# Effects of Avatar Hand-size Modifications on Size Judgments of Familiar and Abstract Objects in Virtual Reality

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Zachary Patterson

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
MASTER OF SCIENCE

Pete Willemsen

June 2019

# Acknowledgements

First I would like to thank Dr. Pete Willemsen for the countless opportunities and support that have been given to me from my research assistantship to the endless hours of research discussion and bad puns. I would like to thank all the fellow students in the SIVE Lab over the last two years: Charles McGregor, Sumanth Kothapalli, Devin Cheek, and my closest friend at UMD Christianah Adigun. Additional thanks to Dr. Peter Peterson and Dr. Ed Downs for being on my committee and talking through my research with me. Extra thanks to all the Master's of Computer Science students at UMD from graduating classes of 2018, 2019, and 2020 for helping shape my experience at UMD. More thanks to all the students I helped teach across my four semesters as a teaching assistant; their kind remarks and positive feedback has meant a lot throughout the last two years. Finally, I would like to thank the MMAD Lab for providing a second Oculus Rift for running my experiment.

## Dedication

I would like to dedicate this thesis to my parents, Jim and Jody Patterson, and sister, Taylor Patterson, for being incredible role models and an excellent support system when I have needed them the most throughout my life. And I would also like to dedicate this thesis to my wonderful girlfriend Avamarie Brueggeman for always believing in me and inspiring me to be my best self.

Abstract

Many research studies have investigated spatial understanding within virtual environments, ranging from distance estimation, size judgments, and perception of scale. Eventually, this knowledge will help us to create virtual environments that better match our spatial abilities within natural environments. To further understand how people interpret the size of virtual objects, we present an experiment that utilizes a proprioceptive-based size estimation measure designed to elicit a three-dimensional judgment of an object's size using a box-sizing task. Participants viewed both abstract and familiar objects presented within action-space in a virtual environment and were asked to make an axis-aligned box the same size as the object they previously observed. A between-subjects manipulation modified a participant's avatar hand size to be either 80%, 100% or 120% of their measured hand size. Results indicate that the avatar hand size manipulation scales various factors of these size judgments in the three dimensions. Additionally, whether an object was abstract or a familiar size object produced distinctly different size judgments.

# Contents

# List of Tables

# List of Figures

# 1 Introduction

It's easy to imagine a future in which people are trained for complicated and dangerous scenarios such as military combat or fire containment through virtual reality simulations. We can even consider using virtual reality on a day to day basis for entertainment where people can explore foreign locations and experience things that aren't possible in our world. While these sound like things that aren't feasible yet with current devices, one of the larger problems in forming these tools is our understanding of how people view virtual environments. For training in virtual reality it would be extremely important to guarantee that the movements and body mechanics used in the training are identical to how the same action would be done in real life. It would also be paramount to make sure that the actions and responses to scenarios that the person is being trained to do are the same as what they should do in real life. Even for entertainment purposes, a feeling of immersion and realistic actions would cement a better user experience. The issue is that, for realistic responses to occur, they require the virtual training environment to be perceived correctly. Often when people make decisions about what action to take next, it relies on the world around them. Even an action as simple and mundane as climbing stairs will be treated slightly differently based on the incline of the stairs. If the stairs were to be perceived incorrectly, people would trip and fall or have unnatural reactions that would harm the immersion of the virtual environment.

While this may sound like a straight forward issue, it becomes more complicated as an experience is being crafted. The chief issue is that virtual environments are not perceived the same as real environments. This then compounds when you consider that perception of one element of an environment changes perception of the other elements. If an object is farther away, you might perceive it to be smaller. But if you were to just make the object bigger, how would that change the perception of how far away it is? Even as a baseline, to be able to answer this question would require a firm understanding on how the size and distance of the object are perceived without manipulation. However, knowing basic information about an environment such as how far things are and how large they are is essential to making decisions and actions involving those objects. To explore these questions I conducted a study that aims to look into how the size of virtual objects is measured and how that is affected by the perception of oneself in an immersive virtual environment(IVE). This had subjects changing the size of a box in terms of width, height, and depth, in order to perfectly fit it around an object previously displayed in front of the subject. The avatar representing the subject has hands that are either 80%, 100%, or 120% of the subject's actual hand size. The objects are categorized as either things the subject has likely seen and interacted with in everyday life, or things that have no discernible real world equivalent. The results showed significant differences between both estimated sizes of familiar and abstract objects, and estimated sizes when hands were enlarged. There were also unexpected results when hands were reduced to 80% size and with an object with no symmetry.

# 2 Background

## 2.1 Virtual Reality Perception

An important component of VR is user perception. How people perceive the virtual objects and environments that they are shown is a direct indicator of how people form an understanding of the scene. A user's understanding of the properties of a virtual object are based on how that object compares to known factors. For example, if a virtual ball is placed next to a real chair in AR the user has an idea of how large the ball is by comparing it to the chair. Similarly, if the shadows cast from the ball lie next to the shadow cast from the chair it helps convey the location of the ball. The perception of these simple properties is how people prepare for actions and make cognitive decisions involving the environment around them. It has been documented that the distance from the user of a virtual object is underestimated[Kel+18][CST15]. While researchers aren't certain why this occurs, it is widely believed to stem from a lack of visual cues that we expect to see based off of real life[Swa+07]. These visual cues are presumed to be[CV95]

- binocular disparity: The difference in location of an object seen by each of a person's eyes.
- binocular convergence: The amount an eye has to rotate in order to focus on an object.
- accommodative focus: The amount of tension an eye undergoes to change its

focal length needed to focus on an object.

- atmospheric haze: When objects are farther away they gain a slight blue tint as a result of water and dust in the air. This is typically seen with mountains, or other very large objects that are visible from great distances.

- motion parallax: When a person's head turns, objects in the background move across their field of view faster than closer objects.

- linear perspective and foreshortening: When parallel lines stretch into the distance they appear to converge at the horizon.

- occlusion: One object physically existing in front of another cues people that the object they can see is closer.

- height in the visual field: As objects move farther away, the viewing angle between the horizon and the object shrinks.

- shading: When light source positions are known, the shadows and illuminations of objects give us information on where they exist in relation to the light sources.

- texture gradient: As you move closer to an object, you can notice more detail in the object's surface.

Not all cues are equally important and some change our perception more than others[Luo+07]. The issue in recreating these cues is that it requires processing power intensive algorithms that need a great deal of environmental understanding. This is an issue of not only software but hardware as well. Providing resolution detailed enough to meld virtual objects with real objects and maintain a high frame rate is an extremely computationally heavy task. Not to mention for tools like AR, devices also need to constantly read and analyze the landscape. Identifying light sources and how they work with the virtual objects becomes increasingly difficult as environments become more complicated.

### 2.1.1 Body Perception

Another aspect of VR perception is the perception of ones own body. In a study by Banakou[BGS13], subjects were placed in either the body of a child or the body of an adult with the same height and asked to judge the size of virtual objects. The study showed a greater overestimation in those who controlled the body of a child, but that was no longer the case when the avatar no longer moved in sync with the subject. They believe this is because the lack of synchronous movement between the subject and avatar broke the illusion. So while one can change a user's perception through various virtual representations, it requires that the user identifies with their avatar.

This concept goes far beyond general perception, and crosses into a complete ownership. A study by Slater[Sla+09] explores a VR equivalent of the rubber hand illusion. The rubber hand illusion sees a fake prosthetic arm placed next to a person, while their real hand is occluded by a wall. The fake hand is placed in the same orientation, and both the real and fake hands are touched with similar objects at the same time on the same parts of the hand. When the visual connection is formed that the feeling they see is being reinforced by the touching sensation on their real hand it creates an ownership of the fake limb. After this, some touch sensations will only be applied to the fake hand. When done properly, subjects still report that they feel the presence on their real hand. In the VR study it was concluded that this ownership can be induced in a few ways. While the mirrored touch sensation with a virtual limb in VR works just as well as the the mirrored touch on a prosthetic in the real world, it was noted through Slater's second experiment that ownership was also formed when subjects saw virtual limbs move similarly to how their real body was moving. When this link was established, it was noted that subjects' physical bodies slightly moved when they saw the virtual limb make a motion.

Furthermore, this ownership can occur to varying extents and is hampered when obvious visual conflicts exist. A study by Schwind[Sch+18] used a variety of virtual hand and arm models in a tactile feedback experiment. What was found was that hands didn't need to represent the person realistically, but glaring inaccuracies lowered limb ownership. Things such as an arm being the wrong color of skin, or missing marks that the subject was familiar with were significant detractors, while more general representations that clearly weren't supposed to be a real human hand received stronger ownership. This being said, instances of large disparities such as a cartoon hand only having four fingers also had worse ownership levels. It is also worth mentioning that subjects played into the visual representation of the hand they were shown. Subjects with mechanical robot hands reportedly pressed on the device harder, and subjects with bulky gloved cartoon hands said they had a more difficult time feeling the bumps and indentations. However, work by Ogawa[ONH18] seems to show that avatar realism may play a bigger role in size perception. The study sees that as avatar realism increased to a more detailed hand the size estimation of a handheld cube became increasingly better.

### 2.1.2 Distance Perception

Distance is one of the most commonly studied spatial properties that people perceive in VR. This is broken down into egocentric(the distance an object is in relation to you) and exocentric(the distance between two foreign objects or points) perception. In the real world, people have a tendency to slightly underestimate distances. In virtual environments this same underestimation occurs, but to a greater extent. The more visual cues that are provided in the virtual environment, the closer the estimation gets to real world estimations.

This concept is further supported through the studies on depth perception in augmented reality(AR). In work by Jones et al.[Jon+08] it is shown that there is less underestimation in AR than VR. This is likely due to the abundance of correct environmental visual cues visible in the real world that provides more context. Current virtual reality headsets don't allow for natural field of view, and environments are often hampered by limitations of current real-time graphic engines. Since AR places virtual objects in the real world, neither of these are a problem. However, a study on how different design details affect depth perception of virtual objects in AR by Diaz[Dia+17] suggests that simplistic representations of virtual objects may currently provide the best results. The research shows that drop shadows had better results than raytraced cast shadows, and among various shaders placed on the virtual object Lambertian shaders(having no light specularity, or no interaction with light sources) often yielded the most accurate estimations. Diaz et al. claim that the design details that interact more with the environment may not be accurate enough to blend in with the real life scenery surrounding the object. This may cause discrepancies among the perception of virtual objects and real objects, creating an uncanny valley effect.

There are a few methods of having people estimate distances, the most common of which is blindfolded walking[WG02; IRA06; Kel+18; CSB15; RVH13], where subjects are told to visualize an area and are then blindfolded and told to walk to where they believe a certain object exists. This can be done in either an open-loop(subjects can visually see where they ended their walk) or closed-loop(subjects are given no visual information about the distance they walked) feedback system. This technique has been proven to have fairly accurate results, and is one of the most frequently used methods of having a subject estimate a distance. This method has also produced positive results when measuring the distance to virtual objects. Other approaches that have been taken include throwing bean bags[Sah+05] as an alternative tool, and triangle walking which is useful for longer distance estimations[Tho+04].

Since egocentric depth perception has been underestimated, there have been attempts to try and change properties of the environment to induce the correct estimations. One such example is a study by Leyrer et al.[Ley+15] where the eye height of the subject was manipulated in attempt to make objects seem closer. When a subject's eye height has been increased, it creates a smaller viewing angle on objects below the horizon. This in turn creates a false sense of size which then leads what would normally be an overestimation into a corrected estimation. This approach was successful, as the average participant produced more accurate results with a shifted eye height. However this technique has usability issues and isn't a perfect fix for a lot of problems, but in some circumstances is an easy way to resolve this problem. It is worth noting a similar example by Zhang[Zha+12] experimented with minimizing a space to produce a similar effect. Minimized spaces did see corrected results, however similar to Leyrer's study, this isn't a solution that works in all cases.

### 2.1.3 Size Perception

While the perception of size in virtual environments has been studied to a lesser extent compared to depth perception, various connections have been made that help us understand the field. Current work suggests an overestimation in size[Ken+08], which makes sense when considered as a counterpart to distance perception. Much work in regards to the size estimation field has been done in relation to the subject's own body proportions. These typically are comprised of a comparison between whether or not a subject could perform a given action in a real life setting and a virtual setting. This has been in seen in work by Geuss et al.[Geu+10] where people were asked if they could walk between two columns without turning their body. In this case a judgement about the size of a space was made in relation to the subject's shoulder size, something they are very familiar with. A second part of the study had subjects separate their hands to represent the distance between the columns. It was shown that there was no significant difference between a virtual environment and a real environment, even though the distance to the columns were underestimated to expected degrees in the virtual environment. This provides additional evidence that egocentric and exocentric distances are perceived differently in virtual environments.

Further work by Stefanucci examined used affordance judgements to compare size estimates of both solid objects and open spaces between a real setting and a virtual desktop display[Ste+12]. Affordance judgements estimations if a subject could complete a certain task given some information, in this case the judgements were could they completely grab the face of different sized cubes, and could they fit their extended hand through square shaped holes of different sizes. It was found that subjects underestimated the size of both the virtual hole and the virtual cube to much larger extents than what was seen without the desktop display. The difference between this study and the work done by Geuss is that these objects were significantly smaller and were at a close range, where the columns were put into a further space in relation to the subject.

While people have many tools to relate their body to width, a slightly different approach is taken with height. Height judgements are often correlated with eye height. Past studies have dictated that there are certain ranges in which people attempt to scale the height of an object as a relation to the horizon (at eye height level). Wraga et. al. [WP00] produced a series of experiments to attempt to map what these zones are. To do this, they had subjects estimate the size of objects(physical/virtual) at various distances with either straight lines on the floor, or lines that were slightly angled to force a false perspective. What was found is that eye height was used for both seating and standing when objects had a height that fell in a range of 20%-250% of the subject's eye height. Outside of this range and the eye height scaling would generally be discarded in favor of context information provided by the linear perspective lines on the floor.

Since much of this work deals with a relation of body size, it shouldn't be a surprise that the perception of a subject's own body can change their estimation of the size of an object. For example, the aforementioned study by Banakou[BGS13], saw that when people were shown to be children they overestimated the size of objects. The other part of this study that is interesting is that without synchronous movement, size perception had less deviation. This links back to the concept of body ownership and how that relates to these size estimations. The idea that without ownership body to object size estimations are worsened is supported in a psychology study by Linkenauger [LRP10] where subjects estimated the size of objects through either a normal, minimizing, or magnifying lens. It was shown that when the subject's dominate hand was in the frame with the object, estimations were still quite good. However, when the subject's non-dominate hand or someone else's hand was used instead, size estimations became increasingly worse. Even tongs were used instead of the subject's hand, if the subject was familiar with using the tongs they had better estimations. This helps demonstrate that the more comfortable a person is with a known object, the better they can use it to scale size, even if it's a foreign object.

# 3 Implementation

## 3.1 Virtual Reality Devices

In this study we opted to make use of the Oculus Rift and it's touch controllers. While there are limitations that come with the head set compared to other commercial devices, we feel this is negated by the benefits of the touch controllers when combined with the simple implementation of avatars.

### 3.1.1 Oculus Rift CV1

For this experiment, we used the Oculus Rift CV1 (seen in 3.1). The Oculus Rift is a fully immersive head mounted display(HMD), this means that the device is secured to the user's head with straps on the sides and blocks out visual stimuli from the outside world, allowing the user to freely move their head without breaking the immersion. The Oculus uses two sensors which are positioned on a table, a

Figure 3.1: Image of the Oculus Rift CV1 and the Oculus Touch Controllers

gyroscope, and accelerometer to find it's position and orientation. The headset has infra-red lights that are are on the face of the HMD. When these blink, the sensors pick up their locations and uses the known layout of the lights to correctly place the HMD in the virtual environment. The sensors are also calibrated with a floor position and eye height, so the current height of the HMD reflects the height of the user. The HMD also sports a pair of headphones which can be positioned over the wearer's ears to play music and audio.

The Rift also has some limitations. Unlike it's commercial counterpart, the HTC-Vive, the Oculus Rift has relatively little space inside the HMD. This can make it too tight for users with large heads or glasses. In addition, since the pair of sensors used by the Rift to track position are always front facing, turning your head beyond 90 degrees to your left or right can impare tracking. Since our study always has subjects facing in one position, this isn't a problem. Also, while there are headphones on the Rift, they do not cover the user's ears fully. This limits the amount of sound isolation which can lead to subjects hearing outside movements and talking. To help prevent this, we play white noise through the head phones which should partially mitigate this problem.

### 3.1.2 Oculus Touch Controllers

The controller of choice for this study is the Oculus Touch Controllers (seen in 3.1). The draw of these controllers is that instead of being a standard gaming controller, the controller is designed in such a way that, when held, your hand naturally makes a fist. This also has a series of triggers and face buttons that allow users to squeeze various parts of their hand to form actions, rather than pressing buttons. This makes actions such as grabbing and moving much more accurate to natural body mechanics, and is therefore more immersive. While this isn't ideal for pointing at objects or moving a character, in our study where subjects are stationary and predominately using the controllers to grab and move objects, this works perfectly. This allows subjects to select objects by "pinching" them, or simply squeezing together your thumb and index finger. Starting and finishing individual trials can be done by pressing the thumbsticks down. This action won't happen by accident since it requires that the subject changes the position of their hand on the controller.

Just like the HMD, the Touch controllers have infra-red lights on them that blink on and off. These are tracked by the sensors as well and uses a similar technique as the HMD. The limitations of this is that the controllers need to be in full view of the sensor. Moving your body too much might block the view of the controller from the sensor, making it stop in place or lag. While this might be problematic for more open areas, in our stationary study this isn't a concern.

### 3.1.3 Oculus Avatar Toolkit

When using Oculus devices in Unity, developers are given access to standardized toolkits and assets which may be useful. One of which allows for full-body avatars. While this is excellent for customization, it provides a bit more than we would like. Luckily, the toolkit also allows for just the hands of the avatar to appear. Since this works in conjunction with the Touch controllers, animations have been made that match the amount the buttons have been pressed on the controller to what position the virtual hand is in. Squeezing all the triggers makes a fist, squeezing just the lower trigger and lifting your thumb and index finger off the controller makes a "finger gun", etc. This works extremely well with virtual avatar ownership via synchronized movement. So long as the extents of accuracy aren't explored greatly, the "close enough" of the animations makes for a compelling experience. Furthermore, the hand avatar used is androgynous and purely white such that it doesn't match any skin tone. This works well with the aforementioned study on body perception[Sch+18] where the hand is clearly recognizable as a human hand, but doesn't try to be too close to a real hand. While the hand is obviously not a real hand, that takes away from the distraction of imperfections, this should theoretically allow for a better ownership across a generalized group of test subjects.

## 3.2 Pilot Study

A pilot study was run in which subjects judged the size of model cubes (24cm, 53cm, 84cm in size) at a distance and then changed a resizable box to be the same size. This study only ran minimal subjects and didn't have any major restrictions in terms of technique or movement. All models rested on the ground and were 3m, 6m, or 9m away from the subject. What was found is that objects were on average underestimated in size and had large variance in all directions and across all conditions. It was concluded that having the object on the ground was too far removed from the subject's eye height, and giving the subject the freedom to measure using any technique they wanted was too inconsistent to get reliable data. So objects should be propped up, and keeping eye height constant and feet stationary removes much of the varibility in measurement techniques.

## 3.3 PC Hardware

This experiment was run on an Alienware Aurora R5 with an Intel i7-6700 as a processor, 16GB of physical RAM, and a GeForce GTX 1070 for a video card.

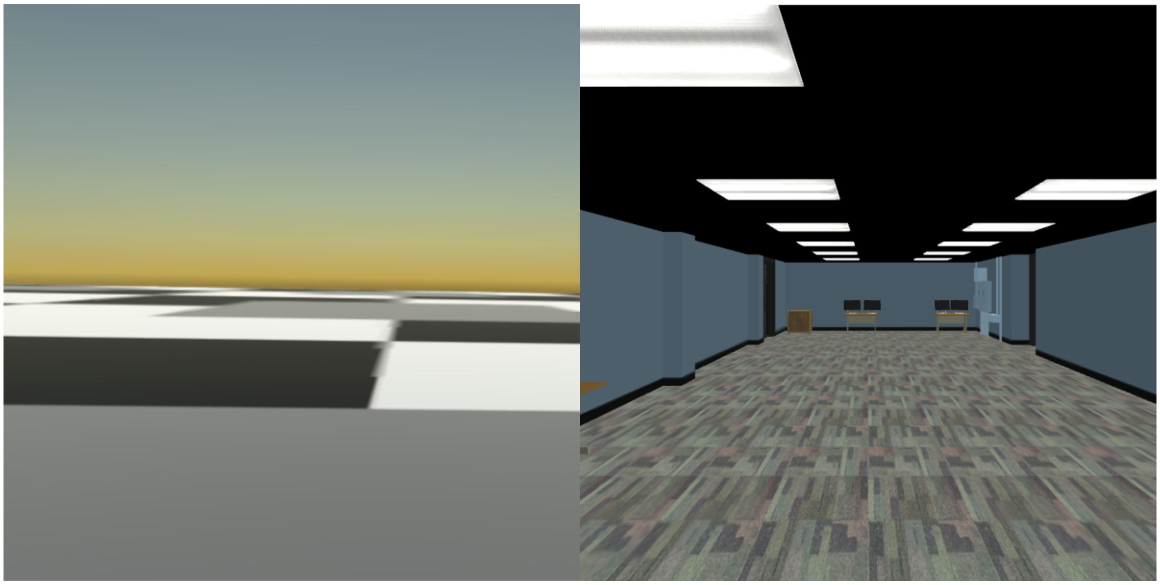

Figure 3.2: View of the two rooms unpopulated(abstract on left, SIVE Lab on right).

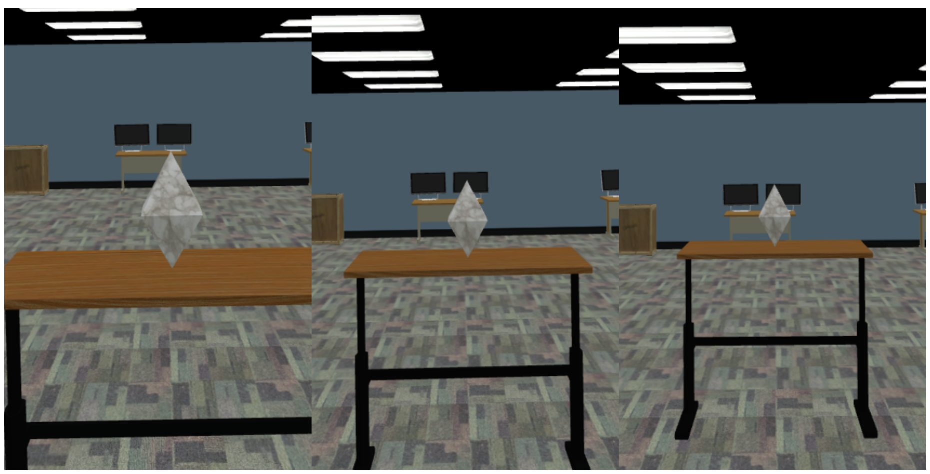

Figure 3.3: Reference for the different object distances. 2.1m on the left, 3.3m in the middle, 4.1m on the right.

## 3.4 Box Sizing Experiment

The primary experiment in this research has subjects in VR observing objects raised on a table, blanking their screen for several seconds, and then having them attempt resize a box so that it is the same size as that object. Subject's are put into abstract rooms to resize the box, and are in a scale model of the lab when they are

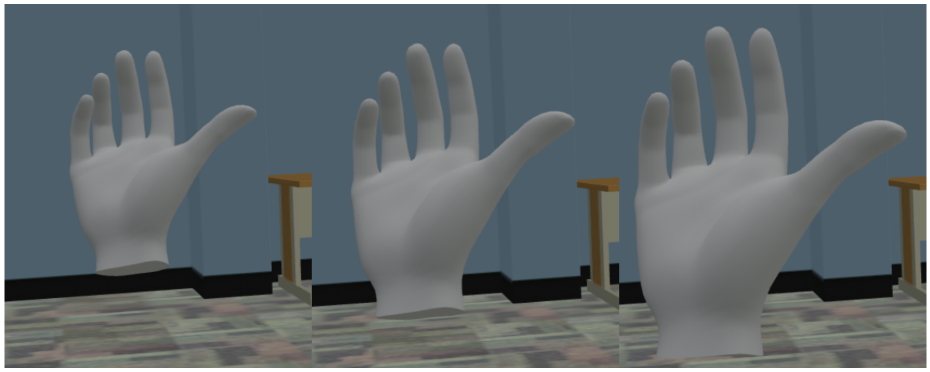

Figure 3.4: Reference for the different hand sizes. 80% on the left, 100% in the middle, 120% on the right.

shown the target object on a table (rooms shown in 3.2). This object is one of 6 different models(cola bottle, coffee mug, volleyball, wooden block, marble diamond, 5 orange spheres) and appears at one of three distances, 2.1m, 3.3m, and 4.1m, these are pictured in 3.3. Each variation is repeated 3 times for a total of 3x3x6 = 54 trials. The height of the table they are shown is a fixed height in relation to the subjects eye height, it is always 50cm below the measured eye height. This was chosen as the tallest object is 40cm, which means even the tip of the tallest object should be comfortably the horizon. As a between subjects variable, the size of the hand shown to the subject is in one of three conditions ( 80% in size, 100% in size, 120% in size) which can be seen in 3.4.

### 3.4.1 Object Selection

The objects chosen for this experiment are split between familiar and abstract objects. The familiar objects (coffee mug, cola bottle, volleyball) are objects the subject has most likely encountered in life and they likely have a good idea of how large the object should be in all dimensions. The abstract objects (marble diamond, wooden block, 5 conjoined spheres) are objects that the subject visually understands,

Figure 3.5: All objects measured by the subject. Marble Diamond is second row left-most, Orange Spheres is second row second to the left, Volleyball is second row second to the right, Cola is second row right-most, Wooden Block is first row left, Coffee Mug is first row right. (Note: this scene doesn't appear in the actual experiment)

| Name | Width(cm) | Height(cm) | Depth(cm) |
|---|---|---|---|
| $OrangeSpheres$ | 32.96 | 34.22 | 33.61 |
| $WoodenBlock$ | 21 | 21 | 21 |
| $MarbleDiamond$ | 20 | 40 | 20 |
| $ColaBottle$ | 10.9 | 29.5 | 10.9 |
| $CoffeeMug$ | 13 | 9.5 | 8.3 |
| $Volleyball$ | 22.9 | 22.9 | 22.9 |

Table 3.1: All the objects (first 3 abstract, last 3 familiar) and their size in each dimension

but has no way of knowing how large the object should be. All objects are pictured in figure 3.5. The dimensions of the objects are shown in Table 3.1.

### 3.4.2 Expectations

Previous work has covered size estimations in terms of width and height, but very little has been done on object depth. This is mostly because depth is hard to judge. without interacting with an object, it's hard to identify how far back it extends to a degree of relative accuracy. So while I expect width and height variations to be more minimal, I expect depth to be the least accurate and have the most variation due to the lack of visual cues and interaction.

Previous works have outlined that people estimate the size of objects well in relation to familiar body proportions, such as hand size, shoulder width, etc. In this study we manipulate the size of the subject's hands, which is the only visible body part in the virtual environment. When using the touch controllers, while the subject isn't completely performing the same actions as what they see, they are using the same muscles. This should help form an ownership connection as they are receiving visual feedback that confirms the familiar muscle movements are causing actions that they would expect. While the sizes of the hand are different, the modification to the hand size is +-20%. I wouldn't expect a size change of this much to have a significant impact on ownership. However, if an ownership connection is still formed regardless of hand size, previous work indicates they will use this hand to scale an object's size.

My hypothesis is that based on familiarity with the object, subjects will enter a different "mode" of perception. If the object is well known, like a coffee mug, the subject will mostly use their hand size when resizing the measuring cube as a reference for how big the object should be. For the more abstract objects, the hands will still play a role but it will be mostly relational to what was seen before. What this means that I expect the hand size modification to have a larger impact on the familiar set of objects than the abstract set. This means that with larger hands, I would expect

Figure 3.6: View of the user when sizing the measuring box

to see larger estimations, and with smaller hands, smaller estimations. For instance, if you judge a coffee mug to be about the height of your hand, and your hand is 20% larger, the estimation would also be 20% larger. I would assume the actual estimation wouldn't quite be one to one, and it would also be relative to how the object would be scaled with no hand modification. For example, if objects with no interference were perceived as being 90% of it's veridical size, you would expect the estimation of 20% larger hands to be 20% larger than the 90% estimate.

The last element of this study is that the objects appear at different distances. I would expect variation and accuracy to get worse as the object moves further from the subject in all three dimensions. Since distance and object size are inversely correlated (objects appear larger in your eye as they move closer to you), and distance is typically underestimated, I would expect sizes to be overestimated. This would be true if the only mis-perceived visual cue was distance.

### 3.4.3 Experiment Procedure

At the start of the experiment the subject is given a consent form and a pre-experiment questionnaire. This has the subject provide information such as affiliation with the university, experience with VR and gaming. After the subject has completed the questionnaire, they are asked to stand straight and look forward and a tape measure is used to measure their eye height. After this, the length of their hand, from the base of the palm to the tip of the index finger, is measured with a ruler. I defined the base as the palm as the upper crease that cross the wrist, directly under the tip of the index finger. Subjects were asked to lay their hand flat for this step. This does mean that only the length of the hand was considered and so subjects with particularly wide or narrow hands may not be represented well.

After measurements were taken, subjects were shown the Oculus Rift headset as well as a debrief on how to put it on comfortably and how to calibrate the lens spacing. After this the subjects were shown the Oculus Touch Controllers and were shown all the buttons and triggers. It was then explained that they would have use two primary actions, the "pinch" action and clicking down the right thumb stick. The pinch action is performed by pressing the top face button down with your thumb, and squeezing the index finger trigger on the same hand. This is used to select objects. Pressing down the right thumb stick moves the subject between environments. After the debriefing, subjects put on an eye mask for sanitation and the HMD which is set to the lens calibration screen. They are instructed to space the lens until the cross is clear, and readjust the hmd if the words are blurry. During this time, the subject's measurements and modification data is written into a text file.

Once the subject has properly spaced the lens, the screen is exited and a Unity program is started. At the start of the experiment, the subjects are placed in the

Abstract Room with no other objects. They are then given the hands and are told to get a feel for them by pressing the button and triggers to make different hand motions. If they do not clearly perform the pinch motion during this time, they are prompted to do so. After this the subject is asked if the performing the hand motions feels correct. Once they respond yes, they are told to click down the right thumbstick and the screen is blanked momentarily (randomly between 2 and 3 seconds), and they are loaded into the Sive Lab environment with a table and example model block placed on top of it. The subjects are then told to look around for a bit and get a feel for the room and their position in it. Once the subject has stopped looking around the room, they are told to focus on the model box on the table. They are told to try and visualize how big that object would be if it were right in front of them, and to think about the size in terms of the lab, their body and their hands. Once they have, they are told to press the right thumbstick down again.

Upon repressing down on the thumbstick, the screen is blanked for a random amount of time on the interval of 2 to 3 seconds. After this, the Abstract Room scene is loaded again but with a grey table with a resizing cube on it. The grey table appears at the same height as the table the model is placed on (50cm below the subject's eye height). The subjects are told how the sizing cube works, that moving your hand onto the cube or any of the knobs will cause them to glow. If you pinch a hand that is on the cube or a knob, you will grab it. Grabbing the box allows you to move it around, grabbing a knob allows you to resize the box in that direction. The subject are asked to try moving the box, grow it, and shrink it. They are then asked if they have a good feel for how the interactions feel. If they say they are comfortable with it, they are then asked to make this box the same size as the model they saw on the table. When they say they are good, they are asked to press the thumbstick again and repeat the process of visualizing the model in the lab, then resizing the cube to be the same size.

After each practice round (must do at least two) the subject is asked if they feel comfortable with the entire process of moving between the lab and abstract room, visualizing the model, and resizing the box. If they are they are told the final rules and details of the experiment specifically:

- You will be shown various objects in front of you, and you will be asked to visualize their size, and then make the resizable box the same size as the object they were shown.
- Not every object is a box, when you see an object that isn't a box, make it so that the object could perfectly fit inside the box, and it would touch all six faces.
- There will be repeating objects, don't try and make it the same size as the last time you saw it. Try to re-visualize the object every time you see it.
- You will estimate the size of about 50 objects. Don't be meticulous with the sizing, make the box the same size as your first gut instinct and don't worry about very small minuet changes.

Once the experiment is described, the subject is told to put the headphones over their ears to help drown out the background with white noise. This noise could be turned up or down if requested. The subject would then use the same controls that were used in practice to navigate through the 54 trials. The subjects were notified when they were half way done, and the objects were trials were done in a psuedo-random order. Every variation had to be shown before anything could repeat. And everything has to appear twice before it can appear a third time. This means hypothetically, it is possible for the same object to appear 6 times in a row, but that never happened. There were a couple instances where an object was shown three times in a row. During execution if subjects were disregarding the rules stated above (such as moving around too much, repeatedly making minuet changes, or making extremely inaccurate estimations) they were reminded of the task and how to preform it (standing still, general size and not super specific, visualize the object as if it were in front of you and try and make the box so the object would touch all 6 sides). No direction, or technique critique was given.

After the experiment was over, the controllers were taken from the subject and they were asked to remove the headset. Experiment statistics were saved to a text and backed up in a separate location on the computer as well as a google drive account associated with a university email address that requires two-factor authentication to access. After every experiment, subjects were asked to fill-out a post experiment questionnaire and the headset lens and controllers were cleaned with a vinegar and water solution.

The post experiment questionnaire uses a series of likert scales and short answers to evaluate how well the experiment was recieved by the subject. This includes questions of what objects were the most difficult to measure, how easily could they perform the experiment, if they felt any nausea or discomfort, if the equipment being used was distracting, and how well they felt the avatar represented them. The questionnaire in it's entirety is in the Appendix section of this thesis.

# 4 Results

Throughout the experiment, 43 subjects were ran. 3 subjects are not included in the data analysis for data collection failure, hardware failure, and health reasons. Out of the 40 subjects that did run, 2 were left handed, but this didn't generate any difference. Out of the six objects the 5 orange spheres was often misidentified as 4 orange spheres, the volley ball was called a soccer ball by several subjects, and a couple subjects referred to the 2 liters of cola as 1 liter of cola. The most common answer for the most difficult object was the 5 orange spheres. This is interesting as it seems to have the most difference based off of hand size. It's also worth mentioning that the other abstract objects have clear symmetry. The cola and marble diamond also had several votes for the most difficult object, and the wooden box had one.

Errors were measured in two different ways; the true error, which is simply the difference between the box sized by the subject and the model they estimated the size of (0 is equal, $> 0$ is overestimate, $< 0$ is underestimate), and percent error, the quotient of the box sized by the subject and the model they estimated the size of (1 is equal, $> 1$ is overestimate, $< 1$ is underestimate). We found that percent error was ultimately a more useful metric as it related better to the different sizes of objects.

Looking into the overall percent error in the different dimensions across hand sizes shown in 4.1. ANOVA shows no significance in hand modification values in X or Y but has Z approaching significance. It's worth mentioning a comparison between the hand modification and a post-experiment question filled out inquiring how well the subject thought the avatar's hand represented their actual hand in terms of size,

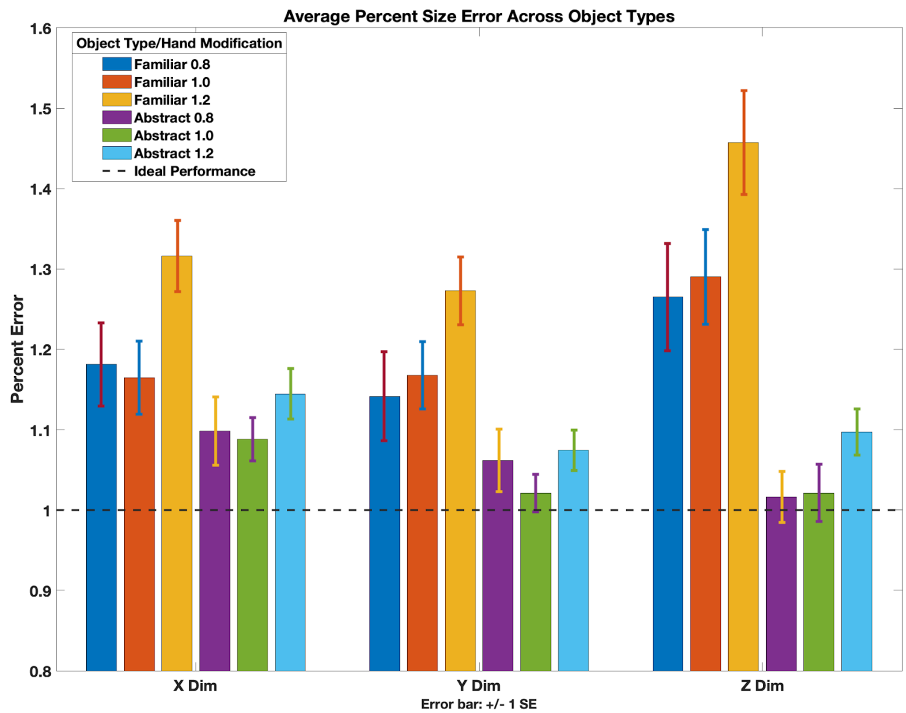


Figure 4.1: Percent error in width, height, and depth estimations across all hand modification conditions.

displayed in 4.2.

An ANOVA of this question with respect to hand size multiplier modifications indicates a significant effect ($F(2, 37) = 6.004$; $p < 0.005$). Post-hoc Bonferroni contrasts shows borderline statistical significance between the 80% and 100% multipliers ($p=0.056$) and stronger effect between 80% and 120% ($p<0.01$). No significance is reported between the 100% and 120% conditions.

## 4.1 Abstract and Familiar

Figures 4.3, 4.4, and 4.5 show the percentage error in their respective dimensions when looking at familiar and abstract objects. ANOVA shows a clear significance of $p < 0.001$ for object type, meaning a clear distinction in measurements of familiar

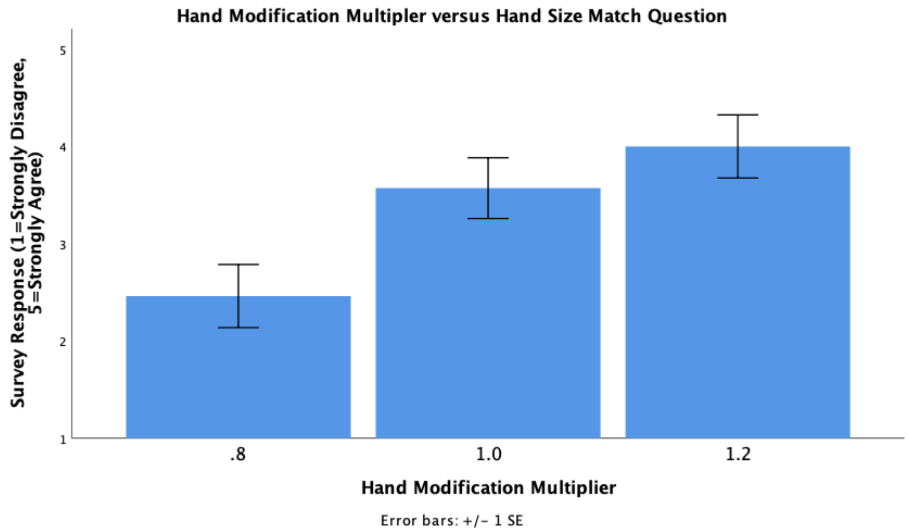


Figure 4.2: Post experiment question on how well subjects feel the avatar's hand represented the size of their own. Responses ranged from 1-5 where 1 was a poor representation, and 5 was a great representation.

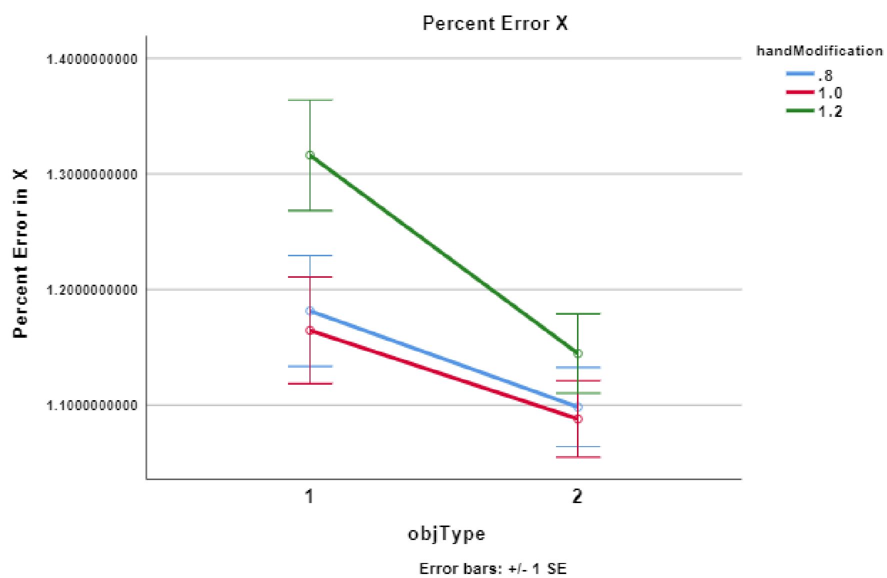


Figure 4.3: Percent error in width estimations when objects are familiar(1) or abstract(2) for all hand modification conditions

and abstract models. When looking at hand modification X object type, or the difference in how hand modifications effected the object types differently, overall an

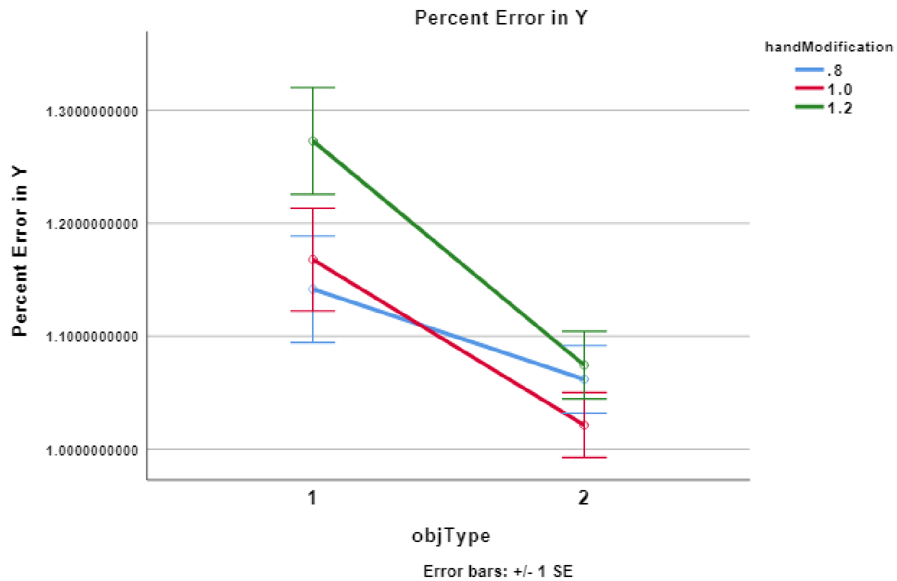


Figure 4.4: Percent error in height estimations when objects are familiar(1) or abstract(2) for all hand modification conditions

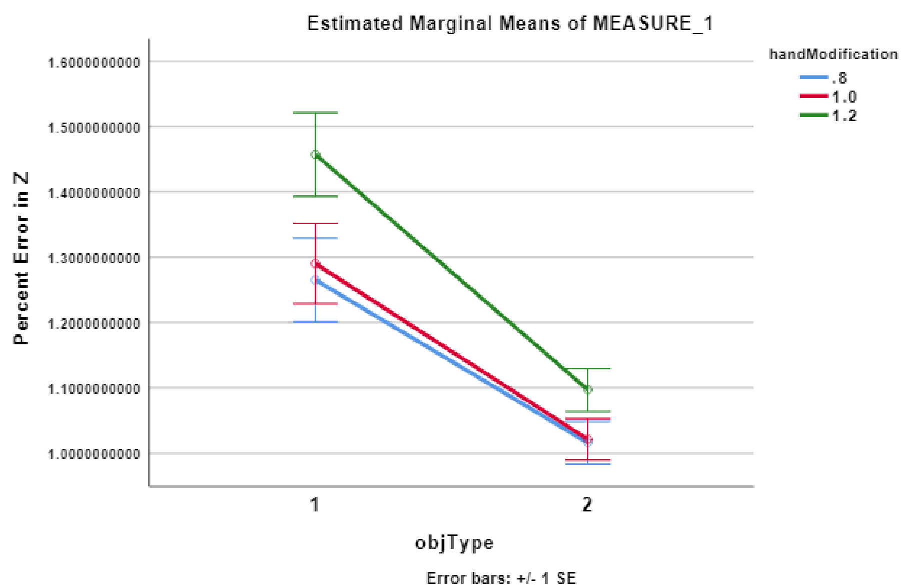


Figure 4.5: Percent error in depth estimations when objects are familiar(1) or abstract(2) for all hand modification conditions

ANOVA $F(2,37) = 2.961$ we find with $p = 0.064$ in the X dimension, $p = 0.036$ in the Y dimension, and $p > 0.168$ in the Z dimension. It's also worth remembering

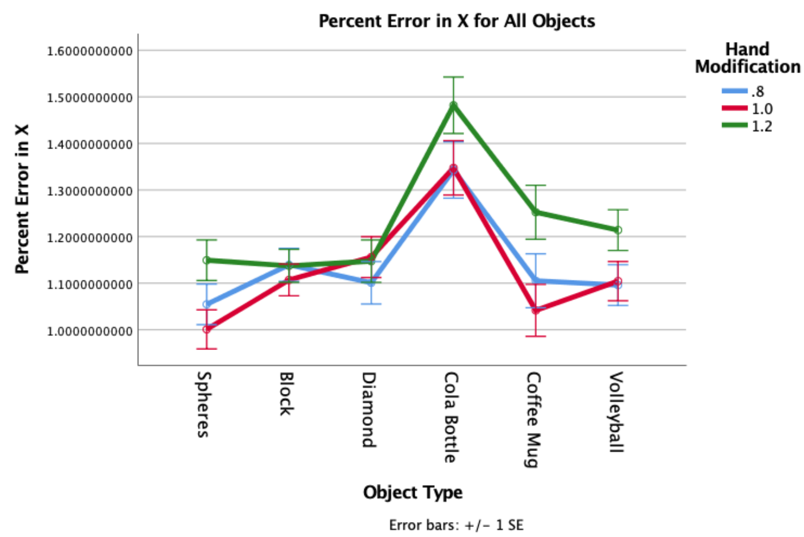


Figure 4.6: Percent error in width estimations for all objects across all hand modification conditions

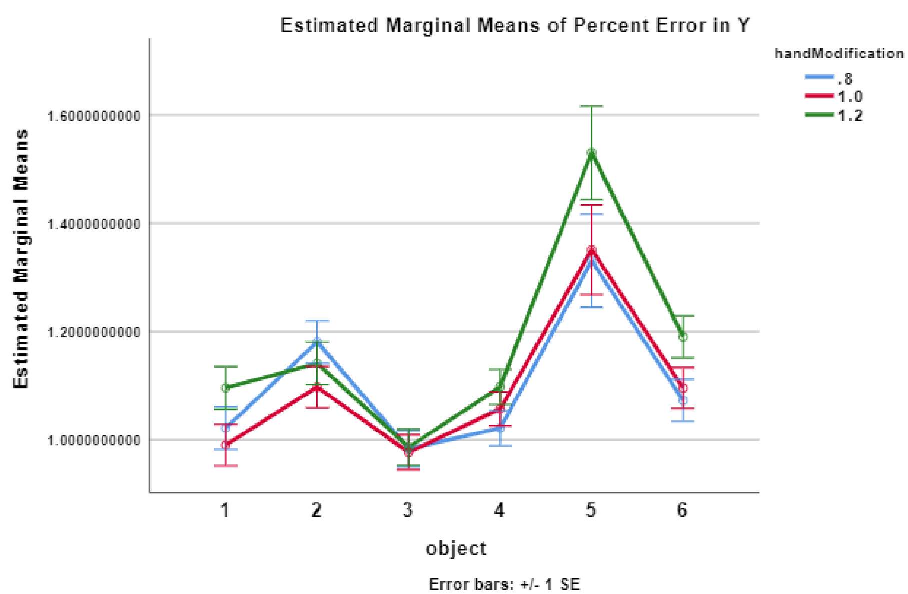


Figure 4.7: Percent error in height estimations for all objects across all hand modification conditions

the distinction between hand size likeliness in 0.8 against 1.0 and 1.2. The individual objects are shown in 4.6, 4.7, and 4.8 for a more detailed breakdown.

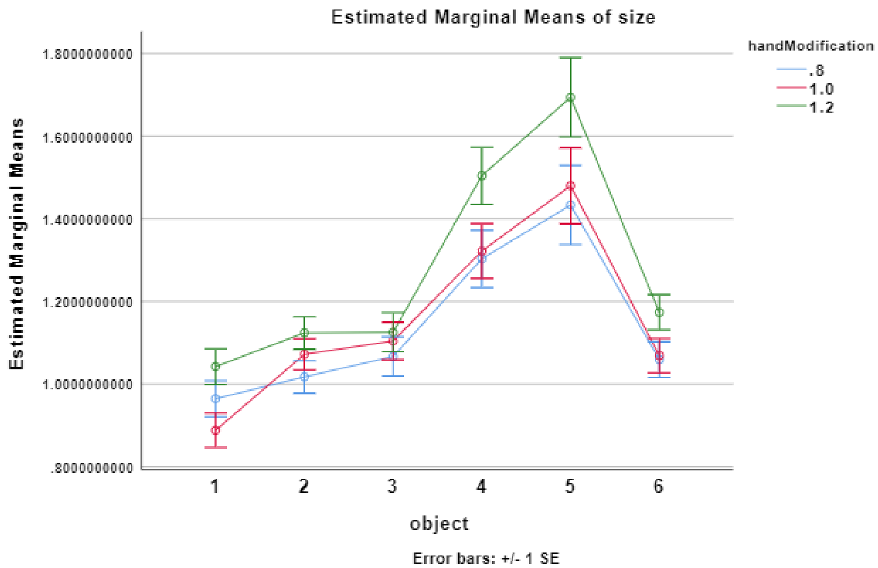


Figure 4.8: Percent error in depth estimations for all objects across all hand modification conditions

## 4.2 Distances

When looking at distances (results shown in 4.9), ANOVA shows a significance of $p < 0.001$, $F(1, 37) = 19.995$, which aligns well with previous work[Ken+08][BGS13]. The more interesting figure is the ANOVA of distance X object type, which shows a significance of $p < 0.05$ for all measures but Lower-bound, which is $p = 0.067$ partial Eta squared $= 0.136$. Showing differences in distance effecting objects with varying familiarity.

## 4.3 Timing and Repetitions

The differences in measurements across different amounts of time or subsequent repetitions is mostly unremarkable. These didn't have a major effect and there isn't anything particularly unexpected shown in the data.

Figure 4.10 shows that error didn't really change much in any dimension as the

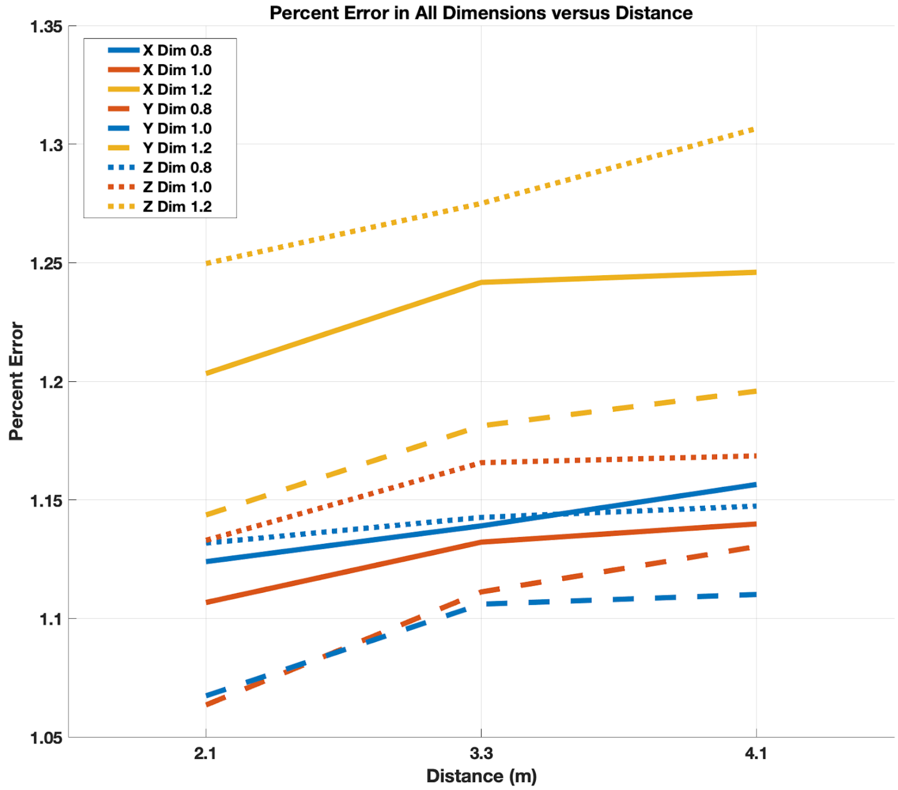


Figure 4.9: Percent error in width, height, and depth estimations when objects are 2.1m, 3.3m, or 4.1m away for all hand modification conditions

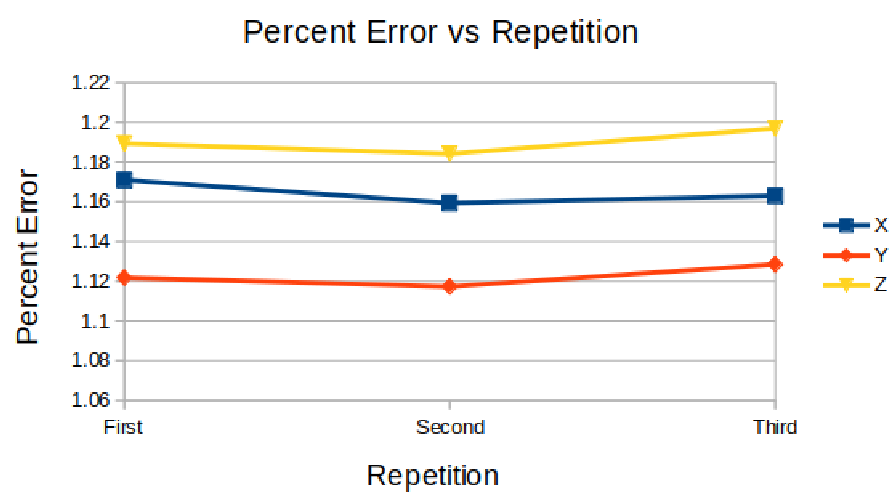


Figure 4.10: Average percent error in all dimensions across object x distance repetition

subject continued the trial. While this could be an indication that previous experience was used as a measuring tool, the same results would be expected if that wasn't the case. Ultimately meaning that the repetition doesn't appear to be an important factor.

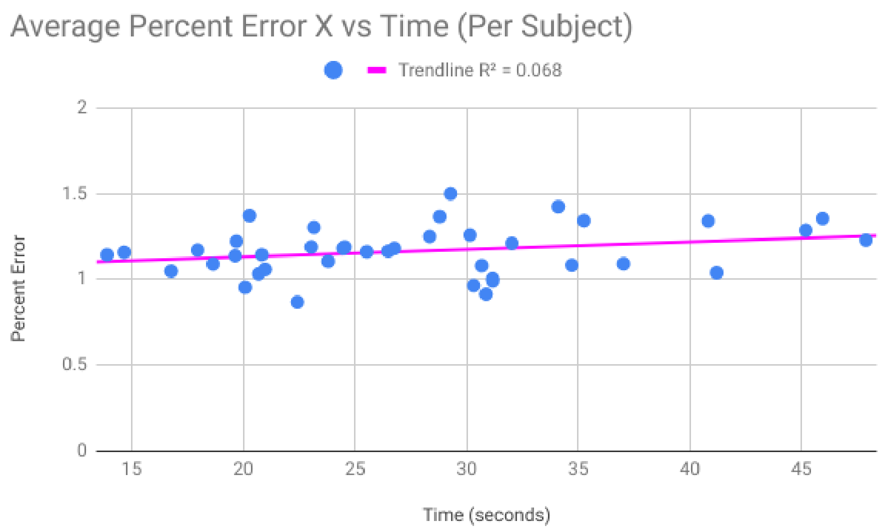


Figure 4.11: Each subject's average percent error in width compared to the time that was taken for them to size the object.

The above graphs (4.11, 4.12, 4.13) show the average percent error in each dimension against how long it took the subject on average to complete a trial. This data was also collected among all samples, and not just the averages of each individual subject, but those had even lower magnitude R-squared values and will be included in the appendix. This shows pretty clearly that subjects didn't necessarily preform any better or worse depending on how much time they took.

Lastly, 4.14 shows the average time taken on each object x distance repetition. As expected, later trials took less time on average as the subjects became familiar with the objects and controls used for the experiment.

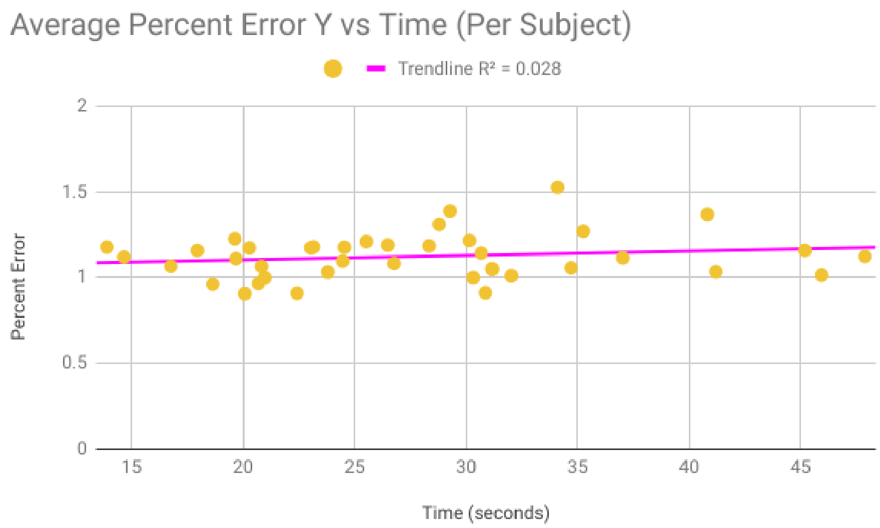


Figure 4.12: Each subject's average percent error in height compared to the time that was taken for them to size the object.

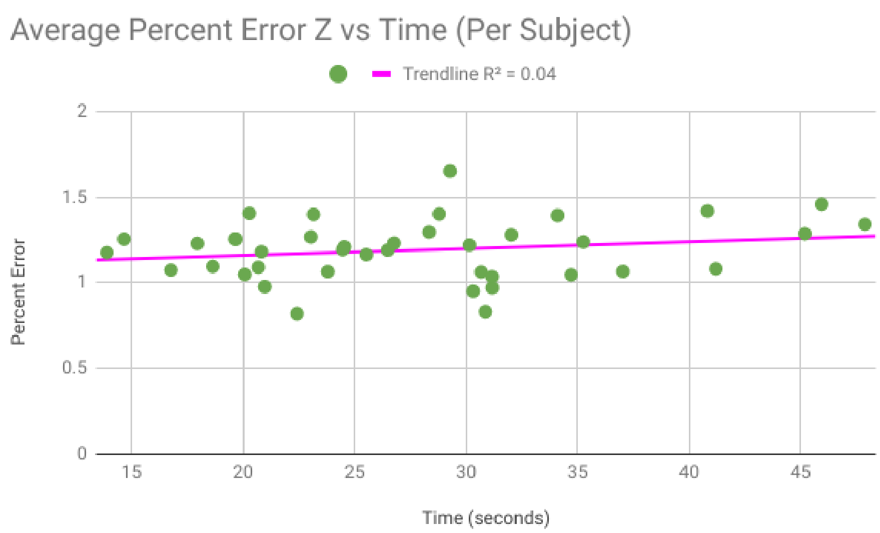


Figure 4.13: Each subject's average percent error in depth compared to the time that was taken for them to size the object.

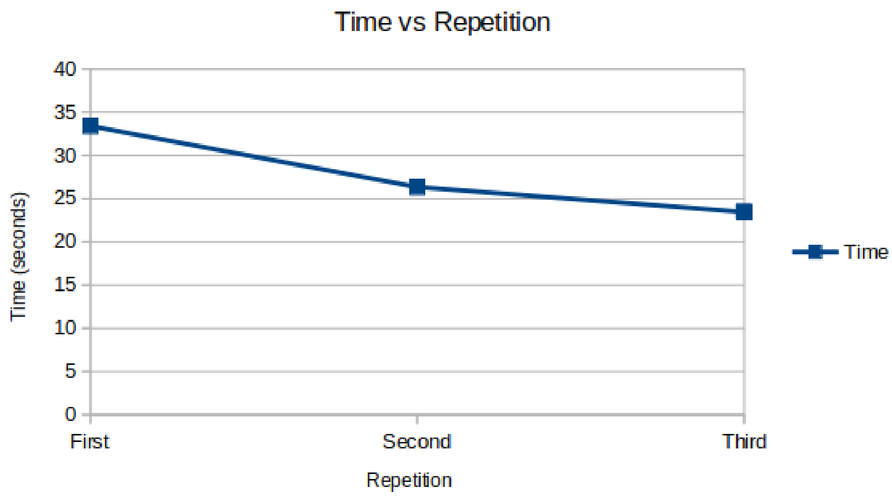


Figure 4.14: The average amount of time taken for each object x distance repetition

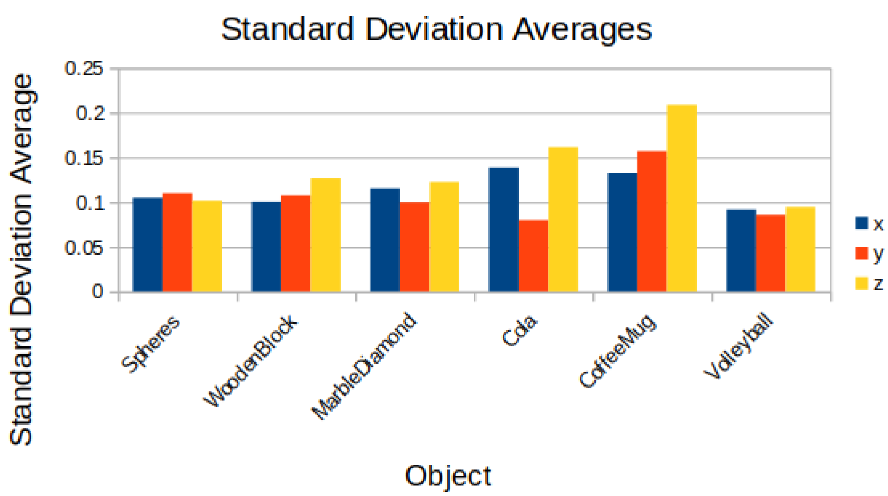


Figure 4.15: The average of standard deviations in percent error across all dimensions for each subject

## 4.4 Subject Range of Measurement

The graph in figure 4.15 shows how much subject's estimates varied on average. While figures 4.16, 4.17, and 4.18 show the variation across all measurements. These can also show potential outliers which are data points separated from the rest of the


Figure 4.16: Every percent error in terms of width across all trials. Object types are 1-Orange Spheres, 2-Wooden Box, 3-Marble Diamond, 4-Cola, 5-Coffee Mug, 6-Volleyball

data.

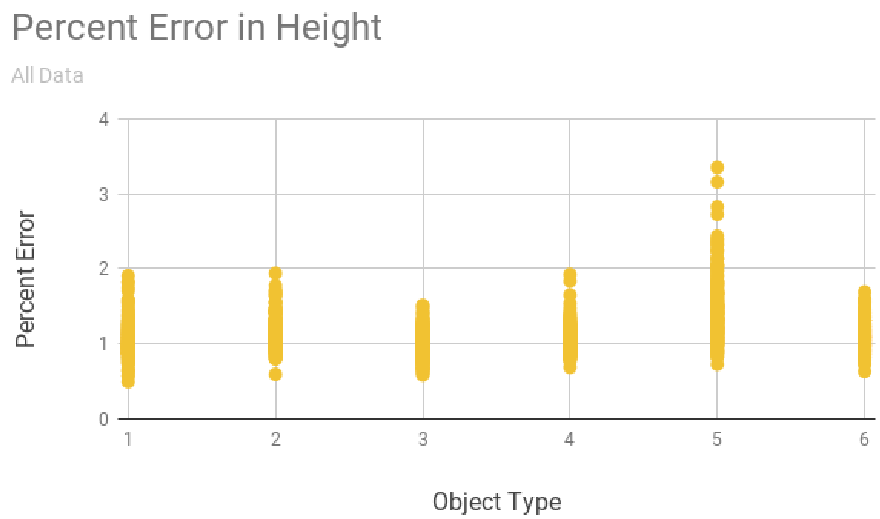


Figure 4.17: Every percent error in terms of height across all trials. Object types are 1-Orange Spheres, 2-Wooden Box, 3-Marble Diamond, 4-Cola, 5-Coffee Mug, 6-Volleyball

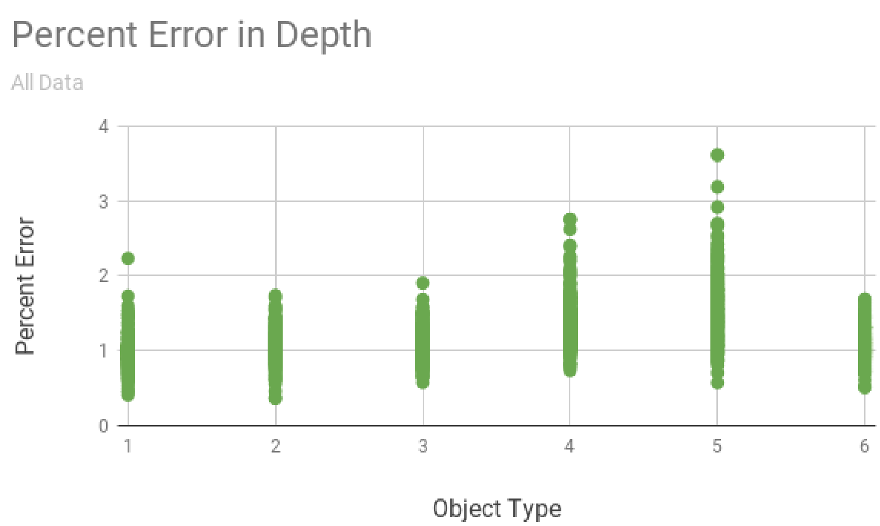


Figure 4.18: Every percent error in terms of depth across all trials. Object types are 1-Orange Spheres, 2-Wooden Box, 3-Marble Diamond, 4-Cola, 5-Coffee Mug, 6-Volleyball

# 5 Conclusions

## 5.1 Discussion

### 5.1.1 General Results

What the results appear to suggest is that the familiar objects were affected by the increased hand size, but were not affected in the same way by the decreased hand size. Although it's not possible to say exactly why, my theory is that people felt that the smaller hands didn't represent their hands well in terms of size 4.2 where 0.8 hand modification condition was the only category that skewed towards disagreement. However subjects in the 1.0 and 1.2 conditions had more positive feelings about how the hands represented their own.

This is interesting when compared to recent work by Lin Et Al.[Lin+19] where a similar process was taken looking into ownership and agency when hand sizes were changed. Our study found that people felt the largest set of hands represented them the most on average, where Lin's study received many comments from subjects who reported that the larger hands (25% increase) were too big and didn't fit them. This may have something to do with Lin's study having subjects use small, fit, and enlarged hands, as they had other hands to compare against.

While the 1.2 condition over-sized objects across the board when compared to 1.0 and 0.8 conditions, the difference was more significant in the case of the familiar objects. We believe this may be to subjects using their hands as the primary sizing tool when estimating the size of the box instead of what was perceived in the actual SIVE Lab environment. Similar to Wraga's[WP00] study where it was found people had different measuring techniques based off of field of view. If the avatar's hands were the primary measuring tool, the distance would play a lesser role as the environment wasn't used heavily in judgement.

This is also supported by how subjects sized boxes. It was common for subjects to make circular motions with their hands on the volleyball, perform the same actions as drinking out of the coffee mug, and simply picking up and placing down the cola. Where as they abstract set of objects seemed to be more body relative judgements rather than action oriented. Since the enlarged hands were still generally perceived to be the correct size of the subject's real hand, these interaction based guesses led to larger estimations. How objects were sized when the small hand was used isn't quite clear, but it is possible that the hands were underestimated in size which gave rise to overcompensation.

### 5.1.2 Dimension Discussion and Concerns

The first thing to address is correlation between the X and Z dimensions. Since 4 of the 6 objects share the same measurement across these two dimensions and it was hard to judge the depth given little context and limited viewing angles, it is likely that the depth estimation was a function of the width estimation. Watching subject's run the experiment, depth was often(but not always) the last dimension to be sized. Since the subjects can use the measure in one dimension as a relative sizing tool, it's hard to say how well subjects actually measured the depth of the object. The two objects that didn't have X, Z symmetry(the orange spheres and coffee mug) have wildly different depth measurements, where the spheres were underestimated and the coffee mug was greatly overestimated. However, the coffee mug was wider than it was deep, so if there was a strong cube like function to the size of the coffee mug that overestimation is explained.

When measuring height, there is a pretty clear correlation between the height of the object and the accuracy of the estimation. The marble diamond, which was the tallest object, had little to no difference across hand modifications and on average was a nearly exact height estimation (see 4.7). The cola bottle was also quite close and was the second tallest object. The volleyball and wooden block are close in height and have similar errors, and the coffee mug was the smallest and had the worst estimations. This is likely due to the top of the object's approaching the horizon, meaning there was much less guess work involved as the angle between the horizon and object is quite small. This being said, it's seems likely that the hands weren't used as a measuring tool as much as correlation to perceived eye height.

### 5.1.3 Object Discussion and Concerns

On a per object level, the most interesting points of data are the cola bottle's width and depth estimates along with the coffee mug's height and depth estimates. Interestingly enough, these are reflected in a subject's own variance as seen by 4.15. I don't have a clear answer on why the cola bottle's width was measured so inaccurately, but the results are consistent with Kenyon's work[Luo+07]. The depth is likely a product of the width estimation, and the over estimate matched previous studies[WP00], but there aren't many clues as to why. The thoughts on the coffee mug are most likely poor judgement in height due to the angle from the horizon, and depth might be a product of width even though the object wasn't the same in those terms. It's also worth mentioning that the coffee mug had the highest errors across any objects with some estimations being more than twice as large as the actual mug. It's likely a few very inaccurate measures (see figures 4.17 and 4.18) pulled up the measurement and increased the deviation, which makes the data somewhat unreliable.

Another interesting note is the errors in the orange spheres. Not only did the orange spheres have the most distinct separation of estimations by the hand modification condition 4.6, 4.7, 4.8, but in the 1.0 hand modification condition the average width measurement was almost exactly 1.0. The lack of clear structure and symmetry made this object more abstract than the other objects and it may show of different levels of abstract objects, which may be judged in different ways.

### 5.1.4 Abstract and Familiar Discussion and Concerns

The chief concern is were the objects actually familiar? Would a cola bottle be familiar to someone who doesn't drink soda? Or only drinks from cans? Is a volleyball a good object in the first place, it was mistaken for a soccerball several times. Also, subjects only reported what they thought the object was after the experiment. If the object was understood half-way through the process, it invalidates the "familiar" aspect until that point. Coffee mugs also vary in size, it might not be clear to some just exactly how big it is. A person's personal experience with coffee mugs might be a strong indicator on how large they estimated the object to be.

It appears clear by the data that the familiar objects were judged differently, but the surplus of outliars definitely pulls up the numbers. The volleyball in particular didn't really stray far from the abstract objects when looking at 4.6, 4.7, 4.8. The objects used in this study are slightly unreliable. There are a few hints of differnce when specifcally looking at the width measurement of the cola bottle and coffee mug, but it isn't exactly clear when combined with the concerns brought up previously.

## 5.2 Future Work

The next steps in this study will focus around exploring the results of the 80% hands and more differences of familiar and abstract objects. What needs to be done first and foremost is to identify what objects are the most familiar to the most people and objects that are definitely abstract. This should be done by modeling many objects (probably about 50) and having subjects say roughly how familiar they are with each one. Then using the most familiar objects for the "familiar" category in future experiments.

The 80% hands is more difficult to test for. This study will likely consist of four conditions, 100%, 90%, 80%, and 70%. This will help breakdown how hands are used as measuring tools and also help identify the limits where the avatar hands are too small for the person to still use as a measuring tool.

Since size estimation has rarely been explored in terms of object depth, a future study should specifically focus on this. It's likely that in this study depth estimations were influenced by being able to see the estimations made in terms of width and height. A similar approach to Stefanucci's[Ste+12] affordance judgements study could be taken to compare our current understanding of distance estimation to object depth estimation. This would however be looking at different sizes of objects than what was performed in this thesis.

Another route for further work would be a new blind-walking study where subjects walk to familiar and abstract objects. Since this showed a clear difference in size estimation, and perceived size is a factor of distance perception, it is possible that it could have significance.

Outside of those, repeat studies and running additional subjects through the current study will be done to help verify the results. With 3 conditions and only 40 subjects, there isn't enough to guarantee the accuracy of all results or confirm suspicions needed to entice other future studies.

# References

Banakou, Domna, Raphaela Groten, and Mel Slater (2013). "Illusory ownership of a virtual child body causes overestimation of object sizes and implicit attitude changes". In: *Proceedings of the National Academy of Sciences*. ISSN: 0027-8424. DOI: 10.1073/pnas.1306779110. eprint: https://www.pnas.org/content/early/2013/07/10/1306779110.full.pdf. URL: https://www.pnas.org/content/early/2013/07/10/1306779110 (cit. on pp. 5, 10, 29).

Creem-Regehr, Sarah H., Jeanine K. Stefanucci, and William B. Thompson (2015). "Chapter Six - Perceiving Absolute Scale in Virtual Environments: How Theory and Application Have Mutually Informed the Role of Body-Based Perception". In: ed. by BRIAN H. ROSS. Vol. 62. Psychology of Learning and Motivation. Academic Press, pp. 195–224. DOI: https://doi.org/10.1016/bs.plm.2014.09.006. URL: http://www.sciencedirect.com/science/article/pii/S0079742114000073 (cit. on p. 3).

Creem-Regehr, Sarah, Jeanine Stefanucci, and William B. Thompson (2015). "Perceiving Absolute Scale in Virtual Environments: How Theory and Application Have Mutually Informed the Role of Body-Based Perception". In: *Psychology of Learning and Motivation - Advances in Research and Theory* 62, pp. 195–224. DOI: 10.1016/bs.plm.2014.09.006 (cit. on p. 7).

Cutting, James E. and Peter M. Vishton (1995). "Chapter 3 - Perceiving Layout and Knowing Distances: The Integration, Relative Potency, and Contextual Use of Different Information about Depth*". In: *Perception of Space and Motion*. Ed. by William Epstein and Sheena Rogers. Handbook of Perception and Cognition. San Diego: Academic Press, pp. 69–117. ISBN: 978-0-12-240530-3. DOI: https://doi.org/10.1016/B978-012240530-3/50005-5. URL: http://www.sciencedirect.com/science/article/pii/B9780122405303500055 (cit. on p. 3).

Diaz, C., M. Walker, D. A. Szafir, and D. Szafir (2017). "Designing for Depth Perceptions in Augmented Reality". In: *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 111–122. DOI: 10.1109/ISMAR.2017.28 (cit. on p. 7).

Geuss, Michael, Jeanine Stefanucci, Sarah Creem-Regehr, and William B. Thompson (2010). "Can I Pass?: Using Affordances to Measure Perceived Size in Virtual Environments". In: *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*. APGV '10. Los Angeles, California: ACM, pp. 61–64. ISBN: 978-1-4503-0248-7. DOI: 10.1145/1836248.1836259. URL: http://doi.acm.org/10.1145/1836248.1836259 (cit. on p. 9).

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. DOI: 10.1109/VR.2006.52 (cit. on p. 7).

Jones, J. Adam, J. Edward Swan II, Gurjot Singh, Eric Kolstad, and Stephen 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*. APGV '08. Los Angeles, California: ACM, pp. 9–14. ISBN: 978-1-59593-981-4. DOI: 10.1145/1394281.1394283. URL: http://doi.acm.org/10.1145/1394281.1394283 (cit. on p. 7).

Kelly, Jonathan W., Lucia A. Cherep, Brenna Klesel, Zachary D. Siegel, and Seth George (2018). "Comparison of Two Methods for Improving Distance Perception in Virtual Reality". In: *ACM Trans. Appl. Percept.* 15.2, 11:1–11:11. ISSN: 1544-3558. DOI: 10.1145/3165285. URL: http://doi.acm.org.libpdb.d.umn.edu:2048/10.1145/3165285 (cit. on pp. 3, 7).

Kenyon, Robert V., Moses Phenany, Daniel Sandin, and Thomas Defanti (2008). "Accommodation and Size-Constancy of Virtual Objects". In: *Annals of Biomedical Engineering* 36.2, pp. 342–348. ISSN: 1573-9686. DOI: 10.1007/s10439-007-9414-7. URL: https://doi.org/10.1007/s10439-007-9414-7 (cit. on pp. 8, 29).

Leyrer, Markus, Sally A. Linkenauger, Heinrich H. Bülthoff, and Betty J. Mohler (2015). "Eye Height Manipulations: A Possible Solution to Reduce Underestimation of Egocentric Distances in Head-Mounted Displays". In: *ACM Trans. Appl. Percept.* 12.1, 1:1–1:23. ISSN: 1544-3558. DOI: 10.1145/2699254. URL: http://doi.acm.org/10.1145/2699254 (cit. on p. 8).

Lin, Lorraine, Aline Normoyle, Alexandra Adkins, Yu Sun, Andrew Robb, Yuting Ye, Massimiliano Di Luca, and Sophie Jfffdfffdrg (2019). "The Effect of Hand Size and Interaction Modality on the Virtual Hand Illusion". In: *IEEE Virtual Reality Conference (VR 2019)* (cit. on p. 36).

Linkenauger, Sally A., Veronica Ramenzoni, and Dennis R. Proffitt (2010). "Illusory Shrinkage and Growth". In: *Psychological Science* 21.9, 1318fffdfffdfffd1325. DOI: 10.1177/0956797610380700 (cit. on p. 10).

Luo, X., R. Kenyon, D. Kamper, D. Sandin, and T. DeFanti (2007). "The Effects of Scene Complexity, Stereovision, and Motion Parallax on Size Constancy in a Virtual Environment". In: *2007 IEEE Virtual Reality Conference*, pp. 59–66. DOI: 10.1109/VR.2007.352464 (cit. on pp. 4, 38).

Ogawa, Nami, Takuji Narumi, and Michitaka Hirose (2018). "Object Size Perception in Immersive Virtual Reality: Avatar Realism Affects the Way We Perceive". In: pp. 647–648. DOI: 10.1109/VR.2018.8446318 (cit. on p. 6).

Renner, Rebekka S., Boris M. Velichkovsky, and Jens R. Helmert (2013). "The Perception of Egocentric Distances in Virtual Environments - A Review". In: *ACM Comput. Surv.* 46.2, 23:1–23:40. ISSN: 0360-0300. DOI: 10.1145/2543581.2543590. URL: http://doi.acm.org/10.1145/2543581.2543590 (cit. on p. 7).

Sahm, Cynthia S., Sarah H. Creem-Regehr, William B. Thompson, and Peter Willemsen (2005). "Throwing Versus Walking As Indicators of Distance Perception in Similar Real and Virtual Environments". In: *ACM Trans. Appl. Percept.* 2.1, pp. 35–45. ISSN: 1544-3558. DOI: 10.1145/1048687.1048690. URL: http://doi.acm.org/10.1145/1048687.1048690 (cit. on p. 8).

Schwind, Valentin, Lorraine Lin, Massimiliano Di Luca, Sophie Jfffdfffdfffdfffdrg, and James Hillis (2018). "Touch with foreign hands: the effect of virtual hand appearance on visual-haptic integration". In: pp. 1–8. DOI: 10.1145/3225153.3225158 (cit. on pp. 6, 14).

Slater, Mel, Daniel Pfffdfffdfffdfffdrez Marcos, Henrik Ehrsson, and Maria Sanchez-Vives (2009). "Inducing illusory ownership of a virtual body". In: *Frontiers in Neuroscience* 3, p. 29. ISSN: 1662-453X. DOI: 10.3389/neuro.01.029.2009. URL: https://www.frontiersin.org/article/10.3389/neuro.01.029.2009 (cit. on p. 5).

Stefanucci, Jeanine K., David A. Lessard, Michael N. Geuss, Sarah H. Creem-Regehr, and William B. Thompson (2012). "Evaluating the Accuracy of Size Perception in Real and Virtual Environments". In: *Proceedings of the ACM Symposium on Applied Perception*. SAP '12. Los Angeles, California: ACM, pp. 79–82. ISBN: 978-

1-4503-1431-2. DOI: 10.1145/2338676.2338692. URL: http://doi.acm.org/10.1145/2338676.2338692 (cit. on pp. 9, 40).

Swan, J. E., A. Jones, E. Kolstad, M. A. Livingston, and H. S. Smallman (2007). "Egocentric depth judgments in optical, see-through augmented reality". In: *IEEE Transactions on Visualization and Computer Graphics* 13.3, pp. 429–442. ISSN: 1077-2626. DOI: 10.1109/TVCG.2007.1035 (cit. on p. 3).

Thompson, W. B., P. Willemsen, A. A. Gooch, S. H. Creem-Regehr, J. M. Loomis, and A. C. Beall (2004). "Does the Quality of the Computer Graphics Matter when Judging Distances in Visually Immersive Environments?" In: *Presence* 13.5, pp. 560–571. ISSN: 1054-7460. DOI: 10.1162/1054746042545292 (cit. on p. 8).

Willemsen, P. and A. A. Gooch (2002). "Perceived egocentric distances in real, image-based, and traditional virtual environments". In: *Proceedings IEEE Virtual Reality 2002*, pp. 275–276. DOI: 10.1109/VR.2002.996536 (cit. on p. 7).

Wraga, Maryjane and Dennis R Proffitt (2000). "Mapping the Zone of Eye-Height Utility for Seated and Standing Observers". In: *Perception* 29.11. PMID: 11219989, pp. 1361–1383. DOI: 10.1068/p2837. eprint: https://doi.org/10.1068/p2837. URL: https://doi.org/10.1068/p2837 (cit. on pp. 10, 37, 38).

Zhang, Ruimin, Anthony Nordman, James Walker, and Scott A. Kuhl (2012). "Minification Affects Verbal- and Action-based Distance Judgments Differently in Head-mounted Displays". In: *ACM Trans. Appl. Percept.* 9.3, 14:1–14:13. ISSN: 1544-3558. DOI: 10.1145/2325722.2325727. URL: http://doi.acm.org/10.1145/2325722.2325727 (cit. on p. 8).

# A Appendix A

This appendix contains the pre-experiment and post-experiment questionnaires.

# Pre-Experiment Questionnaire

Please state your relationship with the University:

- ❑ Student
- ❑ Staff
- ❑ Faculty
- ❑ Duluth Community Member
- ❑ Other

If you are a student, what is your major? ________________

Age __________ Gender _________

To what extent have you experienced motion sickness? ____________________

Approximately how many hours in the last month have you used virtual reality systems? ____________________

How much total experience do you have with virtual reality systems. Please select one option:

- ❑ 0-5 hours
- ❑ 5-30 hours
- ❑ 30-60 hours
- ❑ 60-100 hours
- ❑ 100+ hours

If you have had used virtual reality, please describe the types of experiences (such as games, simulations, movies or other immersive activities) that you have participated in including how you interacted with the virtual objects and environment.

Approximately how many hours each week do you play video games? ________________

If you have had experience with virtual reality, please indicate which virtual reality systems you regularly use:

- ❑ Oculus Rift with Hand Controllers
- ❑ Oculus Rift with Xbox Controller
- ❑ HTC VIVE with
- ❑ PlayStation VR
- ❑ Samsung Gear VR
- ❑ Google Cardboard
- ❑ Other _________________________________

Figure A.1:

**The following will be measured by an experimeter**

Standing Eye Height _________ Seated Eye Height _________

Hand Length (from base of palm to tip of index finger) _________

Figure A.2:

# Post-Experiment Task-Related Questions

**In the Oculus Rift virtual reality condition**

*I felt that I completed the task well*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*When I performed the task, it felt realistic*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*The size of my virtual hands matched my actual hand sizes:*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*Positioning my hands on the visual object was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*Performing the task was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*My virtual hands felt like my real hands:*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

Was the task easier standing or sitting?

Sitting    Standing    About the Same

Please give a brief description (1 - 4 words) of the 6 visual objects that were displayed to you:

1
2
3
4
5
6

Which object was the most difficult to perceive in the virtual reality environment?

Figure A.3:

**In the virtual reality condition in which you held the virtual reality controllers**

*The weight of the visual object that I was holding was important for me to complete the task*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*I completed the task well*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*When I performed the task, it was realistic*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*The objects I held in my hands felt similar to what I actually saw:*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*Positioning my hands on the visual object being held was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*Performing the task was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

**In the virtual reality condition in which you held a stick as a controller**

*The weight of the visual object that I was holding was important for me to complete the task*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*I completed the task well*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*When I performed the task, it was realistic*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*The objects I held in my hands felt similar to what I actually saw:*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|

Figure A.4:

| 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|

*Positioning my hands on the visual object being held was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*Performing the task was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

**In the physical condition in which you held the broom stick,**

*The weight of the visual object that I was holding was important for me to complete the task*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*I completed the task well*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*When I performed the task, it was realistic*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*The objects I held in my hands felt similar to what I actually saw:*

| Strongly Disagree | Disagree | Neutral | Agree | Strongly Agree |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*Positioning my hands on the visual object being held was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*Performing the task was:*

| Very Easy | Easy | Neutral | Difficult | Very Difficult |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

Figure A.5:

## Presence and Experience Questionnaire

*If one's level of presence in the real world is 100%, rate your level of presence in this virtual world.*

| | | | |
|---|---|---|---|
| 0% (none) | 2% - 50% | 50% - 75% | 75% - 100% |

*How realistic were the virtual world's reactions to your actions?*

| | | | | |
|---|---|---|---|---|
| Very Unrealistic | Unrealistic | Neutral | Realistic | Very Realistic |
| 1 | 2 | 3 | 4 | 5 |

*How responsive was the environment to actions that you initiated (or performed)?*

| | | | | |
|---|---|---|---|---|
| Very Unresponsive | Unresponsive | Neutral | Responsive | Very Responsive |
| 1 | 2 | 3 | 4 | 5 |

*How aware were you of events occurring in the real world around you?*

| | | | | |
|---|---|---|---|---|
| Very Unaware | Unaware | Neutral | Aware | Very Aware |
| 1 | 2 | 3 | 4 | 5 |

*How aware were you of your display and control devices?*

| | | | | |
|---|---|---|---|---|
| Very Unaware | Unaware | Neutral | Aware | Very Aware |
| 1 | 2 | 3 | 4 | 5 |

*How much did your experiences in the virtual environment seem consistent with your real-world experiences?*

| | | | | |
|---|---|---|---|---|
| Very Inconsistent | Inconsistent | Neutral | Consistent | Very Consistent |
| 1 | 2 | 3 | 4 | 5 |

*How well could you move or manipulate objects in the virtual environment?*

| | | | | |
|---|---|---|---|---|
| Very Difficult | Difficult | Neutral | Easily | Very Easily |
| 1 | 2 | 3 | 4 | 5 |

*To what degree did you feel confused or disoriented at the beginning of breaks or at the end of the experimental session?*

| | | | | |
|---|---|---|---|---|
| Extremely | Very | Moderately | Slightly | Not at All |
| 1 | 2 | 3 | 4 | 5 |

*How distracting was the control mechanism?*

| | | | | |
|---|---|---|---|---|
| Extremely | Very | Moderately | Slightly | Not at All |
| 1 | 2 | 3 | 4 | 5 |

Figure A.6:

*How proficient in moving and interacting with the virtual environment did you feel at the end of the experience?*

| Extremely | Very | Moderately | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*How much did the visual display quality interfere or distract you from performing assigned tasks or required activities?*

| Extremely | Very | Moderately | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*How well could you concentrate on the assigned tasks or required activities rather than on the mechanisms used to perform those tasks or activities?*

| Extremely | Very | Moderately | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*To what extent did you find the Virtual Environment fascinating?*

| Extremely | Very | Moderately | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*I felt nauseous*

| Extremely | Very | Moderately | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*My eyes felt strained*

| Extremely | Very | Moderately | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*I had a headache*

| Extreme | Strong | Moderate | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*I had problems concentrating*

| Extreme | Strong | Moderate | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

*I felt unpleasant*

| Extremely | Very | Moderately | Slightly | Not at All |
|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 |

Figure A.7:

# B Appendix B

This appendix contains additional graphs not included in the results sectrion.

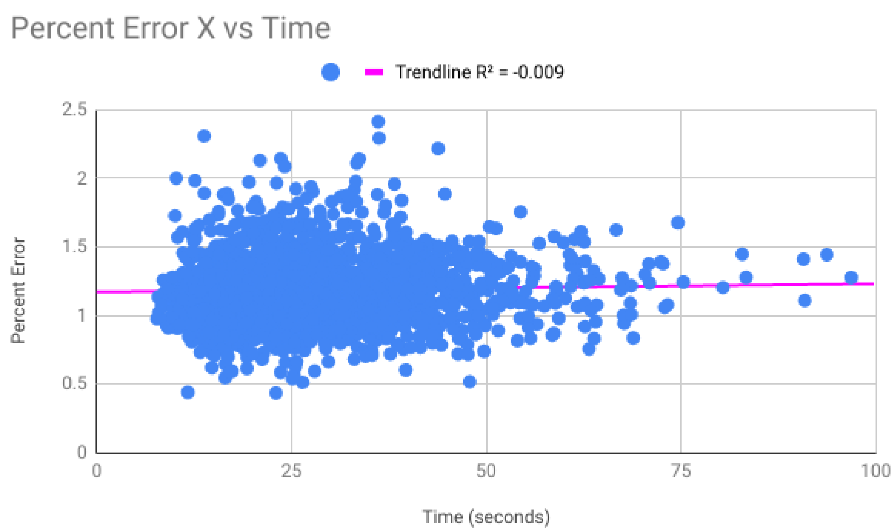


Figure B.1: This is every percent error in terms of width compared the time it took to size the object across all trials. Time is cut off at 100 seconds, there are a couple data points beyond 100.

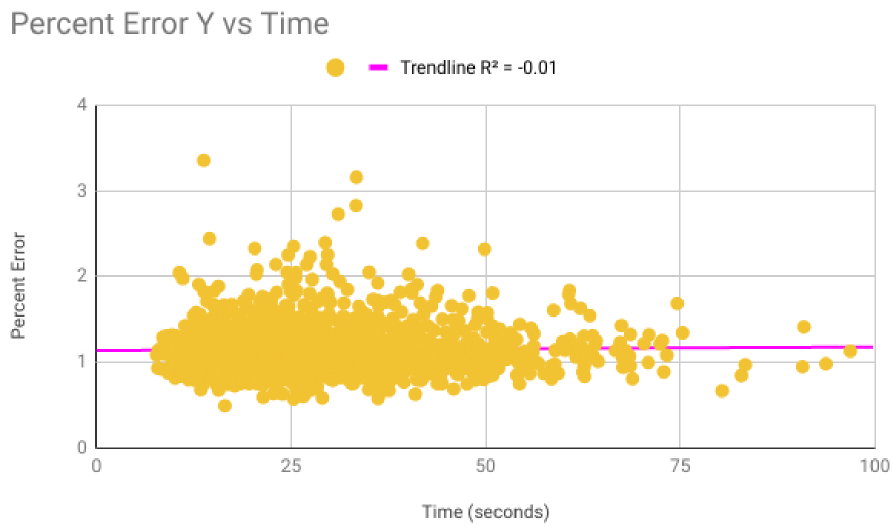


Figure B.2: This is every percent error in terms of height compared the time it took to size the object across all trials. Time is cut off at 100 seconds, there are a couple data points beyond 100.

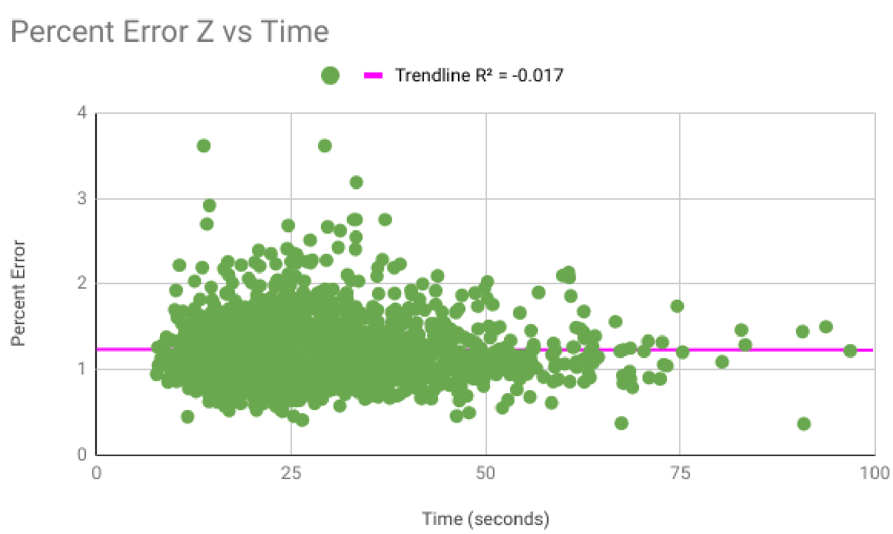


Figure B.3: This is every percent error in terms of depth compared the time it took to size the object across all trials. Time is cut off at 100 seconds, there are a couple data points beyond 100.