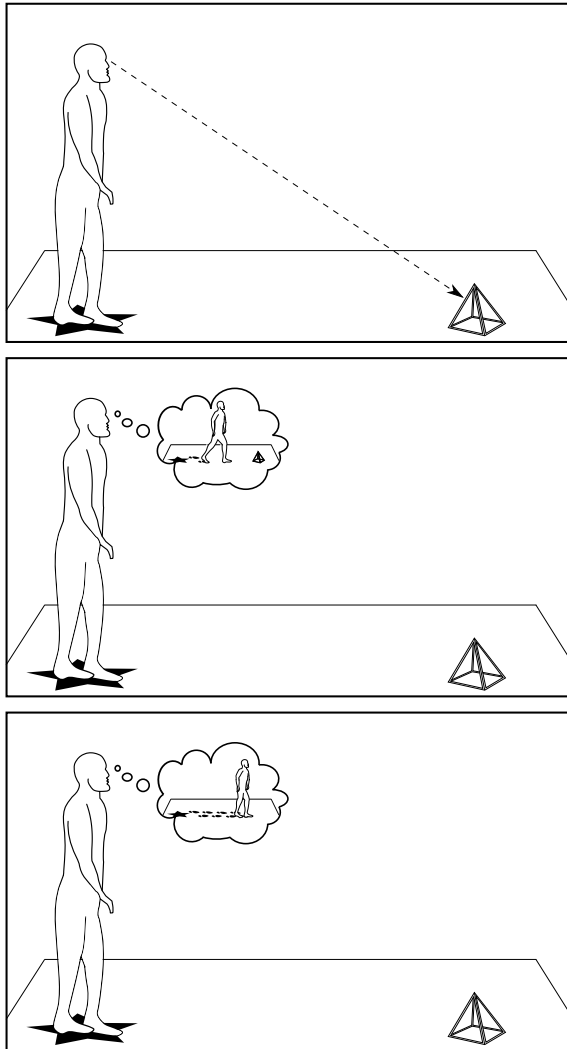
Figure 4.21

Imagined walking
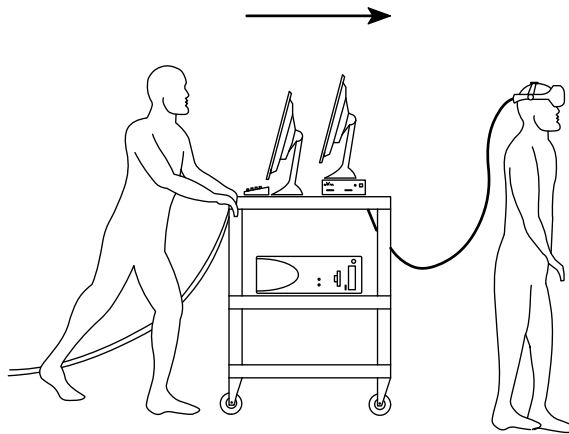
Figure 4.22

Procedure for a typical judgment walk
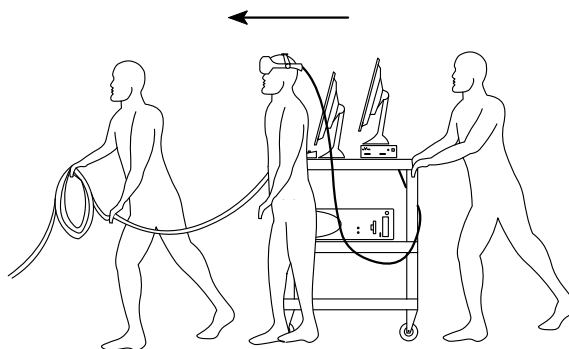
Figure 4.23

Procedure for a typical return walk

fications are depicted in Figure 4.22 and Figure 4.23. The length of the cables associated with the equipment worn by the participant was no longer than 4 meters, but the walkable distances needed for this experiment were nearly triple that length. To make walking these distances possible while wearing the HMD and tracker, the host computer and video control unit were placed on a rolling cart that was connected to a power outlet via a very long extension cable. When participants performed the blind walking task, an experimenter would push the cart behind the participants allowing the extension cable to unroll behind them. On the return walk, one experimenter would walk in front of the participants while rolling up the extension cable, and a second experimenter would push the cart back to its starting position. Also, to keep the participants isolated from audio based cues in the surrounding environment, they were required to wear headphones that played a constant stream of white noise. The headphones were also patched into a wireless microphone system. This microphone system was used to communicate instructions to the participants during the course of the experiment. For hygienic purposes, the earphones were cleaned, disinfected, and placed in a disposable covering prior to running each participant. The actual instructions and scripts used by the experimenters for this procedure can be found in the appendices of this document.

CHAPTER 5

EXPERIMENTATION

5.1 Experiment I

One of the main criticisms of the experiment described in Jones et al. was that the within-subjects, repeated-measures experimental design could potentially lead to transfer effects across conditions, introducing the possibility that exposure to one condition could affect performance in another [14]. This concern was the motivation behind Experiment I, which was a between-subjects replication of the experiment described in Jones et al. [14]. Experiment I's aim was to determine whether or not the unusual lack of underestimation in Jones et al. [14] was due to transfer effects introduced by the within-subjects experimental design [14]. Additionally, the experimenters were curious to determine if there may be a possible link between the participants' dark vergence and the degree by which underestimations occur. This was inspired by other work that investigated perceived visual changes as influenced by dark accommodation and vergence of observers [22–24, 27]. Since the equipment necessary to measure dark accommodation was not attainable for this experiment, the experimenters decided to rely on the well establish accommodative convergence connection. The accommodative and convergent components of the human visual system typically respond in direct relation to the other. This relationship was exploited for this

experiment and dark vergence and dark accommodation were assumed to change proportionally to each other.

5.1.1 Method

A group of 39 naive participants was recruited from the general university population and were monetarily compensated for their participation. Figure 5.1 shows the experimental environment, which was a hallway at the Mississippi State University Institute for Imaging and Analytical Technologies, measuring 1.82m in width and 23.45m in length. Participants were screened for visual dysfunction by self-report and tested for normal stereo vision prior to being allowed to participate in the experiment. Additionally, participants' eye-heights and interpupillary distances were measured prior to beginning the experiment. These measurements were used for individual calibration of the virtual and augmented environments. To present the virtual and augmented environments, a NVIS nVisor ST optical see-through head-mounted display (HMD) equipped with an Intersense IS-1200 motion tracking system was used for the presentation of all computer generated imagery. These devices are discussed in more detail in Chapter 3. Opaque, foam rubber occluders were attached to the left and right sides of the HMD in order to prevent participants from seeing the surrounding environment.

Figure 5.3 depicts the HMD and occluder configuration used in both Experiment I as well as in Jones et al. [14]. Participants performed visually directed blind walking as a method of measuring their egocentric distance judgments [16, 20, 32], as illustrated in Figure 4.19. Participants were instructed to blindly walk until they felt as though the tips

Figure 5.1

Real-world experimental environment

Figure 5.2

Virtual experimental environment
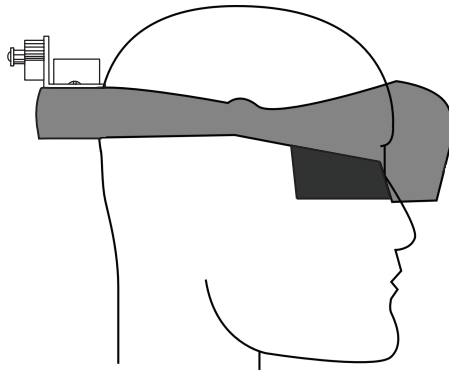
Figure 5.3

HMD in the standard configuration
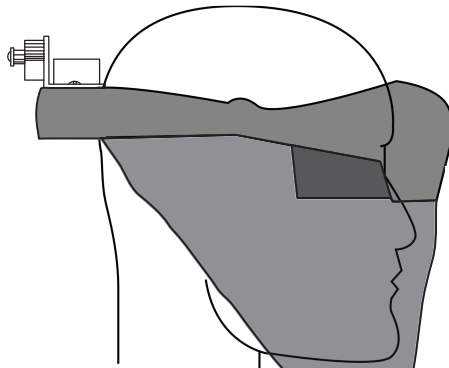
Figure 5.4

HMD in the fully occluded configuration

of their toes were at the target distance. The stimulus used to indicate the target distance was a white, wireframe pyramid measuring 23.5cm in height with a 23.5cm square base.

Prior to beginning the experiment, the dark vergence of the participants was measured using the apparatus described in Section 3.3. After taking this measurement, participants were briefed on the blind walking procedure and were given 5 practice trials of blind walking in an adjacent hallway of similar proportions to the experimental environment. This was done to build the participants' confidence in walking without vision. At this point, participants were escorted to the experimental environment. To prevent miscellaneous auditory cues from influencing the participants' behavior, they were equipped with earphones that played continuous white noise. The volume of the white noise was adjusted until the participants judged it to be subjectively comfortable. Additionally, the earphones were patched into a wireless microphone system through which the experimenters communicated instructions to the participants. The wireless microphone receiver and white noise generating device were stored in a backpack that the participants wore during all experimental conditions. Distance judgments from the blind walking task were measured with a white surveyor's tape that spanned the length of the hallway.

5.1.2 Design & Procedures

This experiment was intended to be a between-subjects replication of the experiment described in Chapter 4. For this reason, four experimental conditions were tested: Real World (Real), Real World seen through the HMD (ReHMD), Augmented Reality (AR), and Virtual Reality (VR). Jones et al. [14] also tested two, crossed viewing conditions:
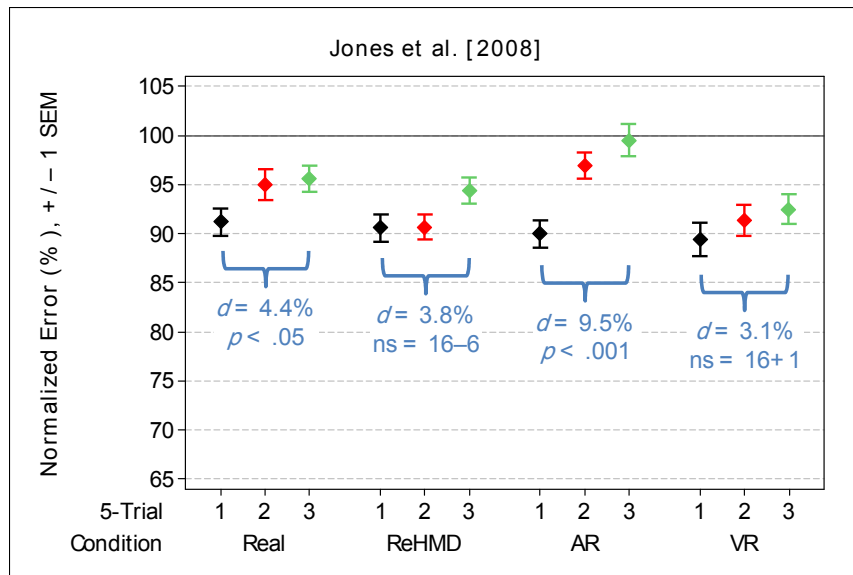
Figure 5.5

Distance judgments from Jones et al. 2008
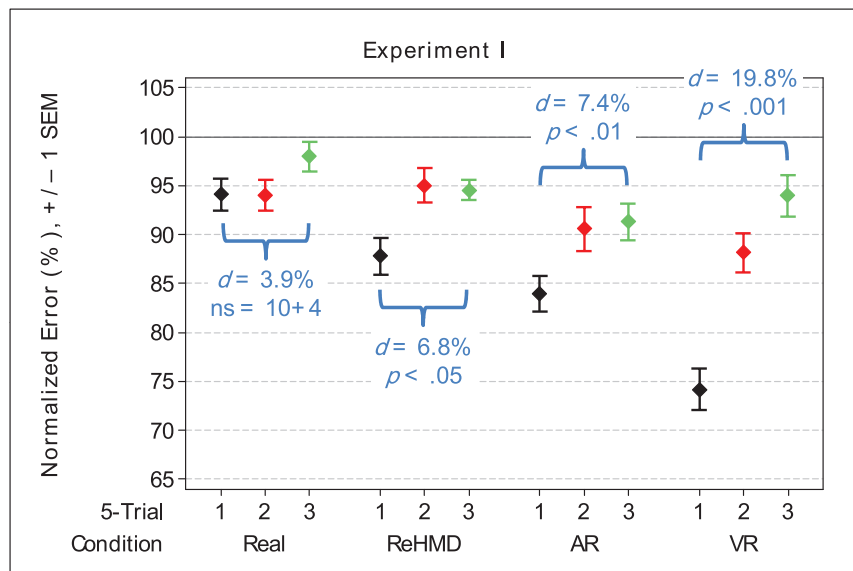
Figure 5.6

Distance judgments from Experiment I

still and motion. Respectively, participants either viewed stimuli while standing stationary (still) or while swaying from side-to-side to induce motion parallax (motion). However, since no consistent effect of the motion condition was observed, it was excluded from this experiment. Participants' movements were not restricted, but they were instructed to look directly at the stimulus during the experiment. Exact computer models of the experimental environment and stimulus were used in the VR condition, depicted in Figure 5.2. An exact computer model of the stimulus was used in the AR condition. Stimuli were presented at one of five distances ranging from 3 to 7 meters in 1 meter increments. Each distance was repeated three times, providing 15 total trials per experimental session. The presentation order of the stimulus distances was determined using a restricted random shuffle, with the restriction that no target distance was repeated in consecutive trials.

Participants were instructed to close their eyes between each trial, at which point the stimulus was placed. Participants were then instructed to open their eyes and observe the stimulus until they felt confident enough to blindly walk to its position. Upon indicating their readiness, the participants were instructed to close their eyes and walk to the object. Once the participant reached their judgment distance, they stopped walking and kept their eyes closed until instructed to turn back in the direction of their starting position. Participants were then allowed to walk back to the starting position with their eyes open. In the Real, ReHMD, and AR conditions the experimental environment was fully visible during the return walk. However, the virtual environment was not displayed and the optical see-through window was closed during the return walk in the VR condition.

The same calibration and alignment procedures discussed in Jones et al. [14] were used prior to beginning each experimental session. These procedures helped ensure that the participants' real-world eye and head positions and orientations matched those modeled in the virtual and augmented environments. Before and after each experimental session, participants were screened for signs of simulator sickness and impaired locomotion.

5.1.3 Analysis

All analyses were conducted with *Normalized Error = Judged Distance/Actual Distance*. Each experimental Condition is subdivided according to *5-Trial*, the mean of 5 consecutive trials, so *5-Trial*$_1$ = mean($trial_1$: $trial_5$), *5-Trial*$_2$ = mean($trial_6$: $trial_{10}$), and *5-Trial*$_3$ = mean($trial_{11}$: $trial_{15}$); in other words *5-Trial* breaks the normalized error into the first, second, and final thirds of the experimental sessions. In addition, this paper reports non-significant hypothesis tests in the form "*ns = N + A*", where "*ns*" denotes a non-significant result, *N* is the number of participants that were run, and *A* is the number of additional participants that an a priori power analysis indicates would need to be run in order to achieve power = .80, assuming the effect size *f* and the correlation among repeated measurements *r* remain constant as additional participants are run, and assuming $\alpha = .05$. Thus the magnitude of *A* relative to *N* quantifies the truth of the null hypothesis. Some results are reported "*ns = N − A*"; these indicate that with *N* participants *power* > .80, and *A* is the number of participants that would need to be removed for *power* = .80, given the same assumptions for *f*, *r*, and α. Power calculations used G*Power software and the techniques discussed by Faul et al. [5].
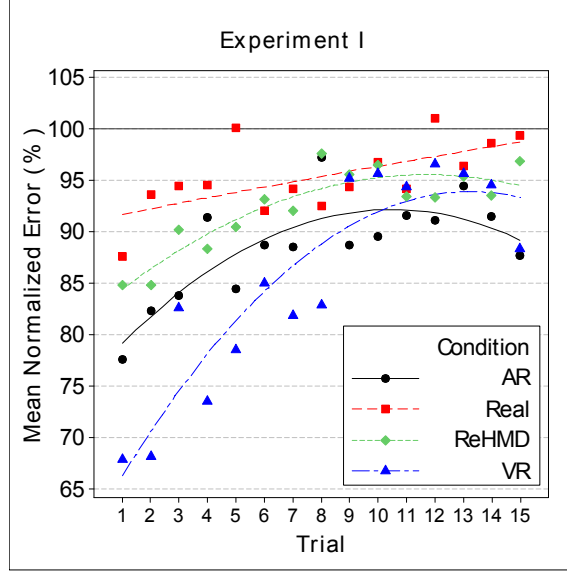
Figure 5.7

Distance judgments by trial fit with quadratic regressions

5.1.4 Results

Figure 5.6 shows the results of Experiment I; here the Real condition served as the control for comparison with the ReHMD, AR, and VR conditions. Distance judgments in neither the ReHMD (92.4%) nor AR (88.6%) conditions significantly differed from those in the Real (95.3%) condition (ReHMD: $F(1,18) = 0.928$, $p = 0.348$, $ns = 20+152$; AR: $F(1,18) = 3.084$, $p = 0.096$, $ns = 20+34$). The VR (85.4%) condition exhibited significant underestimation of distance as compared to the Real condition ($F(1,17) = 7.324$, $p = 0.015$).

As previously discussed, Experiment I was intended to be a between-subjects replication of Jones et al. [14], in order to determine if the unusual results seen in Jones et al. [14] were an effect of that experiment's within-subjects design. Figure 5.5 shows the

mean distance judgments found in Jones et al. [14] [1] : Real (93.9%), ReHMD (91.8%), AR (95.5%), and VR (91.0%). These are very similar to those found in Experiment I, differing by 1.4%, 0.6%, 6.9%, and 5.7% respectively. The low amount of underestimation is especially noteworthy in the VR condition, where underestimation has typically been reported ranging from 50% to 80% of veridical [16, 31, 38, 40, 42, 43]. Disregarding previous exposures and treating each condition from Jones et al. [14] as a unique exposure, an analysis of variance was conducted comparing distance judgments between the two experiments. This analysis reveals that there was no significant difference between the conditions described in Jones et al. [14] and their counterparts in Experiment I (Real: $F(1,24) = 0.170$, $p = 0.684$, $ns = 26+1184$; ReHMD: $F(1,24) = 0.040$, $p = 0.843$, $ns = 26+4220$; AR: $F(1,24) = 2.930$, $p = 0.100$, $ns = 26+46$; VR: $F(1,23) = 1.959$, $p = 0.175$, $ns = 25+79$).

These results seem to counterindicate experimental design as the main factor behind the unusual lack of underestimation seen in both Jones et al. [14] and Experiment I. However, they prompted a thorough reexamination of Experiment I, which revealed a strong trend of improved distance judgments throughout the course of the experiment. Figure 5.7 shows a plot of normalized error means for the conditions by trial and fit with quadratic regressions (Real: $R^2 = 41.0\%$; ReHMD: $R^2 = 81.4\%$; AR: $R^2 = 67.1\%$; VR: $R^2 = 83.3\%$). As Figure 5.6 shows, the effect of improved distance judgments over time becomes even more obvious when examining the data subdivided by *5-Trial*. An analysis of variance was

---

[1]In Jones et al. [14] 16 trials were collected per condition, but in order to allow the two experiments to be directly compared, for this analysis the final trial is dropped.
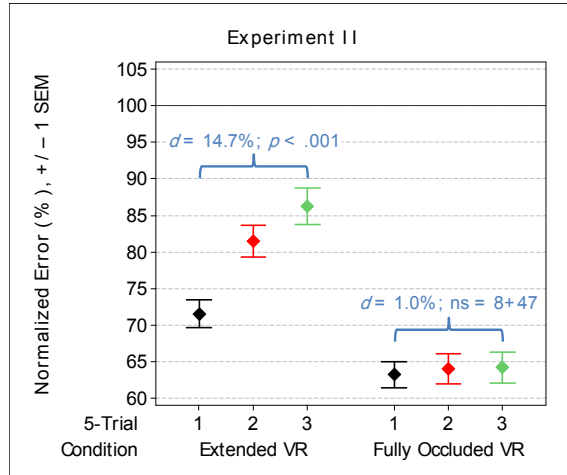
Figure 5.8

Distance judgments in the Extended and Fully Occluded conditions

conducted to examine the effect of time in terms of *5-Trial* on distance judgments. Additionally, an effect size $d = \textit{5-Trial}_3 - \textit{5-Trial}_1$ was calculated between the last and first *5-Trial* to illustrate the size and direction of the adaptation over time. This revealed that all conditions, excepting Real, exhibited significantly improved normalized error between the first and third *5-Trial* (Real: $F(2,18) = 1.029$, $p = 0.378$, $ns = 10+4$, $d = 3.9\%$; ReHMD: $F(2,18) = 3.732$, $p = 0.044$, $d = 6.8\%$; AR: $F(2,18) = 7.176$, $p = 0.005$, $d = 7.4\%$; VR: $F(2,16) = 27.071$, $p = 0.000$, $d = 19.8\%$). As illustrated in Figure 5.6, toward the end of the experimental session, for each condition participants are judging distance within 90%, on average, of the actual target distance. This finding prompted another look at the data from Jones et al. [14] to see if a similar trend existed there as well. An analysis of variance was conducted on the data from Jones et al. [14] to examine the effect of time in terms of *5-Trial* on distance judgments. The effect size, as previously described, was also

calculated to illustrate the size and direction of the adaptation across over time. Figure 5.5 shows the results of this analysis, which are very similar to those found in Experiment I. The Real and AR conditions exhibited significantly improved normalized error over time while the ReHMD and VR conditions did not (Real: $F(2,30) = 3.538$, $p = 0.042$, $d = 4.4\%$; ReHMD: $F(2,30) = 2.376$, $p = 0.110$, $ns = 16?6$, $d = 3.8\%$; AR: $F(2,30) = 17.874$, $p = 0.000$, $d = 9.5\%$; VR: $F(2,30) = 0.995$, $p = 0.382$, $ns = 16+1$, $d = 3.1\%$). Though the ReHMD and VR conditions did not exhibit statistically significant effects, they could, in fact, be masked by the within-subjects design after all. This seems plausible as the effects observed in Experiment I are subtle and time dependant.

Vergence Distance by Repetition
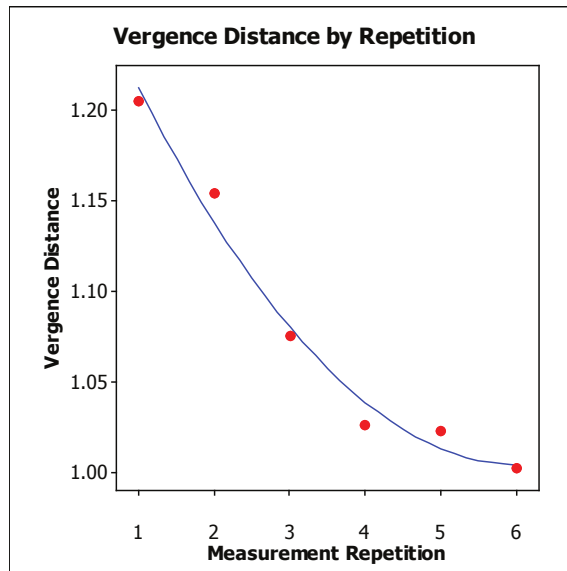
Figure 5.9

The effect of measurement repetition on mean vergence distance

Six total measurements of dark vergence were taken, and upon examining each repetition separately, an interesting pattern emerged. Figure 5.9 clearly shows that the participants' eyes were gradually converging toward a distance of 1 meter as they continued to perform the dark vergence measurement task. Since the physical distance between the participants and the stimulus was 1 meter, it seems that some visual information indiciative of the stimulus' position was gradually being accumulated over the course of the measurements. However, it is important to note that the calibration of the nonius device was not sufficiently validated prior to data collection, therefore the validity of these measurements may be questionable. The first two repetitions of the dark vergence measurements did not significantly differ ($F(1, 82) = 0.39$; $p = 0.536$). Assuming that the participants' eyes were at their dark resting positions at the beginning of the measurement task, an average of the first two repetitions may be more representative of their actual dark vergence. This finding, combined with the adaptation effects observed during the blind walking task, prompted a reanalysis based on the first two dark vergence measurements and the first $5 - Trials$ of blind walking. However, even with an analysis of variance of this subset of the data detected no statistically significant interaction between dark vergence and blind walking results (Real: $F(1, 48) = 0.714$, $p = 0.714$; ReHMD: $F(1, 48) = 3.36$, $p = 0.073$; AR: $F(1, 48) = 6.93$, $p = 0.011$; VR: $F(1, 43) = 0.19$, $p = 0.664$). If an effect of dark vergence or dark accommodation on the perceived distance of projected graphics in either augmented or virtual environments does exist, it seem plausible that these effects would express themselves more powerfully at near-field distances since both accommodation and convergence provide substantially more distance information at these distances.

5.2 Experiment II

Though previous work has demonstrated that participants can significantly improve their performance in the absence of explicit feedback [7, 28], the strong trend of improved distance judgments seen in Experiment I raised the possibility that participants may have been receiving feedback regarding their performance from some uncontrolled aspect of the experiment. This prompted a thorough reexamination of the experimental procedures used in both Jones et al. [14] and Experiment I. After carefully scrutinizing the experimental procedures, we could find no sources of explicit feedback that could give participants knowledge of their performance. However, two possible sources of indirect feedback were identified: (1) proprioceptive feedback from the blind walking task itself and (2) peripheral visual information available via a gap below and between the HMD and the participants' face. The vertical field-of-view of the gap varied depending on the declination of each participant's head but ranged from roughly $35°$ to no more than $50°$ Experiment II attempts to identify which of these potential sources of feedback could be influencing participants' perception of the virtual environment. This experiment compared two conditions: extended walking (Extended) and fully occluded periphery (Fully Occluded). The Extended condition was intended to remove any proprioceptive feedback by forcing observers to perform their return walk from a randomly selected distance further than their judgment distance. The Fully Occluded condition involved wrapping an opaque, black cloth around the bottom and sides of the HMD in order to prevent exposure to any peripheral visual information, as depicted in Figure 5.4. These conditions were tested only in virtual reality, as VR exhibited the strongest adaptation effect in Experiment I.

5.2.1 Method

For this experiment, 16 naive participants were recruited from the general university population and either received course credit or monetary compensation for their participation. Eight participants experienced each condition in a between-subjects design. The procedures for this experiment closely followed the procedures used in Experiment I: the same screening and training protocols were used, but the experimental protocol differed slightly as required by the new experimental conditions. For the Extended condition, participants performed the same blind walking task as in Experiment I, except that the return walk differed. Once the participants completed their judgment walk and their walked distance was measured, they were asked to blindly walk forward until instructed to stop. The extended distance varied randomly from 1 to 4 meters. The participants then performed a normal return walk from the new position. This condition was intended to ambiguate any proprioceptive feedback from walking the judged distance twice: once on the judgment walk and again on the return walk. For the Fully Occluded condition, as depicted in Figure 5.4, participants were required to wear an opaque cloth that wrapped around the bottom and sides of the HMD. This cloth was intended to prevent the participants from viewing any peripheral information that may provide feedback during their return walk. Otherwise, this condition did not differ from the blind walking protocol used in Experiment I.

5.2.2 Results

Experiment II aimed to determine if the improved performance seen in Experiment I was the results of a source of uncontrolled feedback, such as proprioceptive information
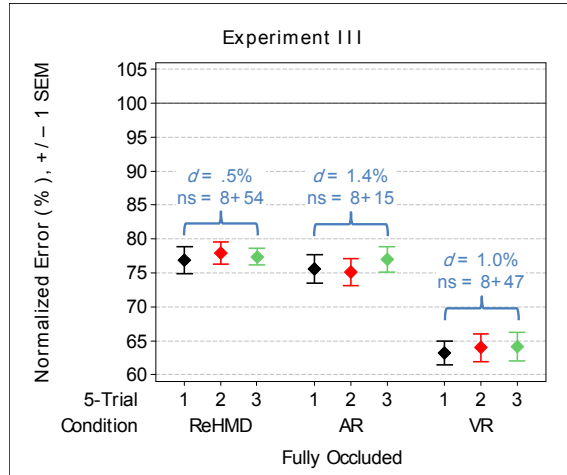
Figure 5.10

Distance judgments in the Fully Occluded ReHMD, AR, and VR conditions.

gained by walking the judged distance twice or peripheral visual information. By systematically removing the possible sources of feedback, one would expect no adaptation to occur in the suspect condition. Otherwise, one could assume that participants were independently modifying their blind walking behavior as reported in [28]. An analysis of variance and effect size calculation reveals that participants in the Extended condition continue to significantly adapt through the course of the experiment, while participants in the Fully Occluded condition did not (Extended: $F(2,14) = 14.496$, $p = 0.000$, $d = 14.7\%$; Fully Occluded: $F(2,14) = 0.111$, $p = 0.896$, $ns = 8+47$, $d = 1.0\%$). Figure 5.8 clearly shows that observers in the Extended condition exhibited significant adaptation, indicating that proprioception is an unlikely source of feedback. This seems to indicate a relationship between the observed adaptation and the presence of peripheral visual information. It is also worth noting that the mean normalized error in the Fully Occluded condition is

63.8%. This puts the underestimation observed in this condition firmly in the range that has been widely observed in numerous other VR studies [16, 31, 38, 40–43].

5.3 Experiment III

Experiment II established that the source of the implicit feedback that influenced the results of Experiment I and likely influenced Jones et al. [14] was peripheral visual information seen through a small gap below the HMD, between the HMD and the participants' face. However, Experiment II only established that this effect occurs in purely VR environments. One of the motivations of Jones et al. [14] was to determine if the underestimation effects typically seen in virtual environments also occur in augmented environments. The relationship between distance judgment errors in augmented environments is not as well studied as virtual environments and is somewhat conflicting [14, 36, 37]. Depth cue theory seems to indicate that the more cue rich an environment is, the more accurately distances should be judged [3]. Given that the augmented environment used in these experiments consisted of a virtual stimulus presented in a real-world environment, one would expect that the available cues would allow for more accurate depth judgments than in a purely virtual environment. This is somewhat indicated, but not significantly so, in the results of Experiment I. However, given the findings of Experiment II, one must ask if these results were also influenced by the presence of the uncontrolled peripheral visual information. The current experiment aims to answer this question by studying a Fully Occluded AR and ReHMD condition.

5.3.1 Method

For this experiment, 16 naive participants were recruited from the general university population and either received course credit or monetary compensation for their participation. Eight participants experienced each condition in a between-subjects design. Participants wore the same opaque cloth depicted in Figure 5.4, and the procedures very closely mimicked those in the Fully Occluded condition discussed in Experiment II. In the AR condition, participants observed a virtual stimulus presented in a real-world environment (Figure 5.1). In the ReHMD condition, participants saw no computer generated imagery, but instead viewed a real-world stimulus placed in the same real-world environment, as seen through the optical see-through window of the HMD. For both conditions, the optical see-through window was closed before the participants performed the return walk.

5.3.2 Results

An analysis of variance and calculated effect size indicated that the improved normalized error observed in Experiment I is not expressed in either the ReHMD or AR conditions when the periphery is restricted (ReHMD: $F(2,14) = 0.119$, $p = 0.889$, $ns = 8+54$, $d = 0.5\%$; AR: $F(2,14) = 0.317$, $p = 0.733$, $ns = 8+15$, $d = 1.4\%$). A somewhat remarkable finding is that the ReHMD and AR conditions did not significantly differ from each other ($F(1,14) = 0.110$, $p = 0.745$, $ns = 16+1100$). These findings are clearly visible in Figure 5.10, which for comparison purposes also shows the Fully Occluded VR condition from Experiment II. When comparing distance judgments in the Fully Occluded AR condition (75.9%) to those recorded in Experiment I for the Real condition (95.3%), we find

that they are significantly different ($F(1,16) = 24.139$, $p = 0.000$). This seems to establish that the underestimation effect exists in augmented environments, but to a lesser degree than seen in virtual environments. Perhaps an even more interesting finding is that the Fully Occluded ReHMD condition (77.4%) also differs significantly from the Real condition (95.3%) in Experiment I ($F(1,16) = 28.129$, $p = 0.000$). These results are consistent with those reported in Creem-Regehr et al. [2] where participants viewed a real-world environment through field-of-view restricting goggles with a horizontal field-of-view of $42°$. This field-of-view exactly matches the horizontal field-of-view of the HMD used in the experiments described in this document. Creem-Regehr et al. [2] found that participants significantly underestimated distances (78.9%[2]) when restricted field-of-view was coupled with restricted head movements. Though, in the current experiment, participants' head movements were not restricted, they were instructed to look directly at the stimulus during the viewing phase of the blind walking task. These findings are also quite similar to those reported in Willemsen et al. [41] where participants significantly underestimated distances (85.4%[3]) when viewing a real-world scene through a mock-HMD.

5.4 Experiment IV

Experiments II and III established that the addition and subtraction of peripheral visual information seen through the gap below the HMD has a strong effect on distance judgments. However, it is unclear if this facilitation was due to participants being able

---

[2]These normalized error values were derived from the figures presented in Creem-Regehr et al. [2].

[3]These normalized error values were derived from the figures presented in Willemsen et al. [41].

to localize their position through this gap or if visual information, such as optical flow, is correcting their spatial or motor perception. Experiment IV aims to answer this question by introducing a Partially Occluded condition. In this condition, the opaque occluder is replaced with a semi-opaque cloth through which luminance changes can be detected but shapes cannot be resolved.

5.4.1 Method

Sixteen naive participants were recruited from the general university population and either received course credit or monetary compensation for their participation. Both AR and VR viewing conditions were studied in Experiment IV. Eight participants experienced each condition in a between-subjects design. Other than the use of a semi-opaque cloth, the experimental procedures used in Experiment IV exactly mimic those used in Experiment III.

5.4.2 Results

An analysis of variance and calculated effect size indicated that participants in the AR condition significantly improved their distance judgments over time, but their VR counterparts did not (AR: $F(2, 14) = 7.399$, $p = 0.006$, $d = 8.3\%$; VR: $F(2, 14) = 0.287$, $p = 0.755$, $ns = 8+24$, $d = 1.3\%$). These results are depicted in Figure 5.11a. The result that no adaptation was seen in the VR condition while it was apparent in the AR condition was somewhat confusing. At the end of all experimental sessions, participants undergo an informal debriefing where they discuss their experiences in the experiment with the
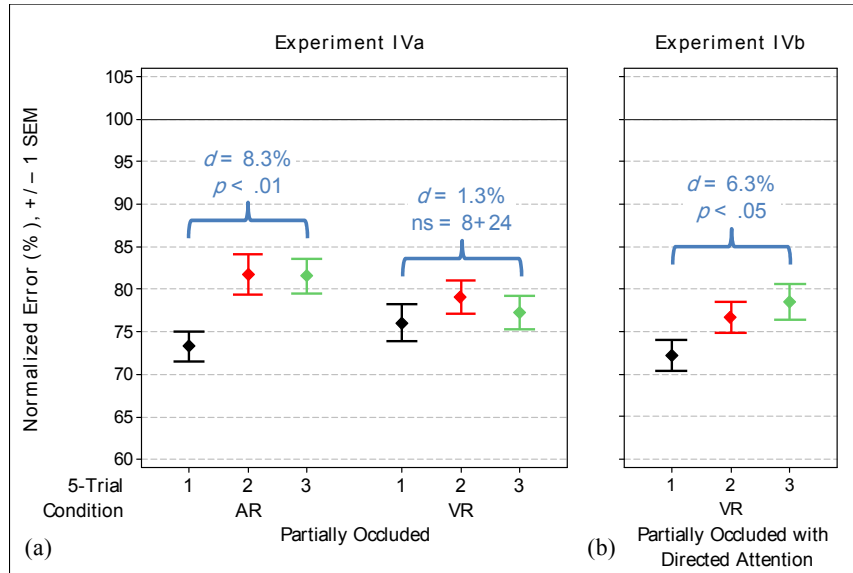
Figure 5.11

Distance judgments in the Partially Occluded and Directed Attention conditions

experimenters. The experimenters noted that participants in the VR condition typically remarked that they noticed the glow of the backlight of the HMD's display elements on the return walk while none of the participants in the AR condition made this remark. It is worth noting that all of the return walk conditions are identical for both the AR and VR conditions; no graphics are displayed and the optical see-through window is closed. This seems to informally indicate that participants in the VR condition may be more narrowly directing their attention to the screen area, possibly due to the novelty of the virtual environment. This hypothesis prompted an extension to Experiment IV where participants in the VR condition were explicitly instructed to attend to their periphery during the return walk. This extension was referred to as Experiment IVb. Eight more participants were recruited for this new condition. As seen in Figure 5.11b, these participants exhib-

ited significantly improved distance judgments with time when directed to attend to their periphery ($F(2, 14) = 4.106$, $p = 0.040$, $d = 6.3\%$).

5.5   Experiment V

Experiments I through IV seem to strongly indicate that cues in the periphery may heavily influence distance judgments. However, there was one important factor to address, the presence of the surveyor's tape used to measure the walked distance. This tape was present during all experiments and was centrally located in the experimental environment. Additionally, this white tape contrasted heavily with the environment's dark brown carpet. The presence of this tape in the periphery could have offered a very strong cue to lateral optical flow. Though this does not invalidate the theory that flow-base cues in the periphery aid in distance judgments, its presence provides us with the opportunity to substantially decrease the saliency of these peripheral optical flow cues by removing the tape. Removing this tape would allow the examination of distance judgments and possible improvements in an even more cue deprived environment than the Partially Occluded condition seen in Experiment IV. This new condition was referred to as the "No Tape" condition.

5.5.1   Method

Seven naive participants were recruited from the general university population and either received course credit or monetary compensation for their participation. These participants viewed a completely virtual environment using the same occluder configuration described in Experiment IV. Since the VR condition consistently yielded the strongest
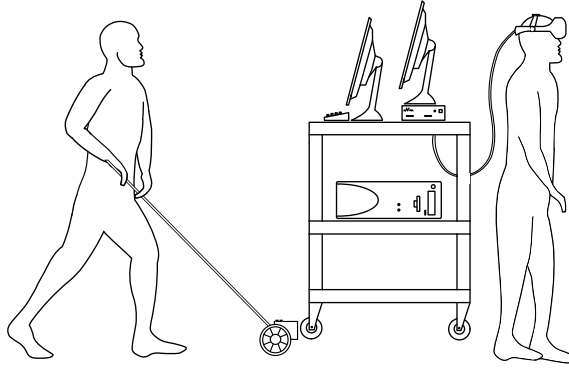
Figure 5.12

Measuring judgment distances with the rolling measuring wheel

adaptation effects, participants in this experiment were only presented with a virtual environment. For this reason, the additional instruction, as described in Experiment IV, to attend to the periphery was included. The only difference between this condition and that presented in Experiment IV is the absence of the white surveyor's tape in the experimental environment. As the surveyor's tape was used to measure the walked distance in all previous experiments, a new measurement method needed to be used. For this experiment, a digital rolling measuring wheel was used to measure the distance walked by the participants, depicted in Figure 5.12. Once the participants completed their blind walk they were asked to stand at that position until instructed to turn around, typically only a few seconds. During this pause, an experimenter would roll the measuring wheel from the participants' starting position to the tip of their toes. This distance was then recorded as the judged distance.
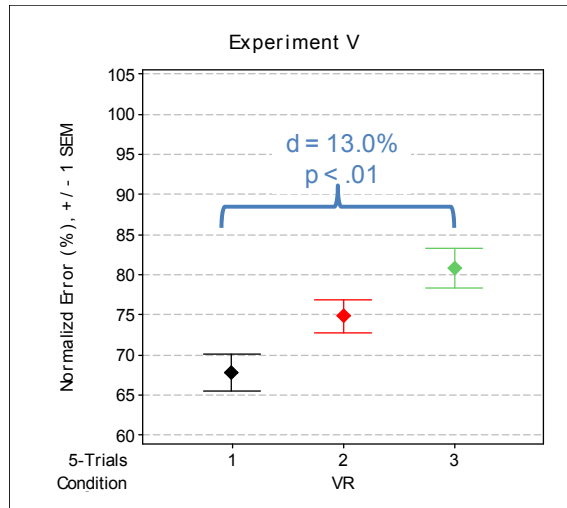
Figure 5.13

Distance judgments with no tape in the periphery
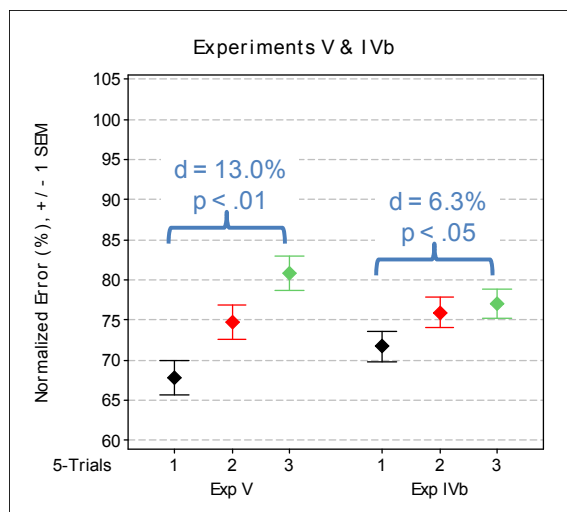
Figure 5.14

A comparison of Experiment IVb and Experiment V

5.5.2 Results

An analysis of variance and calculated effect size indicated that participants in this experiment also exhibited significantly improved distance judgments over time ($F(2,12)$ = 14.43, $p = 0.001$, $d = 13.05\%$). Figure 5.13 shows these results. This is an interesting result as it indicates that even when the most salient feature in the periphery is removed, participants were still able to gather sufficient information to significantly improve their performance. An even more interesting result can be seen in Figure 5.14, comparing the Directed Attention condition of Experiment IV and the current experiment. An analysis of variance comparing these two experiments failed to detect a significant difference between participant distance judgments with an without the additional optical flow information provided by the surveyor's tape ($F(2,28) = 0.024$, $p = 0.976$). This seems to indicate that only a very small amount of optical flow information in the periphery is sufficient to dramatically improve distance judgments in a virtual environment.

5.6 Experiment VI

The previous experiments indicate that visual information in the extreme edges of the periphery have a strong influence on distance judgments. However, the question remains open if these effects persist in a narrow field-of-view with complete occlusion of the extreme periphery. Experiment VI aimed to answer this question by restricting the peripheral information available on the return walk to the central $48°$ x $40°$ in a condition referred to as Restricted FOV.

5.6.1 Method

Sixteen naive participants were recruited from the general university population and either received course credit or monetary compensation for their participation. Both AR and VR viewing conditions were studied in Experiment VI. Eight participants experienced each condition in a between-subjects design. The procedures for this experiment are very similar to the Fully Occluded condition used in Experiment III except that the optical see-through window is opened on the return walk. This was done to restrict the visual cues on the return walk to only those available through the see-through window.

5.6.2 Results
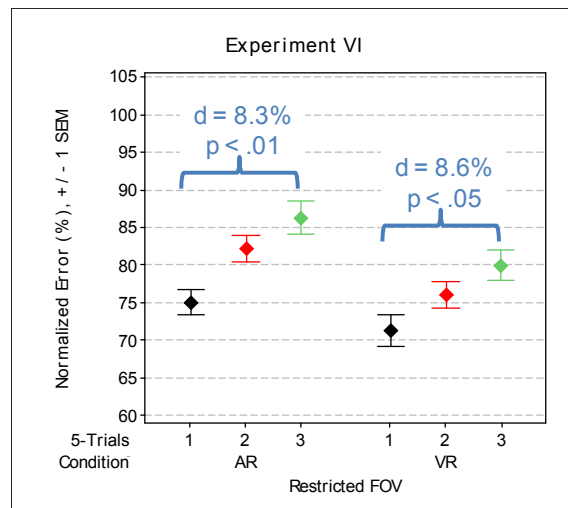
Figure 5.15

Distance judgments in the Restricted FOV condition

The results of this experiment, Figure 5.15, seem to indicate that when vision is restricted to the central field-of-view on the return portion of the blind walking task that participants continue to improve their distance judgments over time. An analysis of variance and calculated effect size clearly demonstrate this effect in both the AR and VR conditions (AR: $F(2,14) = 8.305$, $p = 0.004$, $d = 11.29\%$; VR: $F(2,14) = 4.071$, $p = 0.040$, $d = 8.65\%$;). These results are somewhat comparable to those reported in Bruder et al. [1] where artificially exaggerated optical flow in on-screen periphery in a head-worn virtual environment enable participants to perceive their movement in the VE as equivalent to that in the real-world. This could imply that optical flow cues available in a typical virtual environment, as compared to those available in the real-world, may lack sufficient visual fidelity to accurately convey the motion expressed in the participants movements. This seems especially applicable to those flow cues available in the foveal, parafoveal, and near peripheral regions where participants' visual resolution is greatest.

CHAPTER 6

CONCLUSIONS

Experiment I aimed to determine whether or not the unusual reduced underestimation

seen in Jones et al. [14] were a result of transfer effects due to the within-subjects exper-

imental design. In Experiment I the general trend of reduced underestimation persisted

despite the between-subjects design. However, a striking pattern of increased accuracy

emerged as Experiment I progressed. Since this pattern seems to be time dependant, a

within-subjects design would hamper its detection as a result of presenting multiple en-

vironments in succession. Even so, this pattern was still visible, though to a much lesser

degree, in Jones et al. [14]. Experiment I indicated that between- and within-subjects

experimental designs for exploring cross-environmental distance judgments would likely

yield mutually comparable results but would make time- or repetition-dependant effects

difficult to detect.

The pattern of increase accuracy as a function of time, seen in Experiment I, was an

interesting and somewhat troublesome result, as it indicated that participants were aug-

menting their distance judgments with uncontrolled feedback. Experiment II examined

two possible sources of implicit feedback: the blind walking task itself and a gap below

the HMD. However, neither source seemed a likely candidate. If the walking task was in-

fluencing the participants' judgments, one would expect their performance to decrease in

variability while remaining centered around the originally underestimated position. However, participants' judgments rapidly approached veridical throughout the course of the experimental session, which typically lasted approximately 20 minutes. If the participants were acquiring visual information from the gap below the HMD, there is very little that is visible to use as feedback. Typically, participants would only be able to see the carpet of the experimental environment. Regardless, the ability to see any part of the surrounding environment leaves open the possibility that participants are able to localize their position within the environment during the return walk portion of the blind walking task. Another possibility is that optical flow cues seen in the lower periphery were affecting either the participants' perception of the environment or their movement within the environment. Rieser et al. [32] performed an elegant series of real-world walking tasks where participants were exposed to varying rates of optical flow while walking at different speeds, and this study demonstrated that the calibration of participants' movements can be greatly affected by changing the relationship between optical flow and walking speed.

The results of Experiment II revealed that participants failed to improve their performance when the gap below the HMD was completely occluded. This directly indicate the gap was the source of the uncontrolled feedback. This raised the possibility that observers were simply visually localizing their position during the experiment. Given the amount of the environment which was visible through the gap, this seemed an unlikely possibility. However, there was also the possibility that participants could be calibrating their movements based on peripheral optical flow. Experiment IV to strongly indicated the latter. In this experiment, participants' views were partially occluded, enabling them to detect

luminance changes through the occluder but not resolve their location. In this experiment, participants in the AR condition still exhibited improved performance, but participants in the VR condition only improved when they were specifically instructed to attend to their periphery. This was both an unexpected and exciting finding as it implies that the attention of participants in the VR condition was more narrowly focused than their AR counterparts. All participants were naive and had never experienced HMD-based virtual reality prior to this experiment. Given that this is a very unfamiliar experience, it seems plausible that the novelty of the virtual environment may be narrowing their attention to the screen area, thereby preventing VR participants from utilizing peripheral information as effectively as the AR participants.

Experiment V removed the most prominent feature from the periphery in order to determine whether the flow cues provided by a high contrast stimulus was necessary to significantly improve distance judgments in a virtual environment. This feature was a white surveyor's tape against a dark brown carpet background that was present in Experiments I through IV. The tape was removed for this experiment, which otherwise exactly mimicked the conditions from Experiment IV. Experiment V revealed that removing the white surveyor's tape from the environment did not significantly alter the pattern of improved distance judgments seen in Experiment IV. It is important to note that the only visual cues available in this condition were small luminance variations visible through the partially occluded periphery. This showed that minimal peripheral cues are necessary to facilitate adaptation in virtual environments.

The results of Experiments I through V indicated that visual cues available in the extreme periphery have a strong influence on the accuracy of depth judgments. However, the question remained open whether or not a comparable set of visual information is equally effective when restricted to the same field-of-view as the typical screen area of an HMD, including narrowed-peripheral, parafoveal, and foveal areas. Experiment VI aimed to determine if similar effects appeared under these circumstances. The results of this experiment indicate that the adaptation effects seen in Experiments I through V persist when participants' views, on the return walk, are restricted to the same field-of-view as the screen area.

Experiment III sought to answer a question originally posed by Jones et al. [14]: does the underestimation effects seen in virtual environments also exist in augmented environments? To test this, the gap below the HMD was occluded and participants performed blind walks to a virtual object seen in the real world. Participants did significantly underestimate distances, judging stimuli distance to roughly 76% of their actual distance. This is intriguing, but even more so when compared to distance judgments to real stimuli seen through the HMD. Experiment III demonstrated that distance judgments in an augmented environment was not significantly different from those in a real-world environment when viewed through the HMD. This indicates that the majority of the distance information acquired while viewing an object comes from the surrounding environment and not the object itself. This also implies that augmented environments may not suffer as greatly from the underestimation effects typically seen in virtual environments. The bulk of the underestimation in the ReHMD and AR conditions seems to be caused by viewing the en-

vironment through the HMD. This is likely due to the restricted field-of-view and inability to display visual information in the periphery, which these and several other experiments have indicated to be an important factor in improving distance judgments [2, 41, 44].

Figure 6.1 shows the results of Experiments I through VI ordered by environment and visual cue availability, with the most visually restricted conditions to the left and the least to the right. The most obvious trend that emerges when looking at Figure 6.1 is that distance judgments significantly improve as peripheral visual information is less restricted (VR: $F(4, 35) = 6.397$, $p = 0.001$; AR: $F(2, 23) = 3.824$, $p = 0.037$; ReHMD: $F(1, 16) = 40.853$, $p = 0.000$). It is also interesting to note that, excepting the AR Standard and ReHMD Standard conditions, that the initial underestimation seen across all the experimental conditions did not significantly differ ($F(10, 77) = 1.219$, $p = 0.293$).

Perhaps the most exciting finding, illustrated in Figure 6.2, is that any amount of visual information in the periphery, regardless of the amount, greatly increases depth judgment accuracy in virtual environments as compared the Fully Occluded condition ($F(4, 35) = 6.39$, $p = 0.001$). Additionally, when participants' periphery was fully occluded, but their central field-of-view was completely unobscured on the return walk, they did not perform any better than when the only visual cues available were those seen in the partially occluded periphery ($F(3, 27) = 0.129$, $p = 0.942$). This clearly indicates that visual information available at the extreme edges of a participant's natural field-of-view is extremely important when performing accurate distance judgments. The benefit received, compared the amount of visual information available, is very disproportionate, with slight shifts in luminance at the edges of the periphery having as much influence on distance judgments
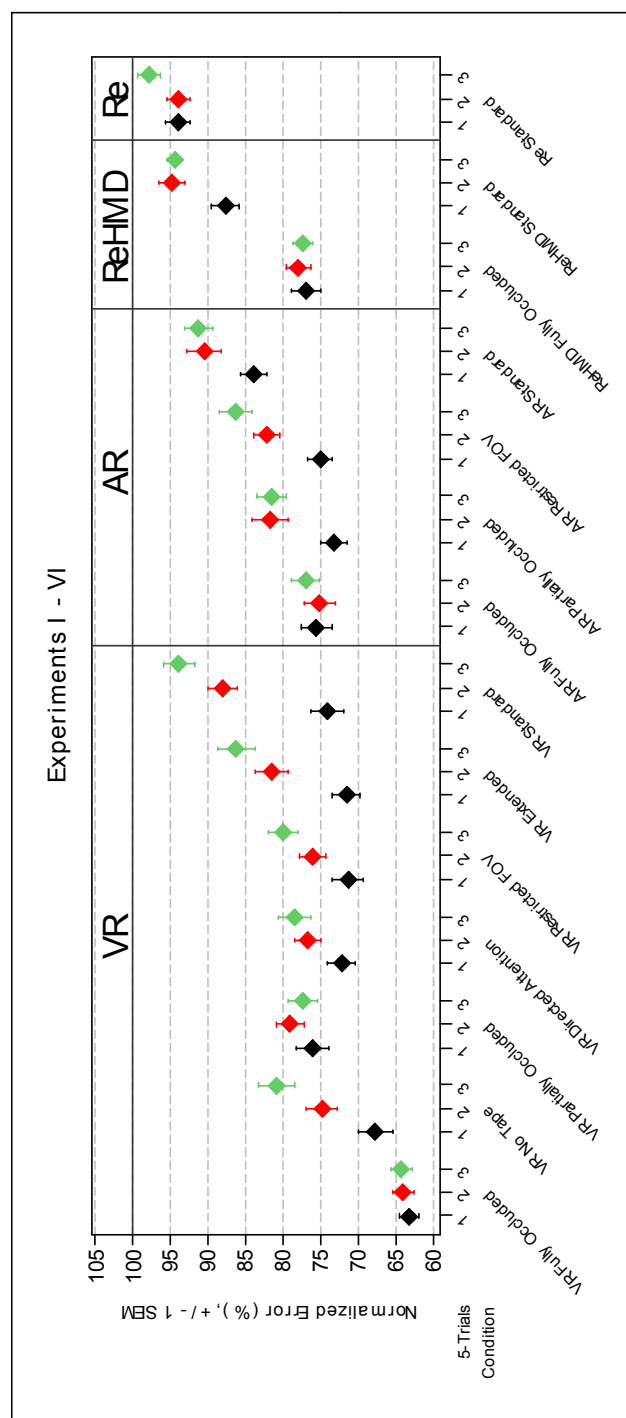
Figure 6.1

Results of Experiments I - VI

as viewing an otherwise unobstructed scene through a $48°$ x $40°$ window. However, when considering that a human's natural horizontal field-of-view is nearly $180°$, it does not seem nearly as surprising that our distance judgments suffer greatly when we lose nearly 75% of the visual information to which we would normally have access.
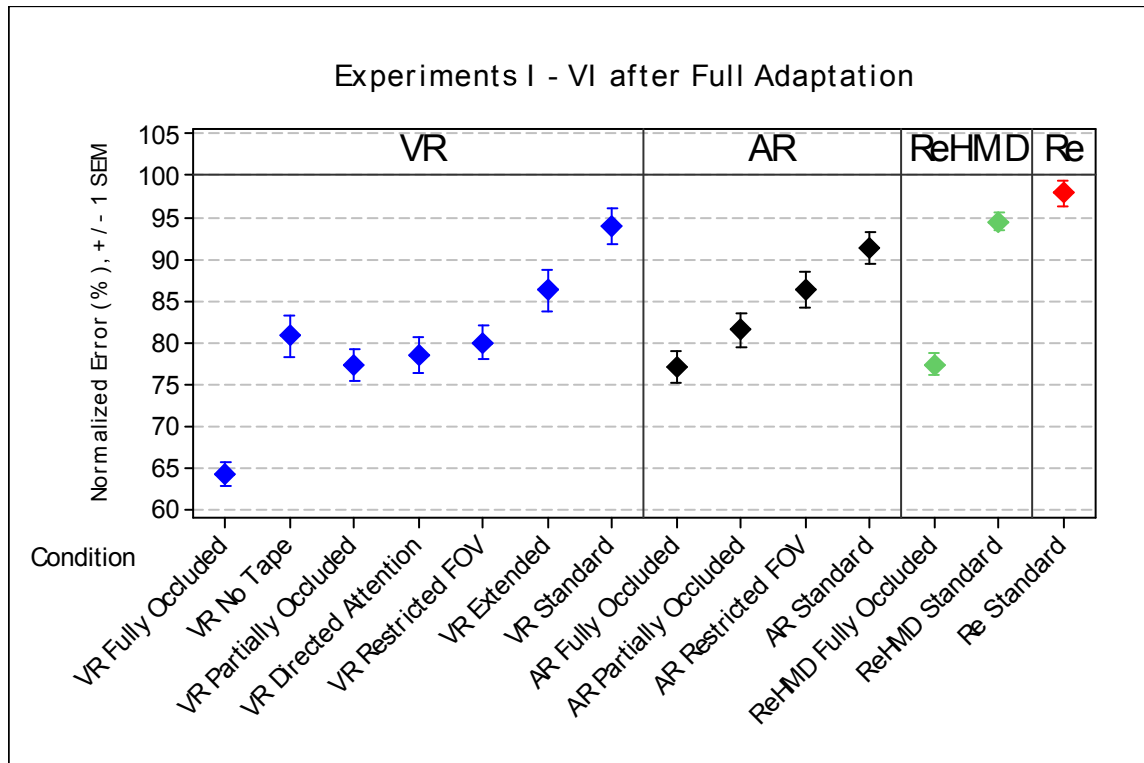
Figure 6.2

Results of Experiments I - VI after adaptation

REFERENCES

[1] G. Bruder, F. Steinicke, and P. Wieland, "Self-motion illusions in immersive virtual reality environments," *Virtual Reality Conference (VR), 2011 IEEE*, march 2011, pp. 39–46.

[2] S. H. Creem-Regehr, P. Willemsen, A. A. Gooch, and W. B. Thompson, "The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual indoor environments," *Perception*, vol. 34, 2005, pp. 191–204.

[3] J. Cutting, "How the eye measures reality and virtual reality," *Behavior Research Methods*, vol. 29, 1997, pp. 27–36, 10.3758/BF03200563.

[4] F. Durgin, F. L.F., J. Lewis, and K. Walley, "Perceptuomotor adaptation: More than meets the eye," *43rd annual meeting of the Psychonomic Society*, Kansas City, USA, 2002.

[5] F. Faul, E. Erdfelder, A.-G. Lang, and A. Buchner, "G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences," *Behavior Research Methods*, vol. 39, 2007, pp. 175–191, 10.3758/BF03193146.

[6] J. M. Foley, N. P. Ribeiro-Filho, and J. A. D. Silva, "Visual perception of extent and the geometry of visual space," *Vision Research*, vol. 44, no. 2, 2004, pp. 147 – 156.

[7] J. J. Gibson and E. J. Gibson, "Perceptual learning: Differentiation or enrichment?," *Psychological Review*, vol. 62, no. 1, 1955, pp. 32–41.

[8] W. C. Gogel and J. D. Tietz, "Oculomotor adjustments in darkness and the specific distance tendency," *Perception & Psychophysics*, vol. 13, 1973, pp. 284–292.

[9] C. D. Harvey, F. Collman, D. A. Dombeck, and D. W. Tank, "Intracellular dynamics of hippocampal place cells during virtual navigation," *Nature*, vol. 461, 2009, pp. 941–946.

[10] H. H. Hu, A. A. Gooch, W. B. Thompson, B. E. Smits, J. J. Rieser, and P. Shirley, "Visual Cues for Imminent Object Contact in Realistic Virtual Environments," *Visualization Conference, IEEE*, vol. 0, 2000, p. 32.

[11] V. Interrante, B. Ries, and L. Anderson, "Distance Perception in Immersive Virtual Environments, Revisited," *Virtual Reality Conference, IEEE*, vol. 0, 2006, pp. 3–10.

[12] *IS-1200 VisTracker Datasheet*, InterSense, Inc., Billerica, Massachusetts, USA, October 2011.

[13] J. A. Jones, *Egocentric depth perception in optical see-through augmented reality*, doctoral dissertation, Mississippi State University, Mississippi State, Mississippi, 2006.

[14] J. A. Jones, J. E. Swan II, G. Singh, E. Kolstad, and S. R. Ellis, "The Effects of Virtual Reality, Augmented Reality, and Motion Parallax on Egocentric Depth Perception," *Proceedings of the Symposium on Applied Perception in Graphics and Visualization (APGV 2008)*, Los Angeles, CA, USA, 2008, pp. 9–14.

[15] E. Klein, J. Swan, G. Schmidt, M. Livingston, and O. Staadt, "Measurement Protocols for Medium-Field Distance Perception in Large-Screen Immersive Displays," *Virtual Reality Conference, 2009. VR 2009. IEEE*, march 2009, pp. 107 –113.

[16] J. M. Knapp, *The Visual Perception of Egocentric Distance in Virtual Environments*, doctoral dissertation, University of California, Santa Barbara, Santa Barbara, California, 1999.

[17] M. Lappe, M. Jenkin, and L. Harris, "Travel distance estimation from visual motion by leaky path integration," *Experimental Brain Research*, vol. 180, 2007, pp. 35–48, 10.1007/s00221-006-0835-6.

[18] *Leica TPS800 Series User Manual*, Leica Geosystems, Canton St. Gall, Switzerland, October 2011.

[19] M. Livingston, C. Zanbaka, I. Swan, J.E., and H. Smallman, "Objective measures for the effectiveness of augmented reality," *Virtual Reality, 2005. Proceedings. VR 2005. IEEE*, march 2005, pp. 287–288.

[20] J. M. Loomis and J. M. Knapp, *Virtual and Adaptive Environments: Applications, Implications, and Human Performance*, chapter 2, CRC Press, Mahwah, NJ, 2003, pp. 21–46.

[21] R. Miller, "Nonius alignment apparatus for measuring vergence," *American Journal of Optometry & Physiological Optics*, vol. 64, 1987, pp. 458–466.

[22] R. Miller, R. G. Pigion, and M. Takahama, "The effects of ingested alcohol in accomodative, fusional, and dark vergence," *Perception & Psychophysics*, vol. 39, 1986, pp. 25–31.

[23] R. Miller and M. Takahama, "Arousal-related changes in dark focus accommodation and dark vergence," *Investigative Ophthalmology & Visual Science*, vol. 29, 1988, pp. 1168–1178.

[24] M. Mon-Williams and J. R. Tresilian, "Some recent studies on the extraretinal contribution to distance perception," *Perception*, 1999, pp. 167–181.

[25] T. D. Nguyen, T. G. Grechkin, J. Cremer, J. K. Kearney, and J. M. Plumert, "Effect of measurement setting in judging traveled distance: additional evidence for under-estimation of distance in virtual environments," *Applied Perception in Graphics and Visualization (APGV 2010)*, Los Angeles, California, USA, 2010, p. 159.

[26] C. B. Owen, J. Zhou, A. Tang, and F. Xiao, "Display-Relative Calibration for Optical See-Through Head-Mounted Displays," *Mixed and Augmented Reality, IEEE / ACM International Symposium on*, 2004, pp. 70–78.

[27] D. Owens and H. Leibowitz, "Oculomotor adjustments in darkness and the specific distance tendency," *Perception & Psychophysics*, vol. 20, 1976, pp. 2–9.

[28] J. Philbeck, A. Woods, J. Arthur, and J. Todd, "Progressive locomotor recalibration during blind walking," *Attention, Perception, & Psychophysics*, vol. 70, 2008, pp. 1459–1470, 10.3758/PP.70.8.1459.

[29] L. Phillips and V. Interrante, "A little unreality in a realistic replica environment degrades distance estimation accuracy," *Virtual Reality Conference (VR), 2011 IEEE*, march 2011, pp. 235 –236.

[30] A. R. Richardson and D. Waller, "The effect of feedback training on distance estimation in virtual environments," *Applied Cognitive Psychology*, vol. 19, no. 8, 2005, pp. 1089–1108.

[31] A. R. Richardson and D. Waller, "Interaction With an Immersive Virtual Environment Corrects Users' Distance Estimates," *Human Factors*, vol. 49, no. 3, 2007, pp. 507–517.

[32] J. J. Rieser, H. L. Pick, D. H. Ashmead, and A. E. Garing, "Calibration of human locomotion and models of perceptual-motor organization," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 21, 1995, pp. 480–497.

[33] A. Rizzo, B. Newman, T. Parsons, J. Difede, G. Reger, K. Holloway, G. Gahm, R. McLay, S. Johnston, B. Rothbaum, K. Graap, J. Spitalnick, and P. Bordnick, "Development and Clinical Results from the Virtual Iraq Exposure Therapy Application for PTSD," *Virtual Rehabilitation International Conference, 2009*, 29 2009-july 2 2009, pp. 8 –15.

[34] J. P. Rolland, W. Gibson, and D. Ariely, "Towards Quantifying Depth and Size Perception in Virtual Environments," *Presence: Teleoperators and Virtual Environments*, vol. 4, no. 3, 1995, pp. 24–49.

[35] F. Steinicke, G. Bruder, K. Hinrichs, S. Kuhl, M. Lappe, and P. Willemsen, "Judgment of natural perspective projections in head-mounted display environments," *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology*, Kyoto, Japan, 2009, pp. 35–42.

[36] J. Swan, A. Jones, E. Kolstad, M. Livingston, and H. Smallman, "Egocentric Depth Judgments in Optical, See-Through Augmented Reality," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 13, no. 3, may-june 2007, pp. 429–442.

[37] J. Swan, M. Livingston, H. Smallman, D. Brown, Y. Baillot, J. Gabbard, and D. Hix, "A Perceptual Matching Technique for Depth Judgments in Optical, See-Through Augmented Reality," *Virtual Reality Conference, 2006*, march 2006, pp. 19–26.

[38] W. B. Thompson, P. Willemsen, A. A. Gooch, S. H. Creem-Regehr, J. M. Loomis, and A. C. Beall, "Does the quality of the computer graphics matter when judging distances in visually immersive environments," *Presence: Teleoper. Virtual Environ.*, vol. 13, October 2004, pp. 560–571.

[39] M. Tuceryan and N. Navab, "Single point active alignment method (SPAAM) for optical see-through HMD calibration for AR," *Augmented Reality, 2000. (ISAR 2000). Proceedings. IEEE and ACM International Symposium on*, 2000, pp. 149 –158.

[40] P. Willemsen, M. B. Colton, S. H. Creem-Regehr, and W. B. Thompson, "The effects of head-mounted display mechanics on distance judgments in virtual environments," *Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, New York, NY, USA, 2004, APGV '04, pp. 35–38, ACM.

[41] P. Willemsen, M. B. Colton, S. H. Creem-Regehr, and W. B. Thompson, "The effects of head-mounted display mechanical properties and field of view on distance judgments in virtual environments," *ACM Transactions on Applied Perception*, vol. 6, March 2009, pp. 8:1–8:14.

[42] P. Willemsen and A. Gooch, "Perceived egocentric distances in real, image-based, and traditional virtual environments," *Virtual Reality, 2002. Proceedings. IEEE*, 2002, pp. 275–276.

[43] B. G. Witmer and J. Sadowski, Wallace J., "Nonvisually guided locomotion to a previously viewed target in real and virtual environments.," *Human Factors*, vol. 40, 1998, pp. 478–488.

[44] B. Wu, T. Ooi, and Z. He, "Perceiving distance accurately by a directional process of integrating ground information," *Nature*, vol. 428, mar 2004, pp. 73–77.

APPENDIX A

BLIND WALKING INSTRUCTIONS

All participants, in all experiments described in this document, were given the following instructions by an experimenter prior to beginning the experimental task. After these instructions, all participants received five practice trials of the experimental task in an adjacent hallway prior to the beginning of data collection.

> Visually directed walking, or blind walking, is walking to an object with your eyes closed. We will show you an object and ask you to look at it until you feel like you have a very good sense of where it is located in space and feel as though you can walk to it with your eyes closed. For instance, when you wake up at night and need to get a drink of water, it's probably completely dark in your bedroom, but you can always walk to the light switch even though you cannot see it. We want you to have the same kind of sense of the target object's position.

> We are going to show you an object on the floor and when you feel that you can walk to the object with your eyes closed, let me know that you are ready and I will ask you to walk forward. I want you to walk until you feel as though the tips of your toes are at the center of where the object was located. The object will be removed from the scene, so there will be no chance of you stepping on or tripping over it. When you stop walking, you may open your eyes, but continue to look forward. There will be a brief pause. After the pause, you will be asked to turn to your right and return to your starting position[1].

> When you reach the starting position you will be asked to turn around to your left and then center yourself in the hallway. There will be a ridge on the floor that you can feel beneath your feet. This ridge marks your starting position. When centering yourself, please stand with the arches of your feet on the ridge.

> At all points during the experiment an Experimenter will be walking in front of you and another will be walking behind you in case you lose balance or become disoriented. We will also stop you before you walk too close to the walls.

> We will also be asking you to close your eyes very frequently. It is important that you do not "squint" your eyes and face when trying to close your eyes. This uses a lot of muscles in your face and they will get tired quickly. This might cause you to accidentally open your eyes. We suggest that you

---

[1]In the Fully and Partially Occluded conditions, the participants are further instructed to reach out to the right and place their hand on the rolling cart, which would guide them back to the starting position

simply relax your face and close your eyelids and perhaps lower your head slightly. Now we will give you some practice with blind walking so you can get comfortable walking with your eyes closed.

APPENDIX B

EXPERIMENTAL TASK SCRIPT

The following script was used to coordinate the blind walking procedure for all experiments discussed in this document.

Experimenter 1:

(In the Real and Real+HMD conditions, Experimenter 1 indicates the position of the target to Experimenter 2)

Experimenter 1:

Open your eyes, observe the object, and tell me when you are ready.

Observer:

(The observer views the object and indicates readiness)

Experimenter 1:

Close your eyes and walk forward.
(Experimenter 1 pushes the cart behind the observer, allowing the extension power cable to unroll behind him)

Experimenter 2:

(Experimenter 2 removes the target object during the Real and Real+HMD conditions)

Observer:

(The observer walks forward with eyes closed)

Experimenter 1:

(Experimenter 1 walks behind the observer, watching for signs of disorientation)

Experimenter 2:

(Experimenter 2 walks in front of the observer, watching for signs of disorientation)

Observer:

(Stops when they believe they are at the target distance)

Experimenter 1:

(Experimenter 1 records the distance walked by the observer)

Turn to your right[1].

(This prevents the observer from getting tangled in the HMD and tracker cables)

Return to your starting position, center up, and close your eyes.

Experimenter 2:

(Experimenter 2 pushes the cart back to the starting position)

Experimenter 1:

(Experimenter 1 rolls up the extension power cable, ensuring that the cart does not roll over the cable)

---

[1]In the Fully and Partially Occluded conditions, the participants are further instructed to reach out to the right and place their hand on the rolling cart, which would guide them back to the starting position