1-1-2019

The Effects Of Differing Optical Stimuli On Depth Perception In Virtual Reality

McKennon B. McMillian

THE EFFECTS OF DIFFERING OPTICAL STIMULI ON DEPTH PERCEPTION IN VIRTUAL REALITY

A Thesis
presented in partial fulfillment of requirements
for the degree of Master of Engineering
in the Computer and Information Science
The University of Mississippi

by

McKennon Brice McMillian

December 2019

ABSTRACT

It is well documented that egocentric depth perception is under estimated in virtual reality more often than not. Many studies have been done to try and understand why this underestimation happens and what variables affect it. While this underestimation can be shown consistently the degree of underestimation can strongly differ from study to study, with as much as 68% to as low as 6% underestimation, Jones et al. (2011, 2008); Knapp (1999); Richardson and Waller (2007). Many of these same studies use blind walking as a tool to measure depth perception. With no standardized blind walking method for virtual reality existing differing blind walking methods may cause differing results. This thesis will explore how small changes in the blind walking procedure affect depth perception. Specifically, we will be examining procedures that alter the amount of ambient light that is visible to an observer after performing a blind walk.

DEDICATION

This is dedicated to my mother Angela Ruth McCluskey, who passed way last summer. She always pushed me to do more and without her I would not be where I am today.

ACKNOWLEDGEMENTS

I would like to thank all the members of the Hi5 lab for the help and support they have given me over the years, the Computer Science Department for giving me the tools to go do something productive with my life, and my committee.

I would like to give a special thanks to Dawn Wilkins and Joey Carlisle. Without their help I would have nowhere near the grab on computer science that I have now. William Panlener, Ethan Luckett, and Lee Jenkins for being close friends and helpingh me every step of the way. My Wife Meghan McMillian for helping me through the last six months and making me keep moving forward.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Underestimation of judged distances along the ground plane in virtual environments is well documented and it has been shown that objects in virtual environments appear closer than they actually are, Jones et al. (2011, 2008); Knapp (1999); Richardson and Waller (2007).

Verbal estimates of egocentric distance are usually too error-prone at distances greater than several meters, yet blindfolded walking to previously viewed targets can be surprisingly accurate. This process is known as blind walking. Blind walking has been found to be significantly more accurate than verbal reports, Andre and Rogers (2006), Swan et al. (2007). Blind walking has become a common measure of perceived target location, and was used in this study to do just that. A typical blind walking procedure consists of showing a subject an object and having them view the object until they feel confident in their ability to blindly walk to the object's position, then blindfolding or otherwise obscuring the vision of the subjects, a common example is simply asking subjects to close their eyes, and having them blindly walk to the object.

This study used blind walking with four different methods of obscuring subjects' vision. A virtual object was presented on the ground plane of a Virtual Reality (VR) environment. A Vive Pro Head Mounted Device (HMD) was used as the VR medium. Subjects' vision was obscured by completely whiting or blacking out the screen of the HMD. Subjects were also instructed to perform the walk with their eyes open 50% of the walks.

Blind walking in VR is a very common method to judge egocentric depth perception int the medium field (from 2m to 20m), but different studies implement blind walking in different ways. Some studies simply black out the screen and do not specify whether the subject's eyes are open. Others keep the environment running or gray-out the screen and ask the subjects to close their eyes. Some do both, some do neither. This is a problem for comparability between VR studies. To make the problem worse, most people simply say the did "blind walking" when writing about their study and give no mention of what the subjects did between trials of the blind walking. Work by Jones et al. 2011 and 2013 provided evidence that the return walk is actually very important, Jones et al. (2011, 2013). The goal of this thesis is to explore how the variables affect blind walking in VR studies.

CHAPTER 2

RELATED WORKS

It is well known that vision is a strong guide for humans to help us preform a large number of tasks. Moving around our environment, being one of the most crucial, Fuchs (1899); Gogel (1993); Gogel and Tietz (1973); Thomson (1983). Humans have evolved so that, visually guided movements, including walking, jumping, pointing, or reaching, are fairly well calibrated with regard to our environment, Ellis and Menges (1998); Gogel (1993); Klein et al. (2009). With vision and movement being so interlocked a large amount of research has become available over the years. A good deal of study has been done on depth cues' effectiveness across the three ranges of distances: near, medium, and far field. Near field extends from 0 to 2 meters away from us, medium field is from 2 to 20 meters, and far field is anything past 20 meters, Cutting (1997).

Human vision is incredibly fine toned to help us navigate in our day to day lives. Due to this, people are quite accurate in judging distances across a large range of distances using a variety of methods, Loomis et al. (1992); Philbeck et al. (2008); Singh et al. (2010). Blind walking is one such technique. Blind walking is a commonly used tool for accurately measuring distance judgement in the medium field, Knapp and Loomis (2004); Philbeck et al. (2008); Thomson (1983). In blind walking, observers are typically shown a target located along the ground plane at a given distance. They then either close their eyes or have their vision blocked and walk until they feel they are standing at the target distance. The walked distance is then used as an indication of the perceived distance of the target position.

Blind walking in the real world is a very accurate measurement of depth perception, Interrante et al. (2006); Knapp and Loomis (2004). However, in virtual reality environments, distances are underestimated, as if the target object appears closer than it is. This has repeatedly been shown to be the case, Interrante et al. (2006); Singh et al. (2010); Creem-Regehr et al. (2005). The reason for this underestimation is unknown.

Studies by Jones et al. (2012, 2011, 2013) have been conducted studying how different variable affect perception in virtual environment. The variables include field of view, differing return walks, and visual stimulation in different areas within subjects' field of view. Visual stimulation was applied in the field of view outside of the area that was considered important in virtual environments. It was found that this did indeed have an impart on depth judgement. When visual stimuli was removed completely a negative affect was shown to occur.

The work done by Jones et al. (2013) presents finding the inspired this thesis. The findings in this paper indicated that when depriving participants of visual information during the return walk appreciably hinders the participant's ability to judge distances with blind walking. This was done by having subjects preform blind walking tasks, but on the return walk subjects had a black veil placed over their faces. It is important to note that this experiment was ran in a real world environment, not a virtual one. Placing the veil over the faces of the subjects reduced judgement distance accuracy to 81.74%. This is much lower then than average reported by several other studies, Interrante et al. (2006); Jones et al. (2011); Knapp and Loomis (2004); Loomis et al. (1992). A variation of this was looked at, in virtual environments, in this study.

CHAPTER 3

SETUP

3.1 Environment Design

The environment used for this study consisted of both a physical location and a virtual model of this location. The physical location was a 10.89x3.89 meter room. When subjects were present in the room all non critical objects were removed to make the virtual environment as close to the physical room as possible. Non critical items include clocks, tools, extra chairs, boxes, and any other objects that were not added to the virtual environment. A physical marker was placed on the floor where subjects needed to stand. The markers consisted of a white piece of paper with a rod taped down the middle of it. The paper was used because it gave an audio cue when stepped on, and the rod was used to help subjects orient themselves. This marker was needed because subjects could not see their surroundings while blind walking and the next target distance was not shown until they were oriented correctly.

The virtual environment was a scale replica of the physical environment. The walls were colored using images taken of the room, the floor was done the same way. Any objects in the room that could not be removed (such as furniture), or were needed, were custom modeled and textured. These included three white tables, a brown desk, four computer monitors, three computer towers, a white board, and three windows. There was a white cube in this environment that was not present in the physical room. This is what the subjects use as a point of reference to walk to. When the environment was first started it was rendered where the subject would stand at the start.

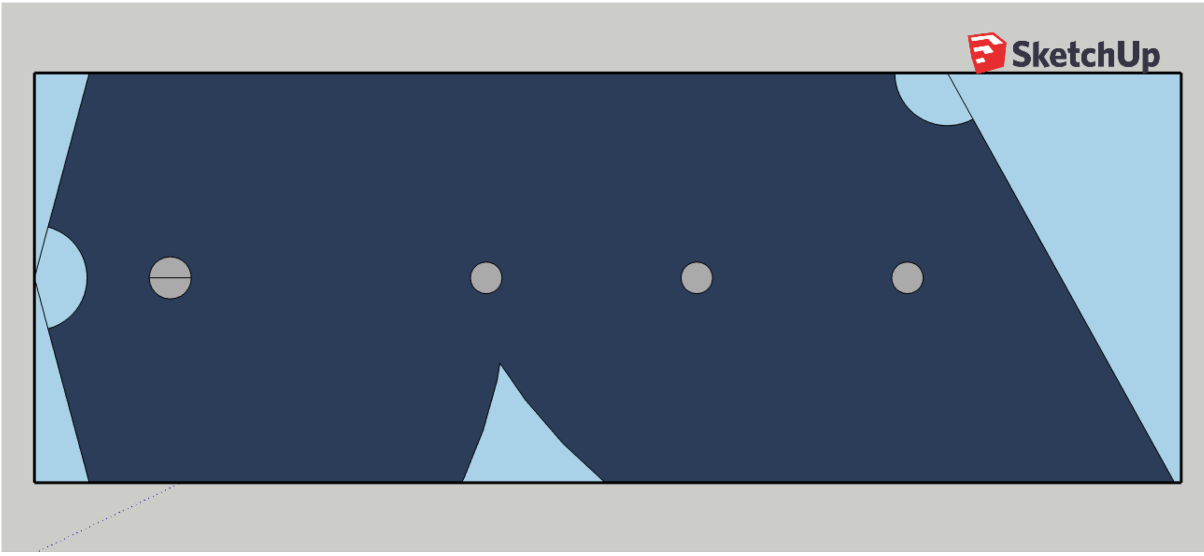

Figure 3.1. This shows the tracking area and basic layout of the physical room. The dark blue area is where the HMD can be tracked reliably. The light blue area is untracked. The larger circle to the left is where the subject starts each trial. The three smaller circles are the target distances of 3, 5, and 7 meters from left to right respectively.

3.2 Equipment Setup

This study was run using the HTC Vive Pro as the head mounted display (HMD). The Vive Pro uses two base stations to motion track the HMD and the controllers. The headset itself is wired into a desktop computer and has a five meter cable.

The area that needed to be tracked was quite large, 8.5x3.89 meters, and required a special base station layout. The full room did not need to be tracked. The first base station was placed directly behind where the subjects would be standing, 1.27 meters back, and raised 2.32 meters to the ceiling and pointed slightly down so that it could track the top of the HMD. The second base station was placed 7.4 meters in front of where the subjects would be standing and 1.4 meters to the left and at a height of two meters. It was then rotated 40 degrees so that it could track a subject even if they walked past it. This layout gave an effective tracking area of 9.25x3.89 meters. The reason this tracking area differs from the the full size of the room is due to limitation of the hardware. This layout gives subjects approximately 2 meters of over walk coverage. See figure 3.1 for clarification.

The desktop was placed in the middle of the tracking area, and out of the subjects walking path, to allow the five meter cable to cover 9.4 meters.

3.3 Virtual Environment Setup

The virtual environment required setup before each subject was tested. This was done to ensure the environment was stable and correct for each subject and to get accurate measurements from the equipment.

An experimenter (myself or an assistant) would wear the HMD and stand on the physical marker placed in the room, they would then make sure they were standing on the cube in the virtual room. A controller was used to make sure the height of the subject relative to the room was correct. This was done by placing the controller on the floor of the physical room and manually adjusting the height (Y-axis) of the virtual floor until the controller model resting on top of the floor in both environments.

The same procedure was preformed for the walls to make sure the physical and virtual environments matched as close as possible. An experimenter would wear the HMD and hold the end of a controller to the corner of a small protruding support column, while another experimenter manually adjusted the virtual space (X-axis, Z-axis) so that the controller tip was at the same point in both environments.

The rotation of the room around the Y-axis would sometime need to be reset. This could happen if a base station was moved or if the Vive had been setup for a different environment previously. When this happened it was corrected by running room setup on the Vive. Room setup is a feature of the Vive Pro, that sets the bounding area. While running room setup you must calibrate the floor by placing two controllers on the level floor, set the forward direction of the environment by pointing a Vive controller in the direction you would like to be forward, and set the bounding area for the environment. The virtual environment was set up so that it would align the back wall with the line that was associated as being forward. By using the walls of the physical room to draw perfectly straight lines for the bounding area, it was always ensured to align the back walls of each environment. This made sure that the rotation was the same.

CHAPTER 4

PROCEDURES

There were 13 subjects ranging in age from 20 to 54, 11 male and 2 female. All subjects had either normal or corrected to normal vision. Each subject ran a total of 36 trials. There were four viewing conditions (eyesopen-black-screen, eyesopen-white-screen, eyesclosed-black-screen, eyesclosed-white-screen), three target distances (3, 5, 7 meters), with three repetitions for each combination of these. The trial order was different for all subjects, randomly generated before tests began, but all subjects saw all possible trials. Each trial was guaranteed to be different then the one before it. No trial was repeated until at least one other trial with a different viewing condition or distance was ran. This was done to help minimize transfer effects between viewing conditions. This is the circumstance where viewing one condition may affect how you respond to the following condition. Though this does not prevent transfer effects, it helps to distribute any such effects evenly between all conditions.

Subjects were instructed to stand on a paper marker placed on the floor. They then put on the Vive Pro (HMD), and adjusted the screen inside it to match their interpupillary distance. This was done with a knob on the base of the HMD and was done to improve the clarity of the virtual environment. Interpupillary distance (IPD) is the distance from the center of one pupil to the center of the other. IPD in this instance was self measured by the subject by adjusting until the surrounding displayed in the HMD was the most clear. This was done in a separate virtual environment to limit exposure to the test environment.

The subjects were told that there would be a cube, on the floor of the virtual environment, placed in front of them at several distances. They were to judge where the cube was placed and blindly walk to where they judged the cube to be after the cube was removed from view with either their eyes open or closed. The distinction between the two was given verbally from an experimenter. The way the virtual environment was removed from view differed from trial to trial. The screen would either whited out or blacked out; either way, the entire virtual world was hidden from the subject (not just the cube).

The subjects were instructed on whether to walk with eyes open or eyes closed for the current trial while they were standing on the paper marker, facing the direction of the cube. When they indicated that they knew where the cube was located and were ready to walk, the screen was either whited out or blacked out and they walked to where they perceived the cube to be. They then indicated verbally that they had reached where the cube was and the actual distance they had walked was recorded using a script running in Unity 3D. The script recorded the starting position of the HMD and the position of the HMD when the subject indicated they had reached the cube. A second experimenter then approached the subject as they turned around, the subjects extended an arm that the experimenter placed on their shoulder. On the return walk subjects were instructed to keep their eyes closed if they had just finished an eyes closed trial or to keep their eyes open if they had just finished an eyes open trial. The subjects were then lead on the return walk back to the paper marker, turned around, shown the next target distance, and the process was repeated.

CHAPTER 5

TECHNICAL

5.1 Python

Python 3.6 was used to generate the input files for each subject. An input file is a Comma Separated Values (CSV) file that consists of 36 line containing the Subjects ID number, line number, occlusion condition, test distance, trial repetition. Each trial was guaranteed to be different then the one before it. This is known as a restricted random shuffle. This script takes a CSV file of all 36 trials that will be ran. This file is then converted into four separate lists based on the occlusion condition. Those list are then sorted by the distance value and split into three lists each, based on the distance value. The script then goes into a loop that runs 17 times. At the beginning of each loop a subject ID is generated with the iteration of the loop. This drops into another loop that builds CSV file using a list called master_list. This loop break when master_list becomes 36 elements long or after 100,000 iterations. If 100,000 iterations were reached, a CSV file for that was loop was generated, but a tag was added that indicated a manual review was needed. This happened 2 times out of 16. The master_list was built by generating a random integer from 1 to 12, each of the integers were mapped to a specific list of occlusion conditions paired with a distance. If the number was ever generated twice in a row then another number was generated. If a number had been generated the times previously, another number would be generated.

Python was also used to help analyze the data after collection was finished. It was used to calculated the error and the number of over and under walks. This was several short custom scripts.

5.2 Unity 3D

Unity 3D is the engine that the virtual environment was rendered in. Unity is a game engine that is used for games design, animation, engineering, and much more. Several custom models were created in Unity, using simple shapes such as cubes and planes. Unity also allowed access to many useful assets such as shaders, models, prefabs, and pre-made lighting options. Along with pre-made prefabs, Unity allows for the creation of custom prefabs. This is a very helpful feature when designing any virtual environment. Unity's Prefab system allows you to create, configure, and store a GameObject (anything inside a Unity environment) complete with all its components, property values, and child GameObjects.

5.2.1 Models

Two large 10x10 meter planes were created. One black and one white; these were the occlusion planes used to black or white the screen. They were placed half a cm away from the near clipping plane of the camera so that nothing could be seen when either plane was rendered. These planes used a special shader that was not affected by light in the virtual environment. Clipping planes let you exclude some of a scene's geometry and view or render only certain portions of the scene. Each camera object has a near and a far clipping plane. Objects closer than the near clipping plane or farther than the far clipping plane are invisible to the camera.

A small 0.5 meter cube was created as a virtual target for the subjects to try and walk to. The cube also served as a guide for where the subjects should stand at the start of the excitement.

The other models in the environment were there because the physical objects could not be easily moved and we wanted to two environments to be as similar as possible. There were three white tables, a brown desk, four computer monitors, three computer towers, a white board, and three windows. All of these were modeled by hand other than the whiteboard; it was an asset that existed previously.

5.2.2 SteamVR Prefab

One of the most important tools offered by Unity was the SteamVR prefab. This prefab is what allows the Vive Pro to work with Unity. The center of the prefab can be considered the virtual subject. It contains Gameobjects for both controllers and for the headset, this is known as the head. The head contains the camera that is used to display the virtual environment onto the screens in the HMD. All tracked subject movement is relative to the starting position of this prefab. This makes for easy adjustments to fix positioning errors in the environment. When setting up or fixing a small alignment problem, experimenters were able to simply adjust the position of the prefab instead of the rest of the environment. The SteamVR prefab also comes with a set of scripts that allow Unity to take in and use the motion tracking data from the Vive Pro and base stations.

The position of the head Gameobject is the position used for the data collection script for this study. This script was attached as a child of the head Gameobject and also allowed for direct control what the subject was viewing.

5.3 C#

The majority of this work was done using C#. This is one of the two languages that are compatible with the Unity 3D engine out of the box, the other being JavaScript. A script was written and attached to the camera object in the SteamVR prefab. The script was responsible for controlling the virtual environment and for collecting and outputting the data as a CSV file.

On startup, the script would temporarily remove the black and white near field occlusion planes. This hid them away at the start of a run so that the subjects could get into position. The input files was then split into and array of strings; each string was a trial.

When N was pressed on the keyboard, position of the headset was stored and the occlusion planes were hidden if either were previously rendered. This also moved the cube to the correct distance in the virtual environment and output a message to the experimenters about the current trial number and condition was.

When the space bar was pressed the appropriate occlusion plain was rendered. This is when the subject would walk to the target object.

When S was pressed the starting location and current location of the headset were passed to a distance function that calculated the actual distance walked. This also updated the trial number, allowing the next trial to be loaded.

5.4 Vive Pro

The Vive Pro is the virtual reality tool used for this study. It consists of a Head Mounted Display (HMD), two base stations, and two controllers. The controllers were only used for calibration. The HMD uses two Dual AMOLED 3.5” diagonal OLED screens with a resolution of 1440 x 1600 pixels per eye (2880 x 1600 pixels combined) and a 90 Hz refresh rate. The head set has a 100 degree field of view. The Vive Pro also contains a G-sensor, gyroscope, and infrared (IR) sensors to help motion track its position.

5.5 Base Stations

The Vive Pro uses devices known as base stations to motion track its exact location. Each base station has a 150-degree horizontal field of view and a 110-degree vertical field of view. These devices work by emitting a vertical and horizontal infrared scan of the physical environment. These scans have a lower bound measured range of 4.55 meters. These scans are picked up by the HMD using special IR sensors places all over the front, top, and sides of the face plate of the HMD. The base stations have a 0.5 meter area directly in front of them where tracking is less than ideal. Two base stations were used in this study.

5.5.1 Issues

One issue with the Vive Pro is that there is no way to get a purely black screen. When the black occlusion plane is active the screen should be completely black; however, there are dark grey artifacts that appear on the screen as a result of the screens being back lit. They are not very pronounced, only being noticed by the experimenters after long viewings of the black occlusion plain. These artifacts are present in other environments when the displays should be completely black as well. These artifacts are not the result of the lighting in the virtual environment; the occlusion plain used a shader that did not reflect any light from the environment.

5.6 Pipe Stat

The data collected over the course of this study was just a list of conditions, line numbers, and distances. Pipe Stat, or | Stat, is a command line stats package written in C, https://garyperlman.com/stat/. Pipe Stat was used to find statically significant changes in the data that allowed for pointed, direct investigation.

Pipe Stat was installed from source on a Linux machine. The original program was written over 30 years ago and the Make file had some issues building correctly. The cflags had to be set to c89 standards and the binary install location needed to be set correctly in order to make the commands global on the machine.

Pipe Stat takes a TSV (Tab Separated Values) file and pipes it in to the a binary package called colex, this is then piped in another binary package called anova. ANOVA stands for analysis of variance and it is a collection of statistical models and their associated estimation procedures used to analyze the differences among group means in a sample. An option grep can be added to help refine the analysis. Pipe Stat then runs a statistical analysis on every possible combination of the columns passed to it.

Colex takes arguments in the form of a list of comma separated numbers. These numbers map to the columns of data in the TSV file. The first number must be the the subject ID column and the last number must be the metric that the user wants to compare the other arguments against. All other arguments are optional.

The anova package takes a list of comma separated string and arguments where each string is a description of the the information contained in a the column, where the first argument of colex must be the subject ID the first argument of anova will describe that.

This is an example of what a typical evaluation would look like.

```
more all.tsv|colex 1 3 4 8|grep "closed"|anova -pp sid rep cond err
```

The output for Pipe Stat is a rudimentary graph of the of the data passed to it. These graphs were just error bars but helped to visualize the data. It also output the mean, f values, and p values for every combination passed to it. The f and p values were used to determine if a statically large variance in the data was present.

5.7 StatsVis

StatsVis is a data plotting tool created by a student at the University of Mississippi for their senior project. It used python to generate a number of different types of graphs.

CHAPTER 6

DATA ANALYSIS

There were a total of 468 trials. Only 79 were over walks. That is 16.9%. 3 meter walks were 19.23% over walks, 5 meters were 12.18% over walks, and 7 meter walks were also 19.23% over walks.

Under walking was very common, with 369 trials being underestimated. That is 78.85%. 3 meter walks were underestimated 74.36% of the time, 5 meter walks 84.62%, and 7 meter walks were underestimated 77.56% of the time. Figure 6.1 - 6.3 illustrates this. This is what was expected and shows that the blind walking procedures follow previous trends in VR depth perception.

Conditions had statically significant effects on judged distance. Figure 6.5 shows there is a clear impact on distance judged by the trial condition, F(1-13) = 15.756, p = 0.013. This figure also shows that the more restrictive, eyes closed, condition was more accurate than eyes open. This would seem to contradict previous findings that optical flow in the peripheral increases the accuracy of distance judgment in the medium field, Jones et al. (2011). But this is not the whole story. Due to the artifacts introduced by the restricted random shuffle we can not look at the raw data; we must instead look at improvement within each condition. We can then compare the amount of improvement. Figure 6.4 show there was improved accuracy over repetitions overall in all conditions. Looking at the improvement over repetitions between white and black viewing conditions, figure 6.6, we see that there is not significant difference between the two. They both show the same rate of improvement. However, when you look at the error over repetitions, shown in figure 6.7, you eyes closed condition improves very little, while the eyes open condition have a very significant improvement. The eyes open

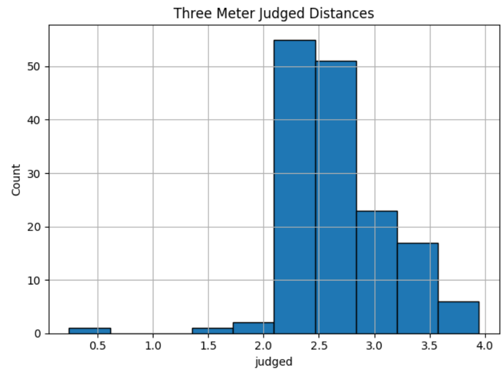

Figure 6.1. The judged distance at three meter target distance.

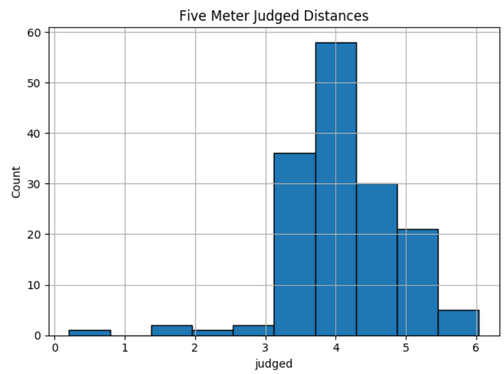

Figure 6.2. The judged distance at five meter target distance.

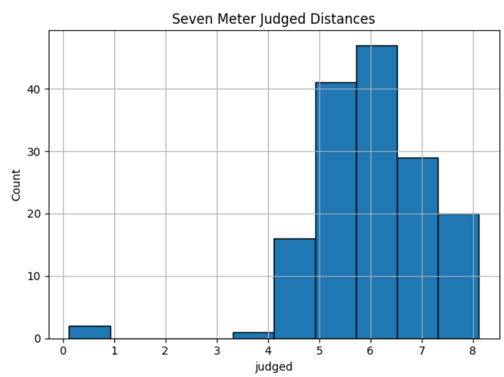

Figure 6.3. The judged distance at seven meter target distance.

condition shows the most improvement, by far out, of the four. This is interesting since it implies that providing exclusively light does not bring about improvements to distance judgement. This may imply there are some physiological or perceptual functions being seen here that are specifically associated with the eyes. Further study would be needed on this

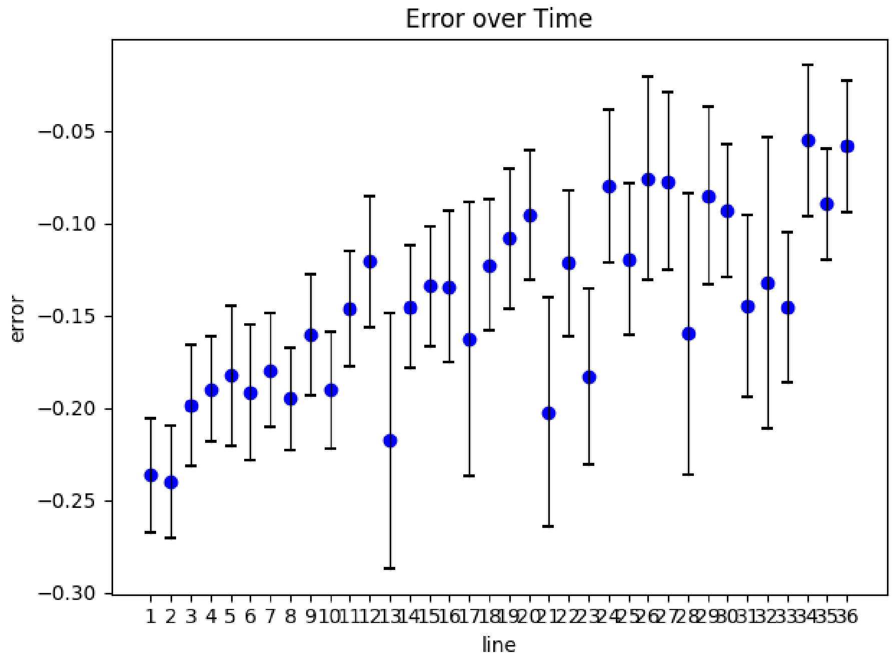

Figure 6.4. This shows a clear overall increase in accuracy over time.

subject to draw any useful conclusions.

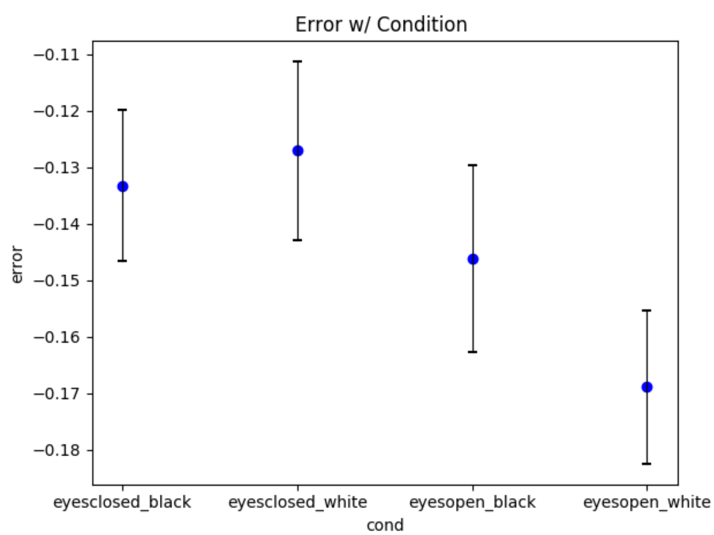

Figure 6.5. The error of judged distance over all subjects with respect to full trial condition. Error bars represent the Standard Error of the Mean for each condition.

Out of all four condition variables (white, black, open, closed) only closed was not measured to have changed significantly over all trials.

- Open: F(1,13) = 12.909, $p < 0.001$
- Closed: F(1,13) = 2.475, $p = 0.105$
- Black: F(1,13) = 4.603, $p = 0.020$
- White: F(1,13) = 5.268, $p = 0.013$

This is shown in figures 6.6 and 6.7. This would suggest that the color of the display does not affect depth perception in VR. At least, not in any measurable way that this study has found.

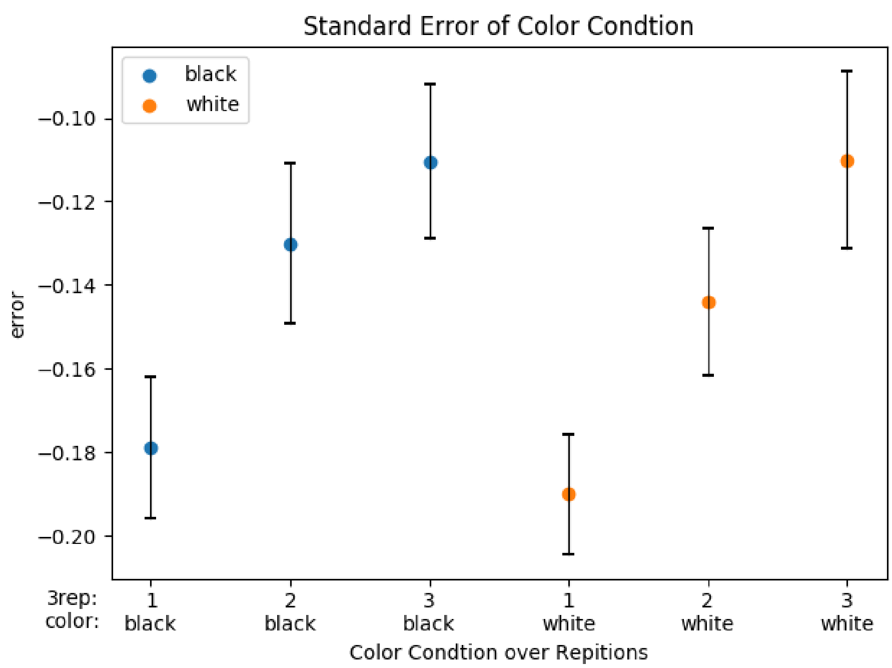

Figure 6.6. The mean error of judged distance over each of the three repetitions of the two screen color conditions. Error bars represent the Standard Error of the Mean for each condition.

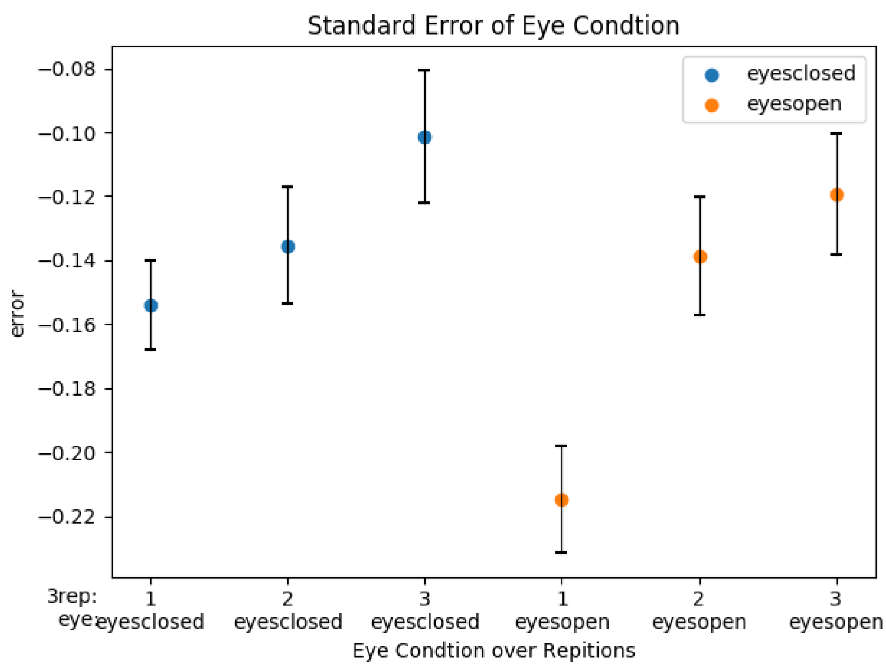

Figure 6.7. The mean error of judged distance over each of the three repetitions of the two eye conditions. Error bars represent the Standard Error of the Mean for each condition.

CHAPTER 7

DISCUSSION

This work shows that there is indeed a significant difference between blind walking with eyes open verses eyes closed and that there is not a significant difference in improvement when differing amounts of light are applied from the display. This could be seen by looking at the improvement over time of each of the four viewing conditions. This is important because regardless of the amount of optical stimuli this difference in improvement is still present. When preforming a blind walking procedure it is important to understand how different variable affect subjects' performance so that you can understand the data you collect and draw the proper conclusions. Previous work has found that this change in blind walking performance when optical stimuli is removed is also present in real world conditions, Jones et al. (2013).

A high level of detail is required when describing blind walking procedures in virtual reality. This is important in making an experiment repeatable by others. Every step should be described in detail, not just novel concepts, from the type of HMD used to the conditions of the return walk. The more detailed the better.

In retrospect a different random ordering algorithm could have been used to produces a better spread of the order of the trials that each subject saw. A potential replacement could be to use Latin Squares. A Latin square is an n x n array filled with n different symbols, each occurring exactly once in each row and exactly once in each column. This would produces n unique orderings. Using this method would show all conditions at different stages of accommodation. For best results n subject should be ran to get a more complete data set. Doing this would allow conclusions to be made base on the raw data points and comparing the actual values represented in the figures.

There are several experiments that that can be ran as future work using details fro. this study.

- Repeat this study using a Latin Square ordering of the trials the subjects saw.
- Remove the screen color condition to better study the condition that affected improvement.
- Reverse the eye condition on the return walks. If the subject had eyes closed for the walk to the target object have them open their eyes for the return walk and vice versa.

These are only a few potential studies that could show intriguing results based off this study.

BIBLIOGRAPHY

BIBLIOGRAPHY

Andre, J., and S. Rogers (2006), Using verbal and blind-walking distance estimates to investigate the two visual systems hypothesis, *Perception & Psychophysics*, 68(3), 353–361.

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

Cutting, J. E. (1997), How the eye measures reality and virtual reality, *Behavior Research Methods, Instruments, & Computers*, 29(1), 27–36.

Ellis, S. R., and B. M. Menges (1998), Localization of virtual objects in the near visual field, *Human factors*, 40(3), 415–431.

Fuchs, E. (1899), Text book of ophthalmology (duane trans.).

Gogel, W. (1993), The analysis of perceived space in foundations of perceptual theory.

Gogel, W. C., and J. D. Tietz (1973), Absolute motion parallax and the specific distance tendency, *Perception & Psychophysics*, 13(2), 284–292.

Interrante, V., B. Ries, and L. Anderson (2006), Distance perception in immersive virtual environments, revisited, in *IEEE Virtual Reality Conference (VR 2006)*, pp. 3–10, IEEE.

Jones, J. A., J. E. Swan II, G. Singh, E. Kolstad, and S. R. Ellis (2008), The effects of virtual reality, augmented reality, and motion parallax on egocentric depth perception, in *Proceedings of the 5th symposium on Applied perception in graphics and visualization*, pp. 9–14, ACM.

Jones, J. A., J. E. Swan, II, G. Singh, and S. R. Ellis (2011), Peripheral visual information and its effect on distance judgments in virtual and augmented environments, in *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization*, APGV '11, pp. 29–36, ACM, New York, NY, USA, doi:10.1145/2077451.2077457.

Jones, J. A., J. E. Swan II, G. Singh, S. Reddy, K. Moser, C. Hua, and S. R. Ellis (2012), Improvements in visually directed walking in virtual environments cannot be explained by changes in gait alone, in *Proceedings of the ACM Symposium on Applied Perception*, pp. 11–16, ACM.

Jones, J. A., J. E. Swan II, and M. Bolas (2013), Peripheral stimulation and its effect on perceived spatial scale in virtual environments, *IEEE transactions on visualization and computer graphics*, 19(4), 701–710.

Klein, E., J. E. Swan, G. S. Schmidt, M. A. Livingston, and O. G. Staadt (2009), Measurement protocols for medium-field distance perception in large-screen immersive displays, in *2009 IEEE Virtual Reality Conference*, pp. 107–113, IEEE.

Knapp, J. (1999), Visual perception of egocentric distance in virtual environments unpublished doctoral dissertation, *Department of Psychology, University of California at Santa Barbara, Santa Barbara, CA*, 93106.

Knapp, J. M., and J. M. Loomis (2004), Limited field of view of head-mounted displays is not the cause of distance underestimation in virtual environments, *Presence: Teleoperators & Virtual Environments*, 13(5), 572–577.

Loomis, J. M., J. A. Da Silva, N. Fujita, and S. S. Fukusima (1992), Visual space perception and visually directed action., *Journal of Experimental Psychology: Human Perception and Performance*, 18(4), 906.

Philbeck, J. W., A. J. Woods, J. Arthur, and J. Todd (2008), Progressive locomotor recalibration during blind walking, *Perception & Psychophysics*, 70(8), 1459–1470.

Richardson, A. R., and D. Waller (2007), Interaction with an immersive virtual environment corrects users' distance estimates, *Human Factors*, 49(3), 507–517.

Singh, G., J. E. Swan II, J. A. Jones, and S. R. Ellis (2010), Depth judgment measures and occluding surfaces in near-field augmented reality, in *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*, pp. 149–156, ACM.

Swan, J. E., A. Jones, E. Kolstad, M. A. Livingston, and H. S. Smallman (2007), Egocentric depth judgments in optical, see-through augmented reality, *IEEE transactions on visualization and computer graphics*, 13(3), 429–442.

Thomson, J. A. (1983), Is continuous visual monitoring necessary in visually guided locomotion?, *Journal of Experimental Psychology: Human Perception and Performance*, 9(3), 427.

VITA

Education:

May 2017, University of Mississippi — Oxford, MS Bachelor of Science Computer And Information Sciences

Experience:

Graduate Teaching Assistant — January 2017 to Current University Of Mississippi, (CSCI Dept)

IT Assistant/Database Architect — April 2018 to May 2019 University Of Mississippi, (School of Engineering IT)

Teaching Assistant — April 2015 to December 2017 University Of Mississippi, (CSCI Dept)