

Distance Perception and Cues to Distance in Virtual Reality

Ross Messing

July 26, 2004

Abstract

Distance perception in virtual reality is reviewed, and investigated in a series of experiments. In experiment 1, the systematic underestimation of distance found in virtual reality displays was investigated. This investigation used a video camera mounted on a head-mounted display to compare a photorealistic "virtual" world to monocular viewing of the real world with restricted field of view and unrestricted viewing. In experiment 2, the use of angular declination from the horizon as a cue to distance was investigated in virtual reality using a virtual world with a variable horizon line height. Subjects performed both a verbal estimation task and a visually directed action task. Subject data were fit to power function models to avoid issues of linear underestimation. In experiment 3, distance perception in indoor and outdoor virtual environments was investigated, using the same tasks as experiment 2. Experiment 1 found significant distance underestimation in the head-mounted display condition, but not in either of the real-world viewing conditions. This suggests that neither restricted field of view, nor level of graphical detail are responsible for the distance underestimation found in virtual reality. Experiment 2 found similar performance between the two tasks, and significantly higher power function exponents when the horizon line was lowered. This supports declination from the horizon as a cue to distance, and is consistent with a theory of distance perception suggesting the two tasks use the same underlying representation of distance. Experiment 3 found similar performance between both tasks in the outdoor environment, but a significantly larger mean verbal exponent than motor exponent in the indoor condition. This dissociation between verbal and motor performance suggests that there might be some fundamental difference between the two tasks, or between the two environments used.

1 Introduction

Understanding distance perception has long been a fundamental problem in the study of human vision. Despite this, distance perception has not yet been fully elucidated, and remains an area of active research.

1.1 How does distance perception work?

In order to investigate distance perception, we have to consider what mechanism underlies it. Because humans can interact with environments that they have seen, even with their eyes closed, we can assume that they have some internal representation of the space around them. Traditionally, this internal representation has been called visual space, and a lot of perception research goes into investigating its properties (J. Loomis, Silva, Fujita, & Fukusima, 1992).

There have been a number of approaches to investigating the structure of visual space. Gibson believed that visual perception was veridical and direct, implying that there is no need for a visual space, since everything visible in a natural, structured environment is directly and immediately available (Gibson, 1986). Gilinsky also investigated the relationship between the real world and our representation of it, seeking mathematical equations to map from physical to visual space (Gilinsky, 1951).

Loomis and his colleagues have done a number of experiments involving visually directed action, meaning that subjects are given visual input, and subsequently asked to perform a task without vision (J. Loomis et al., 1992). This sort of task clearly involves visual space. One theory put forth by Loomis et al. suggested that physical space is mapped to visual space in a way that gives very accurate measures of egocentric distance, but not of exocentric distance. Loomis et al. use this to argue that visual perception and visually directed action draw on the same perceptual resources (J. Loomis et al., 1992).

Another theory, which initially seems to contradict Loomis et al.'s view of a unified visual perception and visual-motor system, comes from Goodale and Milner, who show a range of neurological and behavioral evidence for a separation between the neural pathways leading to conscious visual perception, and those leading to visually guided action (Goodale & Milner, 1992). This theory is at odds with Loomis et al.'s because its suggestion of separate visual and visuo-motor pathways implies some dissociation of visual space for the two. Michaels, an ecological psychologist, attacks Goodale and Milner's definitions of perception and perceptual action, pointing out that the assumption of a purely perceptual information processing stream goes against the tight perception-action coupling advocated by ecological psychologists (Michaels, 2000). Despite this objection, most mainstream perception scientists find Goodale and Milner's assumption of perceptual representation unobjectionable.

While these two theories seem to be in conflict, neither one really excludes the other. Goodale and Milner note that while the two pathways can operate independently, there is definitely some communication between them (Goodale & Milner, 1992). This would allow Loomis et al.'s theorized single, egocentric representation of visual space to reside in one place, and be available to both pathways.

1.2 How is distance perception measured?

One major focus of work by Loomis and his colleagues is the difference between egocentric and exocentric distance perception. Egocentric distance is distance from the observer to some location in the world, while exocentric distance is the distance between two objects. They argue that visual space is organized in terms of egocentric distance, and show evidence of near-veridical egocentric distance perception, and significantly less accurate exocentric distance perception (J. Loomis et al., 1992).

A related concern which Loomis et al. have controlled for, but haven't discussed in detail, is the indicator of perceived distance. Distance can be indicated by subjects either implicitly, or explicitly. Implicit indication of perceived distance involves an action, like walking to a target, which indicates indirectly what the subject believes a distance to be. Explicit indication of perceived distance tends to be more cognitive, and involves a direct report of what the subject believes the distance to be.

Another concern is how to model the data once it's been gathered. One well-established method of modelling this sort of distance perception data is with a power function. The use of a power function generally focuses on the exponent, which indicates the rate at which perceived distance appears to increase as actual distance increases. A pair of studies by Teghtsoonian and Teghtsoonian illustrate this method of modelling distance perception data (M. Teghtsoonian & Teghtsoonian, 1969; R. Teghtsoonian & Teghtsoonian, 1970). In the first of these studies, verbal judgements of indoor distances were given by subjects, and each subject's responses were modelled as a power function. In the second study, the mean exponent found in the original study was compared to the mean exponent found for outdoor viewing. The indoor study found a mean exponent of 1.2, meaning that perceived distance increases more quickly than actual distance indoors. The outdoor study found mean exponents between 0.85 and 0.99, decreasing as the range of distances judged increases, which means that perceived distance and actual distance increase at about the same rate up to a point, after which increased actual distance has a diminishing impact on perceived distance. The popularity of power function modelling has declined in recent years, in favor of simpler, linear models of the relationship between perceived and actual distance, as used by Loomis et al. (J. Loomis et al., 1992). Linear models more directly model absolute distance perception, while power function models are more useful for understanding how that distance perception changes over distance.

1.3 Cues to Distance

Another important consideration is just what cues to depth are available to the human visual system.

There are two kinds of information that cues to distance can give. Most cues give relative distance information, telling the visual system about the distance relationships between different objects in the world. For example, a relative

distance cue might tell the visual system that one object is twice as far away as another object, but not exactly how far away either object is. A few cues give absolute distance information, telling the visual system about exact distance quantities in the world. These absolute cues are often limited, either by effective range, or familiarity with the observed object. The ways in which absolute and relative cues to distance perception can be combined will be discussed shortly.

In order to focus the light that falls on the fovea, the eye uses a lens system. This system involves using ocular muscles to vary the thickness of the lens. The visual system has access to how thick or thin these muscles make the lens, where the thinner the lens is, the farther away the eye is focused. This cue, called accommodation, tells the visual system exactly how far away the object being brought into focus is. Accommodation only provides useful information for objects within about 8 feet of the observer, because at greater distances, the lens is already as thin as it gets, or needs to be to bring distant objects into focus (Palmer, 1999; Cutting & Vishton, 1995).

Some distance cues are given by motion. One such cue, motion parallax, allows an observer to determine the relative distances of two objects based on their relative motion as the observer's viewpoint moves through the world. Objects that are closer to the observer will move more than objects farther from the observer (Palmer, 1999; Cutting & Vishton, 1995).

Another motion cue comes from the motion of texture as the observer moves through the world. This cue is given by the way textures change over time, as either the observer or the world moves. As textures shrink, cover less area, and increase in density, the objects that they're associated with grow more distant (Palmer, 1999).

Texture can also be a depth cue without motion. Since an object's texture density tends to be greater on parts of the object that are further away, the texture gradient can function as a cue for relative depth (Palmer, 1999). Further supporting texture as a depth cue, texture discontinuities have been shown to cause underperception of distance (Sinai, Ooi, & He, 1998). Ferial et al. have hypothesized that this is because the line separating the textures is given some amount of area, and that this separating area is subtracted from the distance estimation (Ferial, Braunstein, & Andersen, 2003).

Another property of the projection from three-dimensional to two-dimensional space is the existence of a horizon line, at eye level, infinitely far away, towards which all parallel lines converge. Basic trigonometry allows the visual system to determine the distance from the observer to an object on the ground plane, as long as the angle between the object and the horizon, and the observer's height, are known (Palmer, 1999).

Ooi et al. recently demonstrated that the angular declination from the horizon line to an object is a cue to distance (Ooi, Wu, & He, 2001). Objects appear closer to an observer when the angle between the object and the horizon is greater. This is because the horizon line is generally constant, as is the observed angle between the horizon and an observer's feet. As a result, when the angle between the horizon and an object increases, this implicitly decreases the angle between the object and the observer's feet, implying that the object is

closer.

Dixon et al. have shown that eye height is a cue to object size (Dixon, Wraga, Proffitt, & G. Williams, 2000). Using the fact that the ratio of an object's height to its distance from the horizon is equal to the ratio of its retinal size to its retinal distance from the horizon, an object's actual size could be extracted from an image of it, as long as eye height is known. This could also provide a cue to distance, as distance can be determined if both observed and actual height are known.

A very simple cue is based on the observed sizes of objects. This cue, relative size, is the visual system's understanding that larger objects tend to be closer to the observer than smaller objects (Palmer, 1999). Another distance cue, familiar size, works for objects of known size. If an object whose size is familiar to the observer is seen at some distance, the observer can use the difference between the object's known size and its perceived size to appropriately scale the perceived distance into actual distance (Palmer, 1999). Familiar size is the only cue to distance that gives absolute, rather than relative, distance estimations without relying on motor signals from the ocular muscles.

Another cue to distance that takes advantage of the structure of the environment is aerial perspective. Aerial perspective refers to the haziness of distant objects caused by photon-blocking particles in the air between the observer and the object (Palmer, 1999; Cutting & Vishton, 1995). This causes objects that are further away in hazy environments to appear to have lower contrast than nearer objects, as less of the light reflecting off of further objects reaches the observer.

The addition of a second eye to the visual system allows two more cues to distance, one ocular, one optical. As discussed before, the visual system has access to information about the ocular muscles. As a result, the different viewing angles of the two eyes, are available to the visual system. Knowing the two viewing angles, and the interpupillary distance, it is possible to determine the point of intersection of those angles, relative to the observer. This cue is called convergence, and while it is only useful within about 8 meters of an observer, after which the angles of the eyes are too close to straight ahead to be useful, it is a very powerful cue to depth within that distance. Like accommodation and familiar size, convergence gives an absolute estimation of distance (Palmer, 1999).

The optical cue given by a second eye is the difference between corresponding points in the two optical images. This difference, called binocular disparity, is a cue that something is close by, as the retinal images of the two eyes differ less and less as the viewed distance increases (Palmer, 1999).

While many cues to distance are well-established, there may well still be cues that we have not taken into account. One method used to investigate this is ideal observer modelling. Ideal observer models attempt to quantify the most accurate possible perceptions given the known physical limits of perceptual cues and optimal Bayesian combination of those cues. Ideal observer models show that we aren't considering the right cues, or misunderstand their limits, when the models perform worse than actual human subjects (Liu, Knill, & Kersten,

1995). While we cannot present an ideal observer model here, we can address how humans actually perform cue combination.

1.4 Cue Combination

While it's important to individually consider the cues to distance, it is also important to consider their combination. Without some way to combine the information from all of these cues, only accommodation, convergence, and familiar size could give estimates of absolute distance. Since the first two cues are only effective within a few meters of the observer, and familiar size is only effective at estimating the size of known objects, using only these cues would be a significant limitation.

One way to combine cues would be a winner-take-all competition, using the distance value reported by the cue that the visual system gives the most weight to, or the one that reports the least uncertainty. This approach is problematic because individual cues are unreliable. The winner-take-all approach also throws out the very powerful technique of using one cue to support another.

Landy et al. describe two ways that cues could be combined to support each other, and offer a model that compromises between the two (Landy, Maloney, Johnston, & Young, 1995). One way, weak fusion, is a weighted linear combination of the values of various cues. This model is appealing because it makes empirical investigation of the weights easy. However, weak fusion is too limited, because not all cues give absolute distance values, and thus couldn't be combined as a weighted linear average of depth maps.

Another model described by Landy et al., strong fusion, involves very complex interactions between different cues (Landy et al., 1995). While this model is very flexible, it effectively eliminates the ability to investigate cues individually. Because of this limitation, strong fusion is seen as an overly flexible, pessimistic model.

A compromise offered by Landy et al., modified weak fusion, involves the use of some cues to "promote" others (Landy et al., 1995). Most cues offer relative information, that can be transformed into absolute information given a single point of absolute distance information from another cue. Other, ancillary cues offer support information that can qualitatively disambiguate cues, providing information like the direction in which an object is spinning. The depth maps generated by combining cues are then combined in a weighted linear fashion.

Hillis et al. show that multimodal cue combination allows access to the original cues, while the combination of cues from the same sensory modality merges them and doesn't allow access to the original cues (Hillis, Ernst, Banks, & Landy, 2002). Hillis et al. demonstrated this using a discrimination task, where three stimuli, two of which were identical, were given, and the task was to pick the odd stimulus. Subjects were shown multi-cue stimuli that appeared to have the same value for a parameter (either slant or size). Some of these stimuli had consistent values for different cues, while others had inconsistent values whose weighted average was the intended value, even though no cue indicated that value by itself. Subjects using a combination of visual and haptic cues were

able to discriminate accurately between individual cue values, while subjects using only visual cues were unable to separate the individual cues from their weighted average.

1.5 Distance perception in virtual reality

Virtual reality is a popular and powerful tool in perception research. It allows experimenters unprecedented control of visual stimuli, while maintaining some measure of ecological validity(J. M. Loomis, Blascovich, & Beall, 1999).

Virtual reality systems have a number of limitations. Because their screens are fixed, they cannot manipulate accommodation. The quality of real-time computer graphics is limited, as is the resolution of virtual reality displays. The field of view available from most head-mounted virtual reality displays is much more limited than a natural field of view. The light from a head-mounted display is collimated, which means that subjects will get the same image on their retina as the position of the display on their head shifts, potentially causing problems if the display is not very securely attached to the subject. Despite these limitations, virtual reality systems are often the best way for experimenters to display complicated, highly-controlled stimuli in something approaching an ecologically valid setting.

One issue that experimenters have run into is a systematic compression of perceived distance in head-mounted displays. Loomis and Knapp thought that this compression might be due to the limited quality of their virtual reality system's graphics rendering(J. M. Loomis & Knapp, 2003). Rowland has hypothesized that using dense, high quality, natural textures on the ground plane would reduce compression(Rowland, 1999). Surprisingly, no significant difference was found in distance compression between very simple line-drawing virtual environments, and photo-realistic ones(Thompson et al., in press). This finding seems at odds with both Loomis & Knapp's, and Rowland's explanations. In addition, Creem-Regehr et al. found that, individually, neither monocular viewing, nor restricted field of view caused underestimation of distance in real environments, which suggests that they are not responsible for compression in virtual environments(Creem-Regehr, Willemsen, Gooch, & Thompson, 2003). Knapp and Loomis also investigated restricted field of view in binocular viewing, and like Creem-Regehr et al., found no statistically significant effect of restricting field of view (Knapp & Loomis, (in press)).

The compression of distance in virtual reality is a confusing phenomenon that has resisted investigation. No variable has been found that affects the magnitude of distance compression, nor has any aspect of virtual reality been found that induces distance compression in real environments.

2 Experiment 1

Thompson et al. showed that the level of graphical detail in a virtual environment does not significantly affect distance judgements(Thompson et al., in

press). The apparatus used in that study was limited by a need to restrict subject head motions. This restriction was necessary because the photorealistic environment was generated by a carefully spaced series of still images, stitched together using computer vision techniques which allow very accurate interpolation of the correct view between images, but only from a fixed point. Thus, while the system could accurately handle rotations, it needed to restrict head translations. This limitation prevents the use of motion parallax, and other motion cues, to determine distance, reducing the ecological validity of the study.

One way to provide a photorealistic level of graphical detail without restricting subject head motions would be to provide live video from a head-mounted video camera. The present experiment uses this technique. The use of a video camera to provide realtime video for a head-mounted display was inspired by an anecdote mentioned by Loomis and Knapp (J. M. Loomis & Knapp, 2003), though our findings are different from theirs.

There is reason to believe that stereoscopic displays may be inaccurately rendered in virtual reality (Wann, Rushton, & Mon-Williams, 1995). In addition, a number of studies have found no difference in distance perception between monocular and binocular viewing of distances greater than 2 meters (Creem-Regehr et al., 2003; Philbeck & Loomis, 1997). As a result of these findings, and the increased complexity of developing a binocular apparatus, this experiment uses only monocular viewing conditions.

Distance perception in the head-mounted display viewing condition was compared to direct viewing of the environment in two conditions. One condition used a pair of goggles that restricted the field of view to the same angle as the head-mounted display condition. This condition attempted to produce direct viewing of the environment with a retinal image as closely matched to the head-mounted display condition as possible. Another viewing condition was only restricted to monocular viewing. This condition was intended to control the restricted field of view condition. This was important because, although restricted field of view has been investigated previously (Creem-Regehr et al., 2003; Knapp & Loomis, (in press)), the simultaneous combination of restricted field of view and monocular viewing in distance perception is novel.

Although the head-mounted video camera provided photorealistic detail while allowing unrestricted head movements, it had some limitations. Practical concerns prevented the system from displaying the retinal image that would be seen from the subject's eye location. Instead, the camera was mounted on the side of the head-mounted display. This led to both a slight translation of the subject's viewpoint, and a shifted axis of horizontal rotation. The viewpoint translation was probably insignificant, because the change in apparent distance of a target several meters away would be very small. The rotational shift may have been more significant, but probably was not responsible for perceptual errors. This is because the head motions involved in viewing a target in depth are primarily translations and vertical rotations. The most significant limitation of the apparatus was the signal delay between the camera and the headset. This delay, between .5 and 1.0 seconds, might be enough to desynchronize motion parallax from actual head motion, which could weaken motion parallax as a cue

to distance.

While Thompson et al.'s study lacked motion parallax as a cue to distance, there is some evidence that it is a weak cue to distance by itself (Beall, Loomis, Philbeck, & Fikes, 1995). While motion parallax is tightly coupled with convergence in natural viewing circumstances, potentially strengthening both cues, motion parallax is a weak cue to distance in virtual environments with fixed convergence (Philbeck & Loomis, 1997).

2.1 Methods

2.1.1 Participants

Subjects were 20 naive undergraduates at Swarthmore College. Some participated as part of an introductory psychology course requirement, and others were volunteers given five dollars and a candy bar as compensation for their participation in the study. 7 participated in the head-mounted display condition, 7 in the restricted field of view condition, and 6 in the unrestricted viewing condition.

2.1.2 Apparatus

In all conditions, subjects wore a piece of foam, held on by an eyepatch, in order to ensure monocular viewing.

In the restricted field of view condition, subjects wore a pair of field of view restricting goggles. The field of view on these goggles was the same as the manufacturer-specified field of view for the head-mounted display, 51 degrees horizontal field of view by 38 degrees vertical field of view.

In the head-mounted display condition, the apparatus consists of a Sharp VL-Z3 digital video camera with a wide-angle lens attachment, whose field of view was approximately equal to the field of view of the head-mounted display. This camera was connected by IEEE-1394 cable to an Apple G4 running custom-written video software, which output video to a Virtual Research Systems V8 head-mounted display. The V8 head-mounted display has an accommodative distance of 1 meter. The camera was attached to the left side of the head-mounted display with a velcro strap, and stabilized with pieces of foam. The camera lens was approximately 10 cm from the subject's left eye. Latency from the camera to the display was determined to be at least 433 milliseconds, by taping an event actually occurring, and being output to a screen. While this method takes into account the delay in a signal being sent from the camera to the computer, it fails to account for a possible difference in display latency between the computer's screen and the head-mounted display. As a result, the 433 milliseconds may not fully account for the latency, but probably serves as a good approximation, and certainly an effective minimum bound.

In all conditions, including the unrestricted monocular viewing condition, subjects wore an eyepatch over one eye. In the restricted field of view condition, because of the design of the field of view restricting goggles, the eyepatch was



Figure 1: The hallway where experiment 1 took place, with a target 7 meters from the observer.

worn over the right eye. In the other conditions, subjects selected their preferred eye to remain unpatched.

In each condition, subjects were shown a target that was a light orange cylinder of foam, 22 cm in diameter and 1.25 cm high. The target is pictured in the hallway used for the experiment in figure 1.

All subjects wore hearing attenuators in order to prevent them from getting audible feedback from the environment.

2.1.3 Design

Subjects were assigned to one of three conditions, which differed in types of viewing. Subjects monocularly viewed the target through a live-video head-mounted display, a pair of field of view restricting goggles, or without restriction.

In all conditions, subjects were shown a target at distances of 2,3,4,5,6, and 7 meters in a hallway with which they had limited familiarity.

2.1.4 Procedure

Subjects performed a visually-directed walking task in a texture-rich hallway. This task required the subject to look at a target while standing in a fixed position. Subjects had no restrictions on time or head movements while looking at targets. Subjects then signaled readiness to begin walking, at which point they closed their eyes. In the head-mounted display condition, the display was also disabled at this point, and was not enabled until the next trial. After closing their eyes, subjects would walk to the target location, and the distance along the hallway from the subject to the target would be measured.

Subjects were told that they were performing a non-cognitive task investigating low-level visual perception, and that as a result, they were not to count

their steps, and were to walk quickly and confidently to the target. Subjects were also advised that the target itself would be removed after they closed their eyes, to prevent them from getting haptic feedback from the foam target itself.

Subjects performed 6 trials with targets at distances of 2-7 meters. Though subjects were not informed of it, the first two trials were recorded as practice, and the targets for those trials were placed, in random order, at 3 and 6 meters from the subject. The next four trials were experimental trials, with, in random order, distances of 4 and 5 meters at the ends of the trial block, and distances of 2 and 7 meters in the middle.

2.2 Results

One subject was removed from analysis of the head-mounted display condition because the subject's average ratio of perceived distance to actual distance was over 5 standard deviations (as measured excluding the subject) away from the other subjects in that condition.

In the head-mounted display condition, subjects walked an average of 77% of the distance viewed (averages were computed in log space) which was significantly different from 100%, $t(5) = 5.87, p < .01$. The 95% confidence interval for the mean distance walked in the head-mounted display condition (computed in log space) extends from 70% to 86%. The distances walked in the head-mounted display were also reliably less than the distances walked with a restricted field of view ($M = 96%$), $t(9) = 3.53, p < .01$, and less than the distances walked in the unrestricted monocular viewing condition, ($M = 96%$), $t(11) = 3.83, p < .01$. Neither of these latter conditions differed reliably from each other, nor did either differ reliably from accurate walking ($p > .10$). These results are illustrated in figure 2.

Fitting each subject's data to a power function, the average exponent in the head-mounted display condition was 0.95. The average exponent in the restricted field of view condition was 0.98. The average exponent in the unrestricted viewing condition was 0.93. None of these values were significantly different from an exponent of 1, and there was no significant difference between any of these values.

2.3 Discussion

We found significant distance compression in the head-mounted display condition, which displayed photorealistic live video. This is consistent with the findings of Thompson et al. (Thompson et al., in press), and shows that the finding persists after eliminating the limitations on translational head motion imposed by their apparatus.

We also found no significant difference between the restricted field of view condition, and the unrestricted viewing condition. This is consistent with the findings of Creem-Regehr et al., and Knapp and Loomis (Creem-Regehr et al., 2003; Knapp & Loomis, in press). Because previous studies have not investigated the combination of restricted field of view and monocular viewing, this

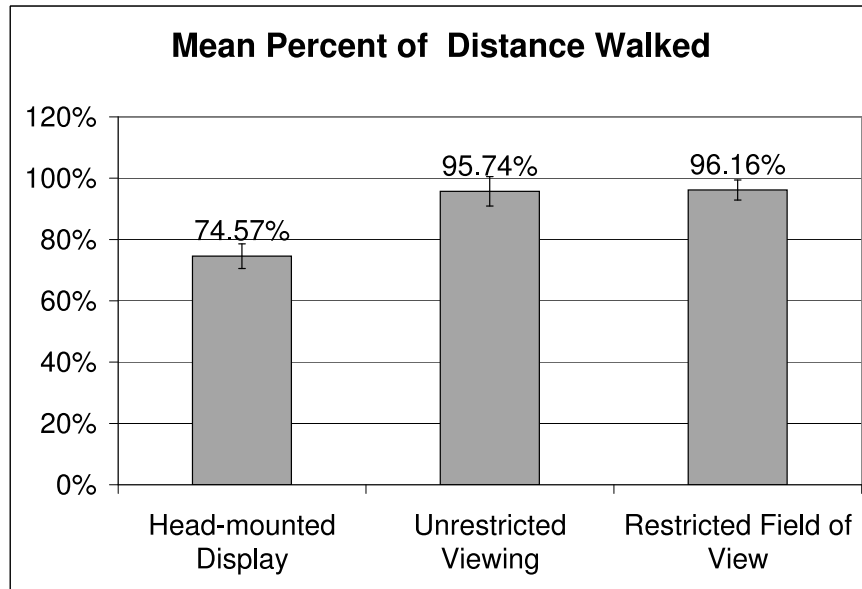


Figure 2: Mean percentage of distance to the target walked, by viewing condition.

study expands those previous findings.

While level of graphical detail and restricted field of view are not the cause of distance underestimation in virtual reality, some other aspect of the system must be responsible. One hypothesis is that the problems lie in part with the optics of the system. Because the system's lens cannot adjust as the virtual world changes, accommodation is always generating the same information for the observer, regardless of the scene being observed. In addition, because no eye-tracking system can track where in three-dimensional space the eye is focusing, virtual reality displays must be in sharp focus at all times. This means that depth blur, which would normally be present in objects not at the distance being focused on, also generates the same information regardless of the virtual scene being observed.

Loomis (personal communication) has investigated the effect of accommodative lenses, and doesn't think that they effect distance judgements at distances as far from the observer as those used in this experiment. This suggests that inaccurate accommodative information is not responsible for the underestimation of distance in virtual environments.

Despite the limitations of virtual reality for providing accurate distance information, it remains a potentially useful tool for investigating distance perception because of the ease with which a virtual environment can be manipulated. The remaining studies reported here will use virtual reality to investigate other hypotheses about distance perception. The problem of distance underestimation will be avoided by modelling subject data with power functions and examining

their exponents, thus ignoring linear effects.

3 Experiment 2

Ooi et al. showed that the angular declination of an object from the horizon serves as a cue to distance (Ooi et al., 2001). Subjects wore prisms which deviated the angle of everything they saw, making objects appear a few degrees closer to their feet, and thus causing underestimation of distance. Subjects were then adapted to the prisms, and when they were removed, subjects perceived the horizon as at a lower than it really was. This had the effect of making objects appear farther away, because as an object's angle of regard gets closer to the angle of the horizon, that is, as its angular declination from the horizon decreases, the object appears closer to the horizon, and thus further away. While the prism method used by Ooi et al. was quite thorough and well-controlled, virtual reality systems have more control over displays, and can vary the horizon line height without effecting the rest of the world.

In rendering virtual worlds, it can be very computationally expensive to render an observer's view as a scene extends infinitely. In order to solve this problem, a computer graphics system must stop rendering a scene as it extends past a fixed distance from the observer. This can result in an artificially low horizon line. However, if the eye height of the observer is known, the amount by which the horizon is lowered can be computed. If a sufficiently simple texture is used, the end of the ground plane can be artificially extended, replacing the horizon at the correct level. We took advantage of this process to investigate the effect of the height of the horizon on distance perception in virtual reality.

The findings of Ooi et al. (Ooi et al., 2001) suggest that distances should appear farther away if the horizon is artificially lowered than if it is not. This is illustrated in Figure 3. In order to lower the horizon by a fixed angle, a , assuming an observer eye height of 1.5 meters, rendering of the world must stop at $1.5m * \tan(90^\circ - a)$ meters from the observer. The correct horizon line should appear at eye level, and object X has angular declination b from the horizon. If the horizon line is lowered by a degrees, then object X has angular declination $b - a$ from the new horizon. If the observer's eye height is 1.5 meters, object X will appear to be $\tan(90 - (b - a)) * 1.5$ meters from the observer.

A horizon a degrees low will cause objects to appear to be a degrees closer to the horizon. Because this constant angular quantity signifies more and more actual distance as an object gets closer to the horizon line, the amount by which a low horizon effects the perception of distance increases with distance from the observer.

Distance perception can be modelled in different ways. Traditionally, distance perception has been modelled as a power function (M. Teghtsoonian & Teghtsoonian, 1969; R. Teghtsoonian & Teghtsoonian, 1970). Modelling distance perception data by fitting a subject's perceptions to a power function characterizes the rate at which perceived distance increases as actual distance increases. Recent work on distance perception has used linear modelling (J.

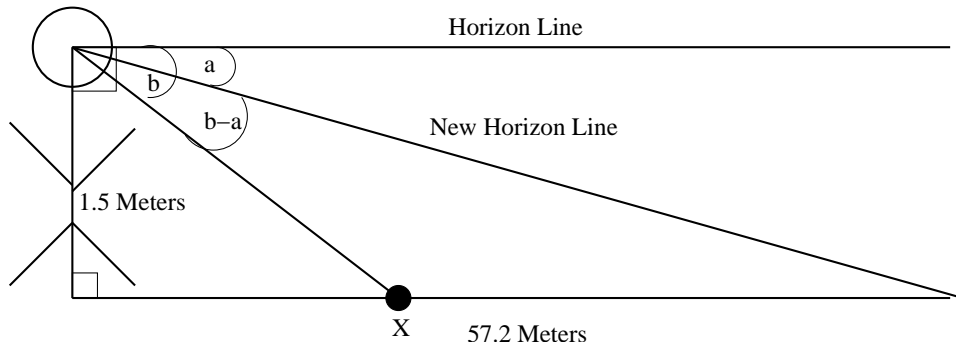


Figure 3: This figure depicts the effects of lowering the horizon line by 1.5 degrees.

Loomis et al., 1992; Thompson et al., in press). Linear models of distance perception have allowed investigators to characterize a number of findings, like the systematic underperception of distance in virtual reality displays. While linear modelling is a useful tool, power function modelling seems more appropriate to the present study. This is because we do not expect a linear finding, like the compression of distance in virtual reality. We expect that the horizon line manipulation will impact distance perception more as distance from the observer increases (See Figure 4).

Perceived distance can be indicated in a number of different ways. One way to implicitly indicate distance, used in the previous experiment, is visually directed action. Another way to indicate perceived distance is explicit verbal judgement. Goodale and Milner hypothesize two separate visual pathways, one for visual perception, and the other for perceptual action (Goodale & Milner, 1992). This theory would suggest and explain differences between a strictly perceptual task, like verbal estimation of distance, and a visually directed action task. Loomis et al., on the other hand, theorize that humans have one underlying egocentric representation of visual space (J. Loomis et al., 1992). This theory suggests that there should be little difference in performance between the two tasks, as they both rely on the same underlying representation of the viewed scene.

While differences in distance perception between the two tasks could be seen as supporting Goodale and Milner's theory, it could also be explained by the differing task demands. Verbal judgements might make subjects more likely to look at the whole scene, while a walking task might make subjects pay more attention to the ground between them and the target. If task strategies differ in this way, verbal judgements might be more greatly effected by variation of the horizon line.

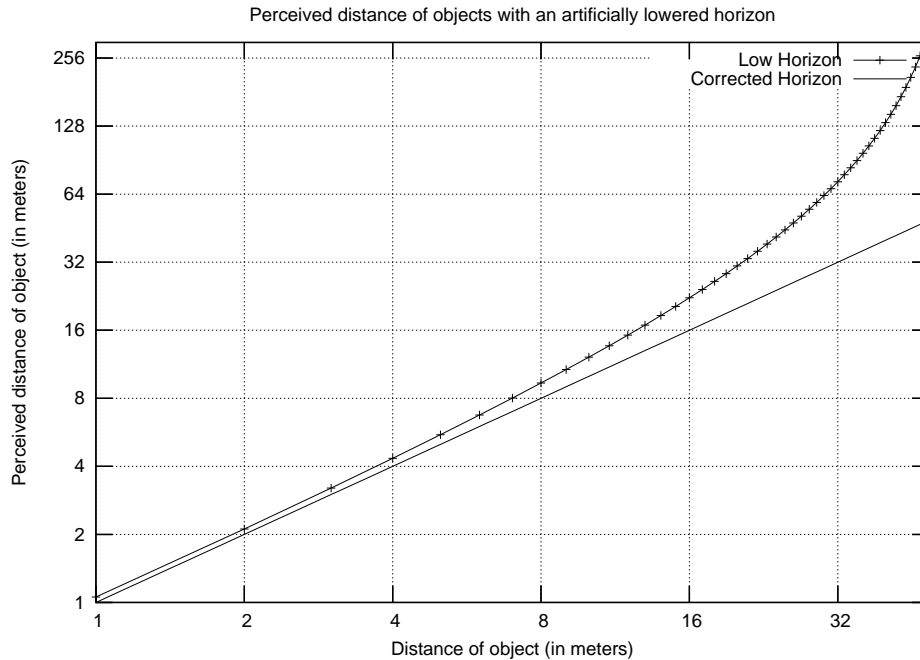


Figure 4: The effect on distance perception of an artificially 1.5° low horizon created by cutting the rendering plane at about 57.28 meters.

3.1 Methods

3.1.1 Participants

Subjects were 20 naive undergraduates at Swarthmore College who participated as part of an introductory psychology course requirement. In each of the verbal response and visually directed action response conditions, 5 of the subjects were male, and 5 female.

3.1.2 Apparatus

Subjects wore a Virtual Research Systems V8 head-mounted display with a 3rd Tech HiBall head tracker. Subjects also wore hearing attenuators. The virtual environments were rendered by two G4 computers. Subjects viewed the virtual environments binocularly.

3.1.3 Displays

Subjects were shown an outdoor virtual environment with a grassy field. This environment included a red circle, 1 foot(0.3048 meters) in diameter, where

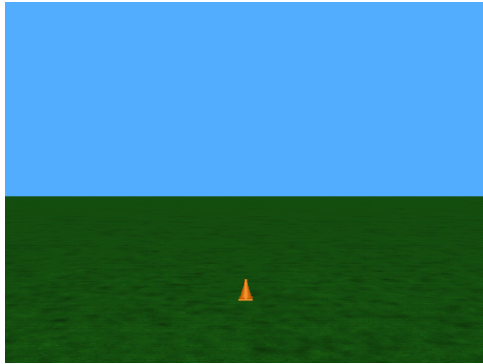


Figure 5: The outdoor virtual environment with a 1.5° low horizon, and a target 7 meters from the observer.




Figure 6: The outdoor virtual environment with a correct horizon, and a target 7 meters from the observer.

subjects were supposed to stand, and an orange cone which served as the target in both conditions. The target remained the same size in the virtual world, so its retinal image varied with its location. The target was rendered without contact cues, but in a world with diffuse lighting and targets at a minimum of 3 meters, contact cues like shadows would probably not have been noticeable on a small object like the target.

Two variations of the outdoor environment were used. In one, the horizon line was positioned 1.5 degrees too low, as in figure 5. In the other, the horizon line was correct, as in figure 6.

Subjects in the visually directed action condition were also shown a featureless indoor environment with a yellow line in the middle in order to aid in repositioning between trials. This environment is shown in figure 7.




Figure 7: The featureless indoor hallway environment used to reposition subjects in the visually directed action condition.

3.1.4 Design

Subjects were assigned to one of two response conditions, either verbal judgement, or visually directed action.

Subjects were shown three practice targets, at distances of 5 meters, 10 meters, and 1 meter, and were then shown 20 experimental targets, which were assigned a distance between 3 and 7.5 meters, and randomly shifted by up to half a meter away from the subject. All of the practice trials and half of the experimental trials that each subject saw were in the corrected-horizon world, while the rest were in the uncorrected-horizon world. The two worlds were randomly interleaved in the experimental trials.

3.1.5 Procedure

Subjects were instructed differently depending upon their task. In the walking conditions, subjects were instructed that they were not performing a cognitive task, and were instructed not to use explicitly cognitive methods of distance estimation, like counting increments of known distances to the target. Subjects were told that the random shifting of target distances would prevent memorizing distances from helping them. Subjects were instructed to look at the world for as long as they felt they needed, and to then close their eyes and walk quickly and confidently to where the target was. Subjects were positioned over a red circle, and instructed that moving away from that circle would blank the display. Subjects were not told that the first three trials were practice.

In the verbal condition, subjects were instructed to give estimations of distance in feet and inches. They were instructed to be as accurate and consistent as possible, and to give as specific estimates of distance as possible. Subjects were also told that distances would be randomly shifted, so that memorization of distances would not be helpful. As in the walking condition, subjects were instructed not to count increments of known distances to the target. This was done in order to vary as little as possible besides report modalities between the conditions. Subjects were positioned on top of a red circle, and informed accurately that it was a foot in diameter, though they were cautioned to only use that as a reference, rather than attempting to count the number of circles out to the target.

In the walking condition, subjects were trained to walk quickly and confidently before the experiment started. Subjects were instructed to hold the headset over their shoulder and walk quickly up and down the hallway. The headset then gave them feedback about whether or not they had walked fast enough.

3.2 Results

Each subject's distance perception data was fitted to a power function. A repeated measures ANOVA was conducted on the power function exponents with horizon (normal or low) as a within-subject variable and sex (male or




Figure 8: Average exponents of power functions modelling subject performance in verbal estimation and visually directed action in an outdoor virtual environment with a correct or artificially low horizon line.

female) and response (verbal estimation or visually directed action) as between-subject variables. Consistent with the predicted effects of the horizon cue, exponents were higher when the horizon was low ($M = 1.095$), than when it was normal ($M = 0.983$), $F(1, 16) = 8.71, p < .01$. This is illustrated in figure 8.

Each subject's average ratio of estimated to actual distance was also computed. A repeated measures ANOVA was conducted on the ratio of estimated to actual distance with horizon (normal or low) as a within-subject variable and sex (male or female) and response (verbal estimation or visually directed action) as between-subject variables. There were no reliable differences found in this measure, though there was a trend for visually directed action to produce greater estimates of distance than verbal estimation. This is illustrated in figure 9.

3.3 Discussion

As predicted, lowering the horizon increased the exponent of the power function fitting subjects' perception of distance. The power function fitting the predicted performance of subjects seeing targets from 3 to 7.5 meters in the low horizon condition had a mean exponent of 1.08, which is very close to the value found. Since Ooi et al.'s investigations took place in indoor environments, the present experiment extends their findings to environments with visible horizon lines (Ooi et al., 2001).

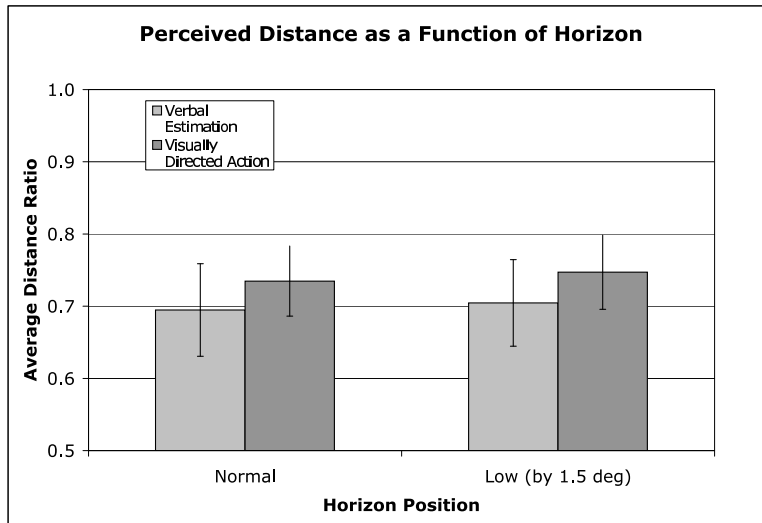


Figure 9: Average proportion of distance walked to distance viewed in verbal estimation and visually directed action in an outdoor virtual environment with a correct or artificially low horizon line.

This experiment showed very similar performance in verbal estimation and visually directed action. This is consistent with Loomis et al.'s theory of a single egocentric representation of visual space (J. Loomis et al., 1992). While there was an unreliable trend for verbal estimation to produce lower ratios of perceived to actual distance than visually directed action, this could be because producing an explicit verbal estimate of distance may require more cognitive processing than an implicit action-based estimation. This extra processing could lead to the added inaccuracy.

One factor that must be considered when working in virtual reality is the level of detail of the virtual environments, and how this can impact ecological plausibility. While Thompson et al. (Thompson et al., in press) and experiment 1 have shown that the level of detail of a virtual environment may be unrelated to distance compression, that level of detail surely impacts the quantity and quality of depth information. The virtual environment used in this experiment is a fairly simple one, with very few objects. Despite having a richly textured ground plane, the environment's paucity of features, and the presence of a visible horizon line, may have emphasized the use of the horizon line as a depth cue more than a more ecologically valid environment would have. Nonetheless, the results we have obtained are consistent with those of Ooi et al., in indicating that even a very subtle manipulation of the horizon line position can produce measurable effects on both motor and verbal judgments of distance.

4 Experiment 3

Experiment 2 showed similar performance from subjects in both the verbal and motor conditions. This similarity of performance seems consistent with Loomis et al.'s theory of a single underlying egocentric representation of visual space (J. Loomis et al., 1992). Differences in performance between the two tasks, on the other hand, would tend to support Goodale and Milner's theory of two separate visual pathways for perception and action (Goodale & Milner, 1992). The next experiment investigated a situation expected to produce dissociation between verbal estimation and visually directed action.

In a series of studies, Teghtsoonian and Teghtsoonian found a difference in distance perception between indoor and outdoor environments (M. Teghtsoonian & Teghtsoonian, 1969; R. Teghtsoonian & Teghtsoonian, 1970). To ensure the robustness of their findings, they compared distance perception across varying ranges of distances. Modelling their results with power functions, they found an average exponent of about 1.2 in indoor environments, regardless of range, and average exponents ranging from 0.99 to 0.85 in outdoor environments, decreasing as range increased. Both of these studies used explicit verbal report of distance, rather than a visually directed action task. These studies clearly show a difference in distance perception, as measured by verbal report, between an indoor and an outdoor environment.

Durgin et al. investigated locomotor adaptation in an indoor virtual environment (Durgin, Fox, Lewis, & Walley, 2002). In pre-adaptation subjects, they found power functions with exponents slightly lower than 1 with a motor task, and slightly greater than 1 with a verbal task. However, their experiment was not designed to investigate this issue, and it is unknown whether the findings of Teghtsoonian and Teghtsoonian will replicate in virtual reality.

4.1 Methods

4.1.1 Participants

Subjects were 26 naive undergraduates at Swarthmore College who participated as part of an introductory psychology course requirement. In the verbal response condition, 8 of the subjects were male, and 6 were female. In the visually directed action response condition, 6 of the subjects were male, and 6 were female.

4.1.2 Apparatus

The same system used in experiment 2 was used here. The virtual environments were the only change.

4.1.3 Displays

Subjects saw both the corrected-horizon world from the previous experiment (see figure 6), and a blue textured hallway. The hallway was 2 meters wide, and



Figure 10: The indoor virtual environment with a target 7 meters from the observer.

2.5 meters tall. The hallway, like the grassy field, had both the red circle at the subject's feet, and the orange cone that served as the target. The hallway is shown in figure 10.

4.1.4 Design

The design of this experiment was similar to the design of experiment 2, but the two worlds were shown in five-trial counterbalanced blocks, rather than interleaved. Subjects were shown the 3 practice trials in whichever of the worlds the subjects would see in their first block. This meant that the first block of trials could serve as a between-subject design, much as Teghtsoonian and Teghtsoonian used a between-subject design in their experiments (M. Teghtsoonian & Teghtsoonian, 1969; R. Teghtsoonian & Teghtsoonian, 1970).

4.1.5 Procedure

The procedure used here was the same as in experiment 2.

4.2 Results

Two male subjects were removed from analysis of the verbal report condition because their average ratios of perceived distance to actual distance were over 3 standard deviations (as measured excluding the subjects) away from the other subjects in that condition. This brings the number of subjects to 6 males and 6 females in both the verbal estimation and visually directed action conditions.

A repeated measures ANOVA was conducted on the power function exponents with environment (indoor or outdoor) as a within-subject variable and sex (male or female) and response (verbal estimation or visually directed action) as between-subject variables. A marginal interaction between response type and world was revealed, $F(1, 20) = 3.445, p = .078$. Separate ANOVAs were therefore conducted for each world. There was no reliable difference between the

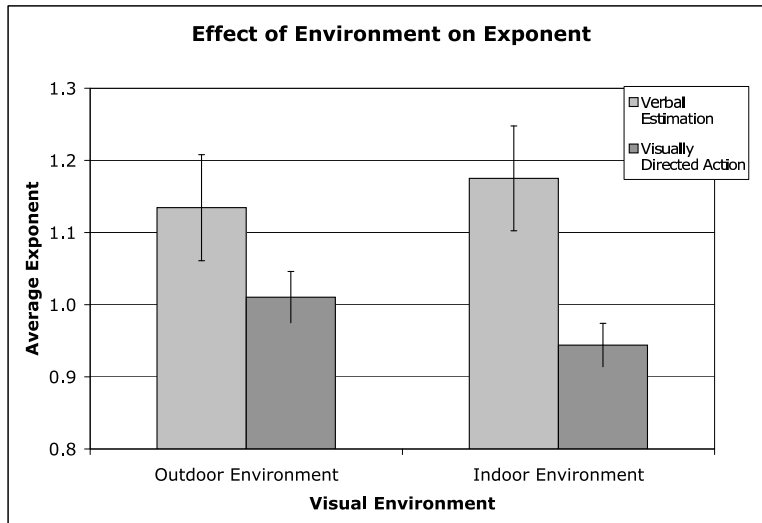


Figure 11: Average exponents of power functions modelling subject performance in verbal estimation and visually directed action in indoor and outdoor virtual environments.

visually directed action exponent (1.06) and the verbal exponent (1.13) in the outdoor environment, $F(1, 22) < 1$. However, in the indoor environment, the verbal exponent (1.18) was reliably greater than the visually directed action exponent (0.98), $F(1, 22) = 5.35, p < .05$. This is illustrated in figure 11. Indeed, only the verbal response exponent in the indoor environment was reliably greater than 1, $t(11) = 2.41, p < .05$.

Each subject's average ratio of estimated to actual distance was also computed. A repeated measures ANOVA was conducted on the ratio of estimated to actual distance with environment (indoor or outdoor) as a within-subject variable and sex (male or female) and response (verbal estimation or visually directed action) as between-subject variables. The indoor environment produced reliably greater estimates of distance than the outdoor environment, $F(1, 20) = 8.12, p < .01$. There was an unreliable trend for visually directed action to produce greater estimates of distance than verbal estimation. This is illustrated in figure 12.

A between subject analysis looking only at the first experimental block showed a similar pattern. In the outdoor virtual environment, the mean verbal exponent (1.10) did not differ from the mean visually directed action exponent (0.98), $F(1, 10) < 1$. However, in the indoor virtual environment, the verbal exponent during the first block (1.33) was higher than the corresponding visually directed action exponent (0.97), $F(1, 10) = 6.21, p < .05$. This is illustrated in figure 13. Perhaps most importantly, the verbal exponent in the indoor virtual environment is reliably greater than 1, $t(5) = 2.65, p < .05$.

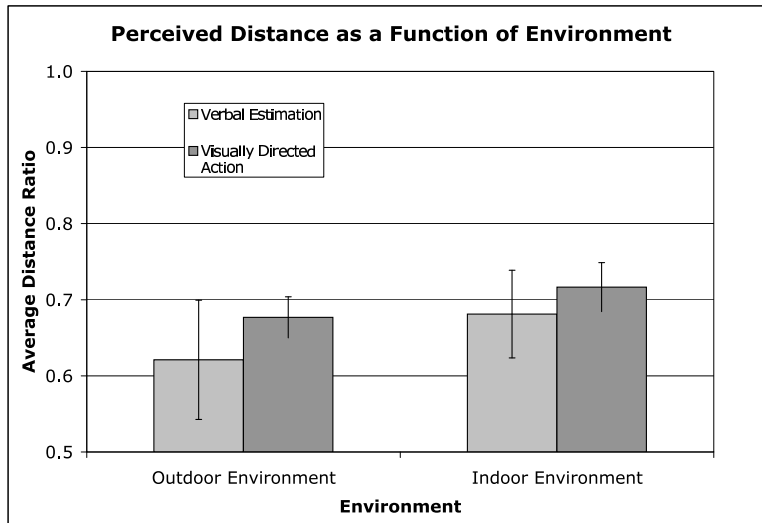


Figure 12: Average proportion of distance walked to distance viewed in verbal estimation and visually directed action in indoor and outdoor virtual environments.

In order to illustrate that an exponent greater than 1 means that distance ratios increase with distance, the average distance ratio given by verbal estimation was calculated for each distance category, as shown in figure 14. The distances are grouped into categories because each distance was randomly extended by 0 to 0.5 meters.

4.3 Discussion

In this experiment, the fitted power functions of subjects performing verbal estimation showed an exponent reliably greater than 1 in the indoor virtual environment, but not in the outdoor virtual environment. Teghtsoonian and Teghtsoonian modelled distance perception in real indoor and outdoor environments with power functions, and found an exponent significantly above 1 in indoor environments, and an exponent very close to 1 in outdoor environments (M. Teghtsoonian & Teghtsoonian, 1969; R. Teghtsoonian & Teghtsoonian, 1970). The results of this experiment are consistent with the findings of Teghtsoonian and Teghtsoonian, and extend their findings to virtual environments.

While the results from the verbal estimation condition were consistent with the findings of Teghtsoonian and Teghtsoonian, the visually directed action condition showed no significant difference in distance perception between environments. In fact, while it was not statistically significant, the visually directed action condition showed a small trend in the opposite direction, with power function exponents greater in the indoor virtual environment than the outdoor

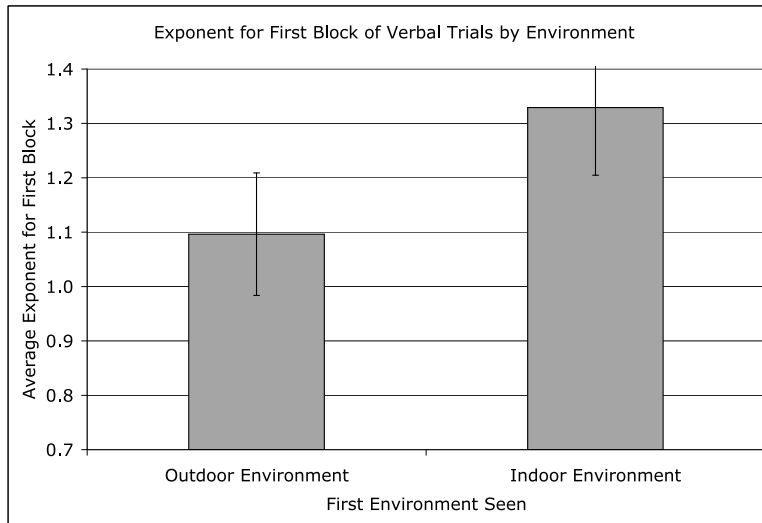


Figure 13: Average exponents of power functions modelling subject performance in verbal estimation during their first block in indoor and outdoor virtual environments.

one. It is likely that the mean exponent for distance perception in the indoor virtual environment is not just a result of small sample size, as the exponent is very close to those found in experiment 1, which also involved a visually directed action task in an indoor environment, although not a virtual one.

The different results seen for verbal estimation and visually directed action seem, unlike experiment 2, to be consistent with Goodale and Milner’s theory of separate visual pathways for perception and action (Goodale & Milner, 1992). However, in the context of the similar results for the two tasks in experiment 2, it is hard to make any strong claims. Since the only point where there was a significant difference between the two tasks occurred in the indoor environment, it seems likely that either a cue that is present in the indoor environment is misleading in some way that is more salient to the verbal task, or a cue present in the outdoor environment is accurate in some way more relevant to the visually directed action task. Another possibility is that the different environments cause different weighting of cues to distance. This is almost certainly the case with the cue of distance by angular declination from the horizon, because the horizon is not present in indoor environments, and thus ought to be ignored.

5 Conclusion and General Discussion

Experiment 1 investigated distance underperception in virtual reality. Distance underperception was found in a live video head-mounted display condition, and not in a monocular restricted field of view condition or an unrestricted monoc-

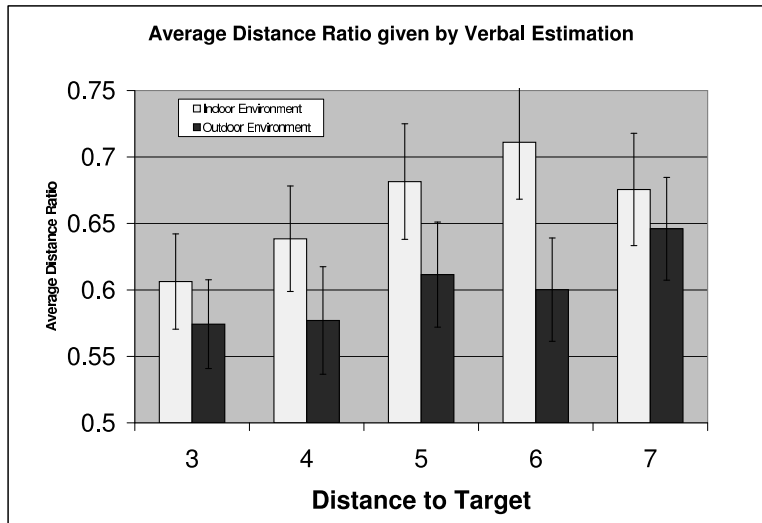


Figure 14: Average distance ratio given by verbal estimation for each distance category. Distances in meters.

ular viewing condition. These findings suggest that the restricted field of view and the level of realism in a virtual reality head-mounted display are not responsible for the underestimation of distance. While there were some limitations of the apparatus used, particularly the latency of the head-mounted display condition, the results strongly suggest that the cause of underperception of distance in virtual reality head-mounted displays lies elsewhere in the virtual reality system. In particular, the cues given by accommodation and depth blur in virtual reality displays can be misleading, if the focal length of the head-mounted display and the distance of the virtual object being viewed differ, and at least one is within the range where accommodations is an effective cue to distance (about 8 feet)(Palmer, 1999). Despite these limitations in virtual reality, it offers very powerful control over displayed environments. The next experiments used virtual reality to investigate distance perception, while using power function data modelling to avoid problems associated with distance underperception.

Experiment 2 used virtual reality to investigate declination from the horizon as a cue to distance. By varying the height of the horizon, it was shown that declination from the horizon does serve as a cue to distance. Subjects in both a verbal estimation task and a visually-directed action task performed similarly.

Experiment 3 tried to investigate a possible dissociation between distance perception as measured by verbal estimation and visually directed action. Teghtsoonian and Teghtsoonian had shown different performance on a verbal estimation task in real indoor and outdoor environments(M. Teghtsoonian & Teghtsoonian, 1969; R. Teghtsoonian & Teghtsoonian, 1970), while Durgin et al. had found differing performance in verbal estimation and visual action tasks in an

indoor virtual environment (Durgin et al., 2002). By using both indoor and outdoor virtual environments, the experiment was able to extend the work of Teghtsoonian and Teghtsoonian to virtual reality. The findings for the verbal task were consistent with the findings of Teghtsoonian and Teghtsoonian, while the motor task showed no significant difference between the indoor and outdoor virtual environments.

In the context of experiment 2, which showed similar performance in both the verbal estimation and visually directed action tasks, and experiment 3, which showed a dissociation between the two tasks, it is hard to make a strong claim about the underlying nature of the visual system. Experiment 2 supports Loomis et al.'s theory of a single underlying representation of visual space, while experiment 3 supports Goodale and Milner's theory of two separate visual pathways for perception and action (J. Loomis et al., 1992; Goodale & Milner, 1992). However, experiment 2 showed trends that, while not reliable, differentiated the two tasks. These differences could be the result of the two tasks using two separate pathways, but they could also be the result of other differences between the tasks, like the degree of cognition required. Experiment 3, on the other hand, showed the two tasks differing in performance in the indoor virtual environment, but not in the outdoor virtual environment. This suggests that the difference observed could be due to differences between the environments, rather than the underlying representation of space. One thing that might differ between environments is the weighting of cues in cue combination. In addition, some cues are only present in some environments. For example, declination from the horizon line and aerial perspective are only present in outdoor environments.

The textures used for the indoor and outdoor virtual environments in experiment 3 have different colors, patterns, and gradients. One way to further investigate the findings from experiment 3 could be to use the same texture in both indoor and outdoor virtual environments. I hypothesize that the dissociation between virtual environments in the verbal report task is not texture-dependant, but rather results from a fundamentally different weighting of cues in indoor and outdoor environments.

6 Acknowledgements

I would like to thank Professor Frank Durgin, my thesis advisor, for the huge amount of help he provided in conceiving, designing, executing, and analyzing the experiments described here, and for helping me draft this document. I would like to thank Joy Mills for helping me run experiment 1. I would like to thank Dr. William Thompson for his helpful comments on an earlier version of this paper. I would like to thank the Swarthmore College department of Psychology for awarding me the Hans Wallach Fellowship, which funded my research during the summer of 2003, and allowed me to begin the work reported here. Last but not least, I would like to thank my parents, Edward and Susan Messing, for their love, confidence, and support.

References

- Beall, A. C., Loomis, J. M., Philbeck, J. W., & Fikes, T. G. (1995). Absolute motion parallax weakly determines visual scale in real and virtual environments. *Proceedings of the SPIE - The international Society for Optical Engineering*, 2411, 288-297.
- Creem-Regehr, S. H., Willemsen, P., Gooch, A. A., & Thompson, W. B. (2003, August). *The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual environments* (Tech. Rep. No. UUCS-03-016). University of Utah.
- Cutting, J. E., & Vishton, P. M. (1995). Perceiving layout and knowing distances: The integration, relative potency, and contextual use of different information about depth. In W. Epstein & S. Rogers (Eds.), (p. 69-117). Academic Press.
- Dixon, M., Wraga, M., Proffitt, D., & G. Williams. (2000). Eye height scaling of absolute size in immersive and nonimmersive displays. *Journal of Experimental Psychology: Human Perception and Performance*, 26(2), 582-593.
- Durgin, F. H., Fox, L. F., Lewis, J., & Walley, K. (2002, November). Perceptuomotor adaptation: More than meets the eye. In *43rd annual meeting of the psychonomic society*. Kansas City, MO.
- Feria, C. S., Braunstein, M. L., & Andersen, G. J. (2003). Judging distance across texture discontinuities. *Perception*, 32(12), 1423-1440.
- Gibson, J. J. (1986). *The ecological approach to visual perception*. Lawrence Erlbaum Associates.
- Gilinsky, A. S. (1951). Perceived size and distance in visual space. *Psychological Review*, 58, 460-482.
- Goodale, M., & Milner, A. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15(1), 20-25.
- Hillis, J. M., Ernst, M. O., Banks, M. S., & Landy, M. S. (2002). Combining sensory information: Mandatory fusion within, but not between senses. *Science*, 298, 1627-1630.
- Knapp, J. M., & Loomis, J. M. ((in press)). Limited field of view of head-mounted displays is not the cause of distance underestimation in virtual environments. *Presence: Teleoperators and Virtual Environments*.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Research*, 35(3), 389-412.
- Liu, Z., Knill, D., & Kersten, D. (1995). Object classification for human and ideal observers. *Vision Research*, 35(4), 549-568.
- Loomis, J., Silva, J., Fujita, N., & Fukusima, S. (1992). Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance*, 18(4), 906-921.
- Loomis, J. M., Blascovich, J. J., & Beall, A. C. (1999). Immersive virtual environment technology as a basic research tool in psychology. *Behavior Research Methods, Instruments, and Computers*, 31(4), 557-564.
- Loomis, J. M., & Knapp, J. M. (2003). Visual perception of egocentric distance

- in real and virtual environments. In L. J. Hettinger & M. W. Haas (Eds.), (p. 21-46). Erlbaum.
- Michaels, C. F. (2000). Information, perception, and action: What should ecological psychologists learn from milner and goodale (1995)? *Ecological Psychology*, *12*(3), 241-258.
- Ooi, T. L., Wu, B., & He, Z. J. (2001). Distance determined by the angular declination below the horizon. *Nature*, *414*, 197-200.
- Palmer, S. E. (1999). Vision science : photons to phenomenology. In (p. 199-253). MIT Press.
- Philbeck, J. W., & Loomis, J. M. (1997). Comparison of two indicators of perceived egocentric distance under full-cue and reduced-cue conditions. *Journal of Experimental Psychology: Human Perception and Performance*, *23*(1), 72-85.
- Rowland, J. H. (1999). *The effects of texture on distance estimation in synthetic environments*. Master's Thesis.
- Sinai, M. J., Ooi, T. L., & He, Z. J. (1998). Terrain influences the accurate judgement of distance. *Nature*, *395*, 497-500.
- Teghtsoonian, M., & Teghtsoonian, R. (1969). Scaling apparent distance in natural indoor settings. *Psychonomic Science*, *16*(6), 281-283.
- Teghtsoonian, R., & Teghtsoonian, M. (1970). Scaling apparent distance in a natural outdoor setting. *Psychonomic Science*, *21*(4), 215-216.
- Thompson, W. B., Willemsen, P., Gooch, A. A., Creem-Regehr, S. H., Loomis, J. M., & Beall, A. C. (in press). Does the quality of the computer graphics matter when judging distance in visually immersive environments? *Presence: Teleoperators and Virtual Environments*.
- Wann, J. P., Rushton, S., & Mon-Williams, M. (1995). Natural problems for stereoscopic depth perception in virtual environments. *Vision Research*, *35*(19), 2731-2736.