Does the Quality of the Computer Graphics Matter When Judging Distances in Visually Immersive Environments?

***William B. Thompson, Peter Willemsen, Amy A. Gooch, Sarah H. Creem-Regehr, Jack M. Loomis*[1]*, and Andrew C. Beall*[1]**

UUCS–02–015

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

December 5, 2002

Abstract

In the real world, people are quite accurate judging distances to locations in the environment, at least for targets resting on the ground plane and distances out to about 20m. Distance judgments in visually immersive environments are much less accurate. Several studies have now shown that in visually immersive environments, the world appears significantly smaller than intended. This study investigates whether or not the compression in apparent distances is the result of the low-quality computer graphics utilized in previous investigations. Visually-directed triangulated walking was used to assess distance judgments in the real world and three virtual environments with graphical renderings of varying quality.

[1]Department of Psychology, University of California – Santa Barbara

Does the Quality of the Computer Graphics Matter When Judging Distances in Visually Immersive Environments?

William B. Thompson, Peter Willemsen, and Amy A. Gooch
School of Computing, University of Utah
Salt Lake City, UT 84112
Sarah H. Creem-Regehr
Psychology Department, University of Utah
Salt Lake City, UT 84112
Jack M. Loomis and Andrew C. Beall
Department of Psychology, University of California
Santa Barbara, CA 93106

Abstract

In the real world, people are quite accurate judging distances to locations in the environment, at least for targets resting on the ground plane and distances out to about 20m. Distance judgments in visually immersive environments are much less accurate. Several studies have now shown that in visually immersive environments, the world appears significantly smaller than intended. This study investigates whether or not the compression in apparent distances is the result of the low-quality computer graphics utilized in previous investigations. Visually-directed triangulated walking was used to assess distance judgments in the real world and three virtual environments with graphical renderings of varying quality.

1 Introduction

The utility of visually immersive interfaces for applications such as simulation, education, and training is in part a function of how accurately such interfaces convey a sense of the simulated world to a user. In order for a user to act in a virtual world as if present in the physical world being simulated, he or she must perceive spatial relations the same way they would be perceived if the user was actually in the physical world. Subjectively, current-generation virtual worlds often appear smaller than their intended size, impacting a user's ability to accurately interact with the simulation and the potential to transfer the spatial knowledge back to the real world.

Controlled experiments done by several different research groups are starting to provide objective evidence for this effect: Distance judgments to targets presented in visually immersive displays are often significantly compressed. There has been much speculation about the cause of this effect. Limited field of view (FOV), the difficulties in accurately presenting binocular stereo using devices such as head-mounted displays (HMDs), errors in accommodation, and limits on sharpness and resolution have all been suggested as potentially contributing to the misperception of distance (Rolland, Gibson, and Arierly 1995; Ellis and Menges 1997; Witmer and Sadowski 1998). Loomis and Knapp (in press) hypothesize that distance judgments are compressed in visually immersive environments because "the rendering of the scenes ... is lacking subtle but important visual cues (e.g., natural texture, highlights).... If this hypothesis is correct, it means that photorealistic rendering of the surfaces and objects in a simulated environment is likely to produce more accurate perception of distance...."

This paper explores the conjecture that image quality affects distance judgments in virtual environments. We start with a discussion of what is meant by a "distance judgment" and point out that different types of distance judgments likely depend on distinctly different visual cues. We next discuss how to experimentally determine perceptual judgments of one type of perceived distance. This is followed by the presentation of

experimental results comparing distance judgments in the real world with judgments based on graphical renderings of varying quality, showing that quality of graphics has little effect on the accuracy of distance judgments. We end with a discussion contributing to the speculation on why distances are incorrectly perceived in visually immersive displays.

2 Background

2.1 Visual cues for distance

Visual perception of distance can be defined in multiple ways. It is often categorized by the frame of reference used. *Egocentric* distances are measured from the observer to individual locations in the environment. *Exocentric* distances are measured between two points in the environment. The distinction is important for two reasons. First of all, the errors associated with the perception of egocentric and exocentric distances are different. Loomis, Da Silva, Philbeck, and Fukusima (1996) suggest that the perception of exocentric distance is dissociated from the perception of location. Although people may accurately perceive an egocentric location, they make systematic errors in perceiving an exocentric interval. Secondly, some depth cues such as shadows can provide information about exocentric distances but not egocentric distances.

Another distinction between types of distance perception is also critical. Distance perception can involve *absolute*, *relative*, or *ordinal* judgments. Absolute distances are specified in terms of some standard that need not be in the visual field (e.g., "two meters" or "five eye-heights"). Relative distances are specified in terms of comparisons with other visually determined distances (e.g., "location *A* is twice as far away as location *B*"). Relative distances can be thought of as absolute distances which have been subjected to an unknown but fixed scaling transformation. Ordinal distances are a special case of relative distances in which it is only possible to determine the depth ordering between two locations, but not the magnitude of the difference.

Finally, distance from the observer affects the nature and accuracy of distance perception. Cutting and Vishton (1995) divide distances into three zones: *personal space*, which extends slightly beyond an arm's reach from the observer, *action space*, within which we can rapidly locomote and extending from the boundaries of personal space to approximately 30m from the observer, and *vista space* beyond 30m from the observer.

The study reported on below deals with absolute egocentric distance judgments in action space, which are particularly relevant to many virtual environment applications. A computational analysis shows that only a few visual cues provide information about such distances (Figure 1). Accommodation and binocular disparity are not effective beyond a few meters. Absolute motion parallax has the potential to provide information about absolute egocentric distance if the velocity of the observer is utilized for scaling, but this appears to be a weak distance cue for people (Beall, Loomis, Philbeck, and Fikes 1995). Within action space, the related cues of linear perspective, height in the field, and horizon ratio are relative depth cues that have the potential for providing absolute depth to objects resting on a ground plane, when combined with information about the observer's eye height above the ground plane (Wraga 1999). These cues can be exploited in visually immersive interfaces if the rendering geometry is correct and both observer and object are in contact with a ground plane having adequate perspective cues. Familiar size – which involves exploiting the relationship between the assumed physical size of an object, the distance of the object from the observer, and the retinal size of the image of the object – can also serve as an absolute depth cue. It is reasonable to assume that the effectiveness of the familiar size cue depends at least in part of the realism of the imagery being viewed, though we are not aware of definitive studies addressing this issue. In the experiment described below, we vary the quality of immersively viewed imagery while fixing the information available from perspective cues in order to determine if image quality affects absolute egocentric depth judgments.

Cue	a	r	o	*Requirements for absolute depth*
Accommodation	x	?	?	very limited range
Binocular convergence	x	x	x	limited range
Binocular disparity	-	x	x	limited range
Linear perspective, height in picture, horizon ratio	x	x	x	requires viewpoint height
Familiar size	x	x	x	
Relative size	-	x	x	subject to errors
Aerial perspective	-	x	x	adaptation to local conditions
Absolute motion parallax	?	x	x	requires viewpoint velocity
Relative motion parallax	-	-	x	
Texture gradients	-	x	-	
Shading	-	x	-	
Occlusion	-	-	x	

Figure 1: Common visual cues for absolute (a), relative (r), and ordinal (o) depth.

2.2 Experimentally estimating judgments of absolute egocentric distance

It is quite difficult to determine the distance to a target that is "seen" by an observer. This is particularly true for absolute distance judgments, since methods involving just-noticeable-differences, reference standards, and equal interval tasks all involve relative distance. Verbal reporting can be used (e.g., "How many meters away is location *A*?"), but verbal reports tend to be noisy and are subject to a variety of biases that are difficult to control. An alternative for evaluating the perception of distance is to have subjects perform some task in which the actions taken are dependent on the perceived distance to visible objects (Loomis, Da Silva, Fujita, and Fukusima 1992; Rieser 1998). Such approaches have the additional advantage of being particularly appropriate for evaluating the effectiveness of interactive virtual environments.

Walking to or towards previously viewed targets has been used extensively to evaluate judgments of absolute egocentric distance. In one form of this task, subjects first look at a target and then walk to the target while blindfolded. They are told to stop at the target location and the distance between their starting and stopping points is presumed to be an indication of the originally perceived distance (Thomson 1983; Rieser, Ashmead, Talor, and Youngquist 1990). A second form of this task involves looking at a target, walking while blindfolded in an oblique direction from the original line of sight to the target, and then pointing towards or walking towards the (now unseen) target (Fukusima, Loomis, and Da Silva 1997). The presumed perceived distance is determined based on the original starting point and the intersection of the original line of sight with the final indicated direction (Figure 2). Triangulated walking or pointing can be used to evaluate perception of larger distances than can easily be tested using direct walking and it has a theoretical advantage over direct walking in that it is less likely to involve some specialized visual-action coupling not related to more generally useful distance perception. High accuracy in distance estimation has been observed in visually directed action studies across many studies.

2.3 Prior studies of distance judgments in visually immersive environments

In the last few years, a number of research groups have addressed the issue of space perception in visually immersive environments. This work has been motivated both by a desire to explore new techniques for probing human vision (Loomis, Blascovich, and Beall 1999) and for quantifying operator performance in virtual environments (Lampton, McDonald, and Singer 1995). Table 1 summarizes the results of three previous

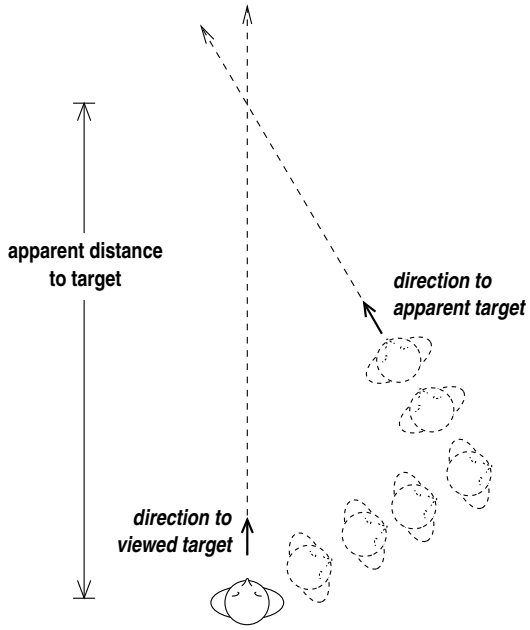

Figure 2: Triangulated walking task: Subjects start walking in an oblique direction from the direction of a previously viewed target. On directions from the experimenter, they turn and take several steps towards where they perceived the previously viewed target to be.

study	*distances*	*real*	*CG*	*task*
Witmer and Sadowski (1998)	4.6m – 32m	92%	85%	treadmill walking
Knapp (1999)	5m – 15m	100%	42%	triangulated walking
Willemsen and Gooch (2002)	2m – 5m	100%	81%	direct walking
Conditions 1 and 2, this study	5m – 15m	95%	44%	triangulated walking

Table 1: Distance judgments based on viewing computer graphics (*CG*) generated imagery using visually immersive displays are compressed from comparable judgments based on viewing real-world environments. The percentages indicate the overall ratio of perceived distance to actual distance.

studies of absolute egocentric distance judgments over action space in visually immersive environments, along with some of the results discussed further in Section 4. In each of these studies, some form of directed action was used to evaluate distance judgments in both real and computer generated scenes. All involved indoor environments and targets situated on a level ground plane. The first study used a Fakespace Labs BOOM2C display with 1280 by 1024 resolution. The second study used a Virtual Research FS5 HMD with 800 by 480 resolution. The final two studies used an nVision HiRes HMD with 1280 by 1024 resolution.

One of the striking results from these studies is that distance judgments in virtual environments were consistently underestimated compared with judgments in the real world. Most of the results in the *CG* column of of Table 1 were based on imagery comparable to that shown in Figure 3b. One potential explanation for this compression of virtual space is that the quality of the imagery is too poor to generate an effective familiar size effect. The experiment described below is aimed at exploring this conjecture.

(a) Section of panorama image, showing target.

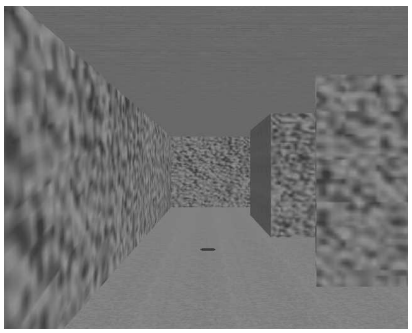

(b) Example of low-quality computer graphics image, showing target. The viewpoint is the same as for Figure 3a.

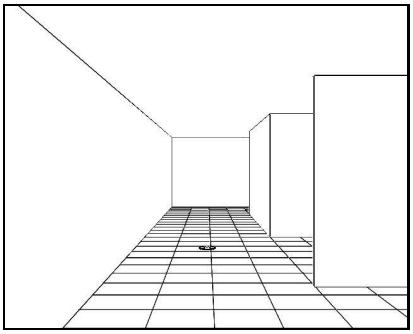

(c) Example of wireframe computer graphics image, showing target. The viewpoint is the same as for Figure 3a.

Figure 3: Sample imagery for conditions 2, 3, and 4.

3 Method

In order to investigate the degree to which image quality affects egocentric distance judgments in virtual environments, we compared distance judgments in the real world with distance judgments in virtual environments utilizing three very distinct styles of graphical rendering: 360° high-resolution panoramic images, intentionally low-quality texture-mapped computer graphics, and wireframe renderings (Figure 3). We probed subjects' perceptions of distance using a directed action task in which subjects indirectly walked without vision towards a previously viewed target. A between subjects design was used in which a given subject viewed trials at three different distances in one of four different environments. Care was taken to make the tasks similar in both the real and virtual environments and to make the scale and general layout of all four environments equivalent.

3.1 Participants

Forty eight college age students participated in this study, with six male and six female subjects in each condition. Subjects either received course credit for participating or were volunteers. All subjects were given a stereogram eye test and had normal or corrected to normal vision. Interpupilar distances ranged from 5.1cm to 7.7cm, with an average of 6.19cm.

3.2 Materials

In the real world condition, subjects viewed a foam-core circular disk approximately 37cm in diameter and placed on the ground at distances of 5m, 10m, and 15m. The experiment was performed in the lobby of an engineering classroom building. Subject positions relevant to computing apparent distance (Figure 2) were determined by measuring foot positions on the floor.

In the three virtual world conditions, imagery was presented using a nVision Datavisor HiRes HMD with interlaced 1280x1024 resolution, full field-sequential color, and a 42° horizontal field of view. The angular resolution of the HMD was on the order of 2 arc minutes per pixel. The nVision has user-adjustable focus. The display was configured with 100% stereo overlap between the two eyes. Head tracking was done using an InterSense IS600 Mark 2 tracker. This tracker uses a mix of inertial, gravitational, and acoustic technologies to provide state-of-the art accuracy and latency. Only tracker rotation was used to update the viewpoint. While translational tracker positions were recorded, the results reported in Section 4 were based on measured foot position on the floor in order to be consistent with the real-world condition. All computer

generated environments were rendered on an SGI Onyx2 R12000 with two IR2 rendering pipelines. One rendering pipeline was used for each eye to provide stereopsis.

Multiple sets of panorama images were produced for different target distances and eye heights, based on photographs acquired by swinging a camera around a fixed axis, located in the same position as the viewpoint for the real-world condition. Targets were placed in the same locations as for the real-world condition. To provide stereo viewing, two sets of images were taken for each panorama, with the camera offset laterally $\pm$3.25cm from the axis of rotation. The two sets of photographs were digitized onto a PhotoCD and then mosaiced into two cylindrical images using the Panorama Factory software package. Each cylindrical image was textured mapped onto a set of polygons forming a cylindrical configuration, providing the ability to generate views over a 360° by 100° portion of the optical sphere. Rendering update rates were no less than 40 frames per second in each eye. The result was a compelling sense of being able to look around in the virtual environment, though no motion parallax was available and the stereo geometry was slightly incorrect. To control for subjects' eye height, multiple panorama image pairs were produced for eye heights spaced at 5cm intervals and then the set nearest to a given subject's eye height was used for that subject's trials. Practical concerns relating to the manner in which the original images were captured precluded a similar control for interpupilar distance.

The second virtual environment condition involved a computer graphics rendering of the same classroom building lobby. The scale of the model was the same as the actual building lobby, but the geometric detail was intentionally kept quite simple. Stereotypical tiled texture maps were used. Simple point-source lighting was used with no shadows or other global illumination effects. Targets were rendered as a red disk, with the size and position corresponding to what was used for the real-world condition. Rendering update rates were no less than 30 frames per second in each eye.

The wireframe virtual environment condition was constructed by rendering feature edges of the model used in the second virtual environment condition. Our software used an OpenGL silhouette drawing algorithm (Raskar and Cohen 1999) to generate the feature edges. The frame rates for this environment were no less than 40 frames per second. The wireframe rendering produced scenes that resemble black on white sketches of the classroom building lobby. The target was rendered with feature edges as well with size and position the same for the previous conditions.

For both the texture mapped and wireframe computer graphics conditions, eye heights were rendered based on the subjects' actual eye heights. Interpupilar distances for stereo rendering were fixed at 6.5cm, consistent with the panorama images.

3.3 Procedure

Subjects were first provided with written instructions that described the triangulated walking task and then given a demonstration of the task in a space both smaller and different from the actual experiment spatial layout. For all conditions, both real and virtual, subjects were instructed to obtain a good image of the target and their local surroundings while first facing the target. Subjects were told that a "good image" is obtained if after closing their eyes, they would still be able to "see" the environment, and most importantly, the target. Subjects were allowed to rotate their head about their neck but were instructed not to move their head from side to side or back and forth. This was done to minimize motion parallax cues in the real world condition so as to make it as comparable as possible to the virtual world conditions.

Once a good image was achieved, subjects were instructed to physically turn their bodies approximately 70° to the right to face a junction of two walls in the environment. After subjects turned, they were instructed to turn their head back toward the target to obtain a final view and reaffirm their mental image of the environment. Then, subjects either blindfolded themselves (real-world condition) or closed their eyes while the HMD screen was cleared to black (virtual-world conditions). Subjects were then directed to walk purposefully and decisively in the direction their body was facing. After walking approximately 2.5m, an

(a) Real-world condition

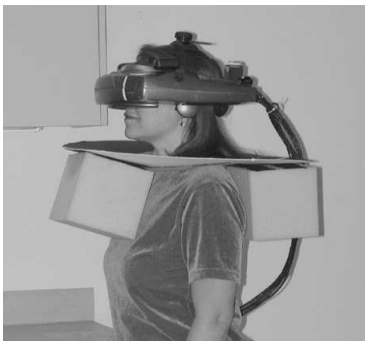

(b) Virtual-world conditions

Figure 4: Viewing collar to hide viewer's body and floor close to standing position.

experimenter would give the verbal command "TURN", signaling the subject to turn towards the target and walk a step or two in it's direction until they felt they were facing the target. Subjects were instructed to perform this turn as if they were turning a corner in a real hallway to make the movement as natural as possible. At this point, the subject's position was marked and they were directed to "Take two steps in the direction of the target". Again, the subject's position was marked and recorded. The subject was then led without vision to the starting location by an experimenter. In all conditions, the apparent location of the target was assumed to lie at the intersection of the line of sight to the (visible) target from the initial vantage point and a line corresponding to the subject's trajectory on the final walk towards the presumed target location (Figure 2).

The user's own body is seldom rendered in immersive visual environments. This is a potential problem when investigating absolute egocentric distance judgments, since eye height is an important scaling factor which could conceivably be affected by looking down at the user's feet and the floor on which she or he is standing. Rendering avatar feet may not be sufficient, since it is difficult to achieve a high degree of realism. We controlled for this potential problem by having users wear a circular collar in both the real world and virtual world conditions (Figure 4). The collar had the effect of occluding users' view of the floor out to about 2m, hiding the area around their feet in all four tested conditions.

Prior to the experiment trials, subjects practiced blind walking for 5 minutes. During this practice, subjects walked blindfolded in a hallway and responded to verbal commands to start and stop walking. The training is helpful in building trust between the experimenter and the subject (Rieser 1998), but more importantly accustoms the subject to walking blind. During both the training session and the actual experiment, subjects wore headphones fed by an external microphone to help limit the effects of sound localization in the environment. A remote microphone worn by the experimenter allowed subjects to hear instructions. After the training session, subjects were led, still blindfolded, to either our laboratory or the real lobby. This last step was performed to help ensure that the subject's movement during the experiment would not be inhibited by *a priori* knowledge of the location of the walls in our lab. The sound masking headphones remained on during this time. For the virtual world conditions, when subjects arrived in the laboratory the HMD was placed on their head while their eyes remained closed. Once on, subjects were allowed to open their eyes and adjust the fit and focus of the HMD, after which the orientation of the virtual world was aligned with the the natural resting position of the HMD on the subject.

4 Results

Figures 5–8 show the average judgments for each of the four conditions: real world, high-quality panorama images, low-quality texture mapped computer graphics, and wireframe. Error bars indicate one standard error above and below the mean. The intersection computation used to compute apparent distance (Figure 2) results in asymmetric variability around the mean, since a turn of δ° too far to the right produce an overshoot

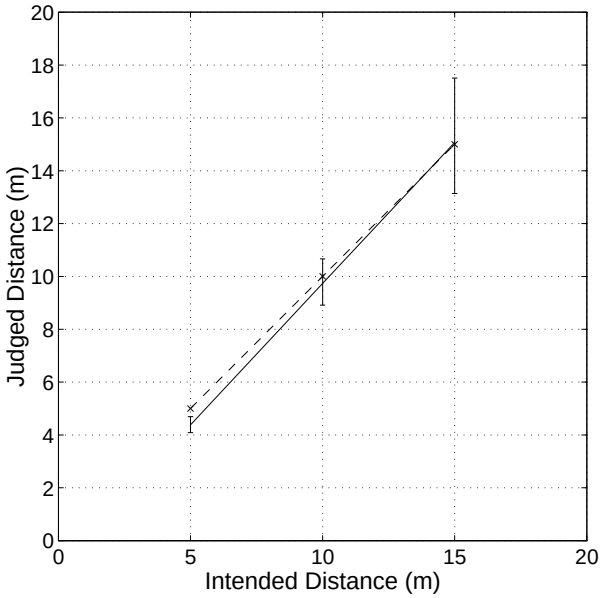

Figure 5: Distance judgments: Real world.

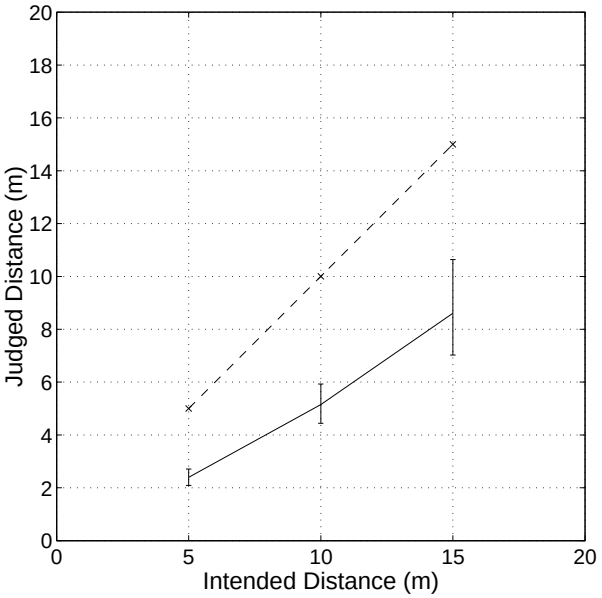

Figure 6: Distance judgments: Panorama images.

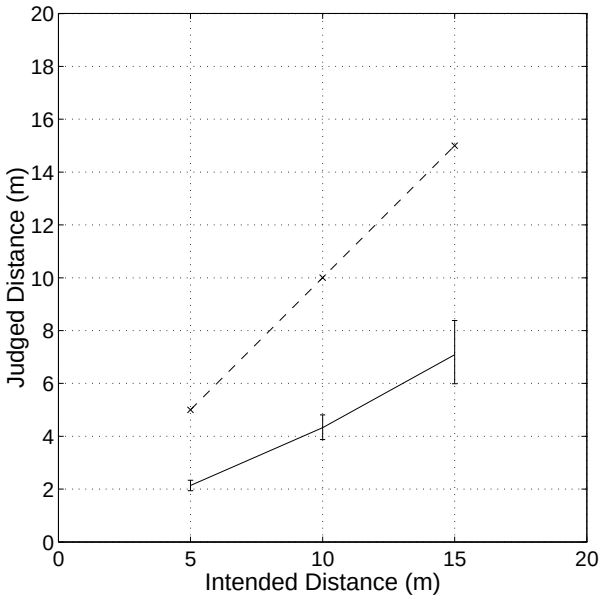

Figure 7: Distance judgments: Low quality computer graphics.

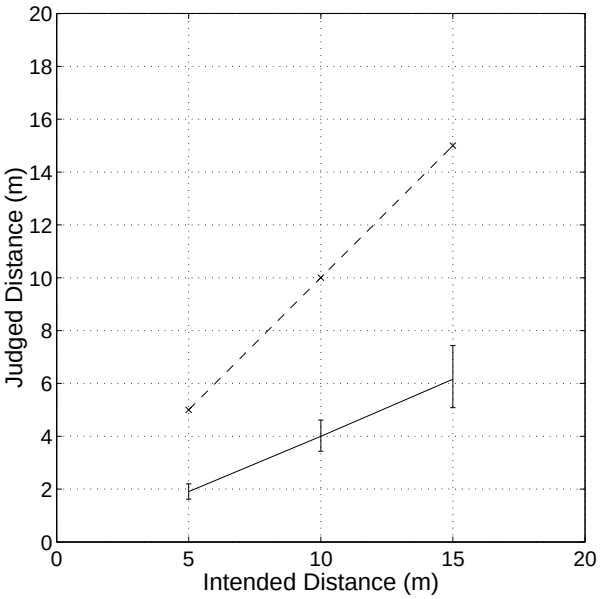

Figure 8: Distance judgments: Wireframe graphics.

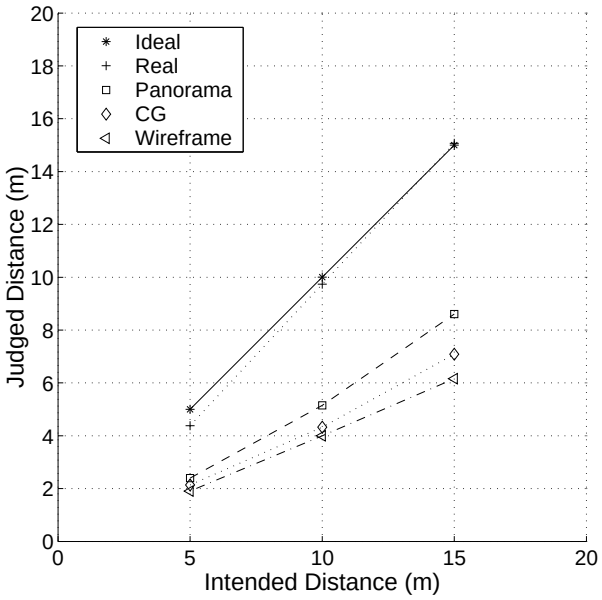

Figure 9: Distance judgments: Comparison of all conditions.

in distance larger than the undershoot in distance produced by a turn of δ° too far to the left. An arctangent transform was applied to the data to reduce this effect. Averages, error estimates, and measures of statistical significance were calculated in the transform space. The inverse transform was then applied to the calculated averages and errors in order to allow presentation of the final results in terms of judged distance.

Figure 9 allows easy comparisons between results for all four conditions. The experiment confirmed previous studies showing that for standing observers viewing ground level targets in action-space range, distance judgments in the real world were near veridical while distance judgments based on computer graphics were significantly compressed. The surprising result was that the amount of compression was nearly the same for all three graphical displays. That is, distance judgments were almost unaffected by the quality of the imagery presented to subjects.

A 4 (environment) x 3 (distance) x 2 (sex) repeated measures ANOVA with distance as a within-subject variable and environment and sex as between-subject variables was performed on the transformed, average distance judgments and indicated a significant effect of environment, $F(3, 40) = 10.77$, $p < .001$. Collapsed across distance, Scheffe post hoc comparisons showed that distance judgments in the real-world were greater than those given in each of the other environments ($p < .01$) and that performance in the other three environments did not differ ($p > .48$ for all comparisons). Although the means at 10m or 15m suggest differences between the virtual conditions, post-hoc univariate ANOVAs (with three environmental conditions) at each distance indicated that these differences were negligible ($p > .4$ for the effect of environment). The ANOVA also indicated an effect of distance, $F(2, 80) = 183.84$, $p < .001$. Judged distance increased as a function of physical distance for all environments. In all, the analyses demonstrated that perceived distance was significantly more accurate in the real world compared to the virtual environments and that distance judgments in the virtual environments did not vary much from each other.

5 Discussion

The results presented above are a strong indicator that compressed absolute egocentric distance judgments in visually immersive environments are not caused by a lack of realistic graphics rendering. The phenomenal

experience of realism in the panoramic environment is best expressed by the comments of several subjects. When looking into a glass window in the rendered display, they commented, "why can't I see my reflection in the glass?". Despite this subjective experience, judgments based on wireframe renderings were as good as judgments based on actual images presented with the same display system. In all virtual environments there was a large compression of egocentric distance. As a result, absolute egocentric distance judgments in virtual environments are not likely to be aided by photorealistic improvements in computer graphics, such as better texturing and illumination. From a theoretical standpoint, this suggests that familiar size may be a relatively minor contributor to the sort of distance judgments which were investigated, though it is important to note that all four conditions involved hallway-like scaling and geometry. The similarity between judged distances to targets on the floor in the three types of virtual displays is consistent with the hypothesis that the declination of visual angle to targets dominates distance egocentric perception (Ooi, Wu, and He 2001). However, this does not explain the large differences observed between distance judgments in the real and virtual conditions.

The present experiment used a methodology that involved a stationary viewer and an action-based judgment task to address specific questions about judgments of distance in visually immersive environments. Our intent was to determine whether observers would judge egocentric distance in the simulated environment in a similar manner as in the real-world without the experience of active exploration. Thus, we restricted the observer's movement while viewing the environments. Previous visual-motor adaptation studies (Rieser, Pick, Ashmead, and Garing 1995; Pick, Rieser, Wagner, and Garing 1999) have demonstrated that active observers will quickly adapt to a new mapping between visual input and their own movements, leading to the result of modified motor output that corresponds to the visual world (recalibration). We might predict that allowing active exploration of the virtual environments would lead to a similar adaptation and recalibration effect so that observers would learn to walk and turn an accurate distance to virtual targets. While this prediction addresses an important question, it is a different question than the one presently asked in this paper. Our goal was to test whether egocentric distance judgments would replicate the accurate performance demonstrated in the real-world, not whether these judgments could become accurate after interacting within a compressed perception of the world. Future studies should consider both the extent of veridical (real-world) perception in visually immersive environments, as well as the role of actions in making immersive environments useful despite a potential lack of veridical perception.

What might explain the compression of absolute egocentric distance judgments, if not image quality? We suggest several possibilities, but no solid evidence supporting any of the potential explanations has yet been published. While the realism of the panorama images used in this study far exceeded any of the computer graphics employed in distance judgment experiments by other investigators, resolution and apparent sharpness were still limited compared to natural viewing of the real world. This may have influenced a familiar size effect or may have degraded the sense of presence while wearing the HMD. Dixon, Wraga, Proffitt, and Williams (2000) found that visual immersion was needed for eye height to appropriately scale linear perspective cues. Perhaps a full of sense of presence, not only visual immersion, is needed for distance judgments to be comparable to what is seen in the real world. Limited field of view is often suggested as a cause of distorted spatial vision in HMDs, but Knapp (1999) found that limiting FOV did not affect real-world egocentric distance judgments, at least if the observer was free to move his or her head to visually explore the environment. Motion parallax was not present in our virtual display conditions, but motion parallax appears to be a rather weak absolute distance cue (Beall, Loomis, Philbeck, and Fikes 1995). In addition, subjects performed veridically in our real-word condition with at most very limited translational head motion. Focus and stereo convergence are not well controlled in HMDs (Rolland, Gibson, and Arierly 1995; Wann, Rushton, and Mon-Williams 1995), and incorrect accommodation cues are known to affect distance judgments (Andersen, Saidpour, and Braunstein 1998; Bingham, Bradley, Bailey, and Vinner 2001). It seems unlikely, however, that accommodation and convergence would have an effect this large at the distances we were investigating. Finally, there may be some sort of ergonomic effect associated with wearing

an HMD (Lackner and Dizio 1989).

Future research that manipulates factors other than the image quality, such as FOV, stereo, and physical effects of the HMD, is needed to begin to answer these questions. A sense of presence is more difficult to define and manipulate, but is likely to be an important component in accurate distance perception in virtual environments.

Acknowledgment

Support for this research was provided by National Science Foundation grants 9623614 and 0080999.

References

Andersen, G. J., A. Saidpour, and M. L. Braunstein (1998). Effects of collimation on perceived layout in 3-D scenes. *Perception 27*, 1305–1315.

Beall, A. C., J. M. Loomis, J. M. Philbeck, and T. J. Fikes (1995). Absolute motion parallax weakly determines visual scale in real and virtual environments. In *Proc. of the SPIE – The international Society for Optical Engineering*, Volume 2411, pp. 288–297.

Bingham, G. P., A. Bradley, M. Bailey, and R. Vinner (2001). Accomodation, occlusion, and disparity matching are used to guide reaching: A comparison of actual versus virtual environments. *Journal of Experimental Psychology: Human Perception and Performance 27*, 1314–1334.

Cutting, J. E. and P. M. Vishton (1995). Perceiving layout and knowing distance: The integration, relative potency and contextual use of different information about depth. In W. Epstein and S. Rogers (Eds.), *Perception of Space and Motion*, pp. 69–117. New York: Academic Press.

Dixon, M. W., M. Wraga, D. R. Proffitt, and G. C. 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.

Ellis, S. R. and B. M. Menges (1997). Judgments of the distance to nearby virtual objects: Interaction of viewing conditions and accommodative demand. *Presence: Teleoperators and Virtual Environments 6*, 452.

Fukusima, S. S., J. M. Loomis, and J. A. Da Silva (1997). Visual perception of egocentric distance as assessed by triangulation. *Journal of Experimental Psychology: Human Perception and Performance 23*(1), 86–100.

Knapp, J. M. (1999). *The Visual Perception of Egocentric Distance in Virtual Environments*. Ph. D. thesis, University of California at Santa Barbara.

Lackner, J. R. and P. Dizio (1989). Alterned sensor-motor control of the head as an etiological factor in space-motion sickness. *Perceptual and Motor Skills 68*, 784–786.

Lampton, D. R., D. P. McDonald, and M. Singer (1995). Distance estimation in virtual environments. In *Proc. Human Factors and Ergonomics Society*, pp. 1268–1272.

Loomis, J. M., J. J. Blascovich, and A. C. Beall (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., J. A. Da Silva, N. Fujita, and S. S. Fukusima (1992). Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance 18*, 906–921.

Loomis, J. M., J. A. Da Silva, J. W. Philbeck, and S. S. Fukusima (1996). Visual perception of location and distance. *Current Directions in Psychological Science 5*, 72–77.

Loomis, J. M. and J. M. Knapp (in press). Visual perception of egocentric distance in real and virtual environments. In L. Hettinger and M. Haas (Eds.), *Virtual and Adaptive Environments*. Hillsdale, NJ: Erlbaum.

Ooi, T. L., B. Wu, and Z. J. He (2001, November). Distance determination by the angular declination below the horizon. *Nature 414*, 197–200.

Pick, Jr., H. L., J. J. Rieser, D. Wagner, and A. E. Garing (1999). The recalibration of rotational locomotion. *Journal of Experimental Psychology: Human Perception and Performance 25*(5), 1179–1188.

Raskar, R. and M. Cohen (1999). Image precision silhouette edges. In *Proc. ACM Symposium on Interactive 3D Graphics*, pp. 135–140.

Rieser, J. J. (1998). Dynamnic spatial orientation and the coupling of representation and action. In R. G. Golledge (Ed.), *Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes*, pp. 168–190. Baltimore, MD: Johns Hopkins Univesity Press.

Rieser, J. J., D. H. Ashmead, C. R. Talor, and G. A. Youngquist (1990). Visual perception and the guidance of locomotion without vision to previously seen targets. *Perception 19*, 675–689.

Rieser, J. J., H. L. Pick, Jr., D. Ashmead, and A. Garing (1995). Calibration of human locomotion and models of perceptual-motor organization. *Journal of Experimental Psychology: Human Perception and Performance 21*, 480–497.

Rolland, J. P., W. Gibson, and D. Arierly (1995). Towards quantifying depth and sizer perception as a function of viewing distance. *Presence: Teleoperators and Virtual Environments 4*, 24–49.

Thomson, J. A. (1983). Is continuous visual monitorying necessary in visually guided locomotion? *Journal of Experimental Psychology: Human Perception and Performance 9*(3), 427–443.

Wann, J. P., S. Rushton, and M. Mon-Williams (1995). Natural problems for stereoscopic depth perception in virtual environments. *Vision Research 35*(19), 2731–2736.

Willemsen, P. and A. Gooch (2002). Perceived egocentric distances in real, image-based, and traditional virtual environments. In *Proc. IEEE Virtual Reality*, Orlando, FL.

Witmer, B. and W. Sadowski, Jr. (1998). Nonvisually guided locomotion to a previously viewed target in real and virtual environments. *Human Factors 40*, 478–488.

Wraga, M. (1999). Using eye height in different postures to scale the heights of objects. *Journal of Experimental Psychology: Human Perception and Performance 25*(2), 518–530.