

Distance Perception and the Visual Horizon in Head-Mounted Displays

ROSS MESSING
University of Rochester
and
FRANK H. DURGIN
Swarthmore College

Can distance perception be studied using virtual reality (VR) if distances are systematically underestimated in VR head-mounted displays (HMDs)? In an experiment in which a real environment was observed through an HMD, via live video, distances, as measured by visually directed walking, were underestimated even when the perceived environment was known to be real and present. However, the underestimation was linear, which means that higher-order space perception effects might be preserved in VR. This is illustrated in a second experiment, in which the visual horizon was artificially manipulated in a simulated outdoor field presented in immersive VR. As predicted by the claim that angle of declination from the horizon may serve as a strong cue to distance, lowering the horizon line produced “expansive” judgments of distance (power function exponents greater than one) both in verbal and in motor estimates.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Virtual Reality*

General Terms: Experimentation, Human Factors

Additional Key Words and Phrases: Virtual reality (VR), distance perception, space perception, head-mounted displays (HMD)

1. INTRODUCTION

Immersive virtual reality (VR) seems like an ideal tool for studying perception because it is so easy to manipulate aspects of the world that might otherwise be difficult to control. In particular, the study of factors affecting space and distance perception is facilitated by being able to conduct experiments involving subtle manipulations of a large-scale visual environment, while preserving a sense of immersion. However, a number of researchers have reported a pervasive underestimation of perceived distance in VR using HMDs [Loomis and Knapp 2003; Rowland 1999; Durgin et al. 2002; Creem-Regehr et al. 2005]. The basis for these distance errors are not understood, but they are not corrected by using photorealistic rendering [Thompson et al. 2004]. Here we will show that although the underestimation of distance persists even when properly scaled near-live video is presented through the head-mounted display (HMD), the compression is linear over the range of distances tested. Studies of space perception

Direct correspondence to: Frank Durgin, Department of Psychology, Swarthmore College, 500 College Avenue, Swarthmore, Pennsylvania. 19081.

R. M. was supported by the Hans Wallach Fellowship from the Swarthmore College Department of Psychology; This work was supported by a Swarthmore College Faculty Research Grant to F.H.D.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.
© 2005 ACM 1544-3558/05/0700-0234 \$5.00

ACM Transactions on Applied Perception, Vol. 2, No. 3, July 2005, Pages 234–250.

may, therefore, circumvent the problem of underestimation in HMDs by looking at higher-order distortions of space. This insight is applied to a study of the effects of altering the position of the visual horizon to illustrate how higher-order analyses can overcome the problem of linear distance compression.

Although the goals of the two experiments we will present are diverse, they represent two parts of a single argument. In order to study perception using VR HMDs, it is necessary to characterize the kinds of distortion that are likely to occur and determine their likely consequences on the interpretation of results. Having characterized those distortions, we can then illustrate the value of VR as a tool for perception research by directly manipulating the visual horizon in a way that is essentially impossible to do in the real world in real time. The results of this manipulation of the visible horizon are consistent with theoretical predictions that we make based on other recent work implicating the role of the perceived (but unseen) horizon in the perception of distance [Ooi et al. 2001].

In both of the experiments reported here, we analyze not only the linear magnitude of distance underestimation, but also compute the slope of distance judgments in log–log space. This classic technique (equivalent to computing power function exponents for perceived distance [e.g., Teghtsoonian and Teghtsoonian 1969]) can help to elucidate something more about the geometry of the visual space than simply its quantity of compression. A visual space is “compressed” if all distances are underestimated. It is “compressive” if the underestimates for longer distances (expressed as a ratio) are more drastic than those for shorter distances. This is evident if the exponent of the power function is less than 1. A visual space is “expansive” if the exponent is greater than 1. We find that egocentric distance judgments in VR are compressed, but do not find evidence that they are compressive. If the visual horizon is used as a distance cue, lowering the apparent location of the visual horizon (by actually truncating an apparently infinite ground plane), should make distance judgments “expansive.” That is, perceived distance should increase in a manner consistent with an exponent greater than 1. This prediction will be borne out experimentally using HMD VR.

1.1 Underperception of Visual Space in Head-Mounted Displays

Virtual reality seems like an excellent tool for studying perception because it allows experimenters unprecedented control of visual stimuli, while maintaining the measure of ecological validity that comes with an immersive environment [Loomis et al. 1999]. However, virtual reality systems have a number of limitations. Some of the most obvious limitations have to do with the size and quality of the displays. The field of view available from most VR HMDs is much more limited than a natural field of view, for example (typically less than one-third of the horizontal field of view [FOV] of normal vision) and screen resolution is often quite limited as well. Moreover, the geometry of the peripheral parts of the display is typically distorted by the optics, which may, in turn, grossly alter peripheral optic flow during head movements. There are subtler problems, as well, that are particularly relevant to the representation of visual distance. The light from a HMD is collimated to allow it to be focused at close range, but, as a consequence, eye movements produce no parallax, which may defeat a sense of depth. In addition, all distances represented in HMDs have the same accommodative signal, which is determined by the optics of the HMD. Although accommodation is not thought, by itself, to be a strong cue to distance, it is tightly yoked to the convergence system, which is [Foley 1980; Owens 1987]. Finally, displays that are updated quickly and accurately in response to head motions can provide a strong sense of immersion, but there always remain subtle lags that may interact with the very high temporal resolution of the neural systems involved in computing structure from motion. All of these limitations may signal problems for researchers hoping to take advantage of the technology.

In fact, there is a systematic compression of perceived distance in HMDs that has yet to be fully explained [Loomis and Knapp 2003; Rowland 1999; Creem-Regehr et al. 2005]. Loomis and Knapp [2003] thought that this compression might be due to the limited quality of their VR system’s graphics

rendering. Rowland [1999] has hypothesized that using dense, high-quality, natural textures on the ground plane would reduce compression. However, in a direct test of this idea, no significant difference was found in distance compression between very simple line-drawing virtual environments, and photorealistic ones [Thompson et al. 2004]. In addition, Creem-Regehr et al. [2005] 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. Knapp and Loomis [2004] also investigated restricted field of view in binocular viewing, and like Creem-Regehr et al., found no statistically significant effect of restricting FOV. Recent findings by Willemsen et al. [2004] suggest that mechanical aspects of HMDs may be responsible for some distance compression.

The studies we have mentioned have used a variety of measures to assess perceived distance, but the most common task is walking, without visual feedback, to a visually previewed target. Such “visually directed” walking tasks have been shown to produce excellent performance in normal environments [Loomis et al. 1992]. Correspondingly, the underperception of distance in virtual reality can also be demonstrated with a variety of tasks. However, these studies have consistently reported only the amount of compression without analyzing whether the compression varies with distance.

2. EXPERIMENT 1

Distance might be underperceived in HMD’s primarily because of optical or mechanical properties of the HMD itself, or because of the quality and compellingness of the rendering. Thompson et al. [2004] showed that the level of graphical detail in a virtual environment does not significantly affect distance judgments. Although there were limitations in the interactivity of their displays (head translations were discouraged because of limitations in the graphics), these are probably not sufficient to explain the lack of any effect of photorealism. This suggests that the effects are probably due to intrinsic limitations of the medium and that characterizing the effects of the medium on distance perception is a reasonable strategy.

To further investigate the possible role of realism, we used a visually directed walking task to measure distance perception in the real world when viewed, by means of live video, through an HMD. Because our participants knew that they were acting in the real world, this eliminated any cognitive nonrealism that might have remained in photorealistic virtual worlds. Although we had also hoped to improve on the motion-based depth cues available, by allowing translational head movements, as well as head rotations, a substantial lag in our video signal probably rendered motion parallax cues ineffective in our video display as well. Nonetheless, our participants knew that they were acting with respect to a real environment.

Our apparatus was monocular, primarily for simplicity. However, there is reason to believe that stereoscopic displays may be inaccurately rendered anyway in virtual reality [Wann et al. 1995]. In addition, a number of studies have found no difference in distance perception between monocular and binocular viewing of distances greater than 2 m [Creem-Regehr et al. 2005; Philbeck and Loomis 1997].

Although the head-mounted video camera provided realism, it had several limitations. Practical concerns prevented the system from displaying the image that would be seen from the subject’s actual eye location. Instead, the camera was mounted at the side of the HMD. This led to both a slight translation of the subject’s viewpoint and a shifted axis of horizontal rotation. The viewpoint translation was probably of no consequence, given that the targets were several meters away and translation was along an orthogonal axis. The rotational shift could have produced small, but probably unimportant image distortions for head rotations about a vertical or a horizontal sagittal axis. The primary direction of head rotation during target distance assessment was sagittal, where the geometry would have been

correct. The lag in the video delay, which was nearly 500 ms, was probably sufficient to desynchronize motion parallax from actual head motion and, therefore, weaken motion parallax as a cue to distance. In addition to optical limitations, the mounted camera added weight to one side of the HMD, which could also have had an effect on distance perception [Willemsen et al. 2004; Proffitt et al. 2003].

To confirm that any residual errors in distance perception were due to the HMD apparatus, distance perception in the HMD viewing condition was compared to direct viewing of the environment in two conditions. One condition used a monocular goggle that restricted the FOV to the same angle as the HMD condition. This condition allowed for direct viewing of the environment with a retinal image that was matched to the HMD condition as closely as possible. This is important given recent evidence that restricted FOVs may require subjects to use serial scanning to judge distances [Wu et al. 2004]. This condition, in turn, was compared to an unrestricted monocular viewing condition. It has previously been shown that performance in a monocular viewing condition does not differ from unrestricted binocular viewing [Creem-Regehr et al. 2005; Knapp and Loomis 2004].

2.1 Methods

2.1.1 Participants. Subjects were 20 Swarthmore College students who were naive concerning the purpose of the experiment. Some participated as part of an introductory psychology course requirement and others were paid to participate. Seven students were in the HMD condition, 7 in the restricted field of view condition, and 6 in the unrestricted monocular viewing condition.

2.1.2 Apparatus. The experiment was conducted in a carpeted hallway (2.4 m × 22 m) outside the lab. A selection criterion for subjects in the study was that they be unfamiliar with this particular hallway so that they could not be acting on preexisting cognitive maps.

In the main (HMD) condition, subjects wore a Virtual Research Systems V8 HMD (640 × 480, 60 Hz, LCD, 60° diagonal FOV, accommodative distance of 1 m.). A digital video camera (Sharp VL-Z3) with a wide-angle lens that provided the same FOV was mounted on the left side of the HMD. It was connected by IEEE-1394 cable to an Apple G4 running custom software that output video to the HMD. The camera lens was approximately 10 cm to the left of the subject's left eye. Latency from the camera to the display was determined to be at least 433 ms, by videotaping the output on the display and measuring the relative onset of an event and its display.

In the restricted field of view condition (rFOV), subjects wore a monocular goggle with a rectangular cardboard tube that restricted the FOV to approximately the same size as that of the HMD.

In the unrestricted monocular condition (MONO) subjects wore an eyepatch over one eye.

In the rFOV condition, the eyepatch was worn over the right eye. In the other conditions, subjects were allowed to patch the eye of their choosing.

The location to be walked to was marked with a light orange foam cylinder (the target), 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-attenuating ear plugs (NR 31) in order to prevent them from getting auditory feedback about their location in the environment.

2.1.3 Design and Procedure. All subjects performed the visually directed walking task in a texture-rich hallway for distances of 2, 3, 4, 5, 6, and 7 m. Subjects were required to view each target while standing still. There were no restrictions on time or head movement while viewing the target. Subjects were to signal readiness to begin walking and close their eyes. In the HMD condition, the display was turned off at this point and was not turned on until the next trial. After closing their eyes, subjects would walk to the target location and stop. The distance along the hallway between the subject and the target was then measured surreptitiously. One experimenter gave instructions and recorded measurements; a second experimenter placed and removed the targets and measured the distances.



Fig. 1. Video image of the hallway where Experiment 1 took place. The target disk is 7 m from the observation point.

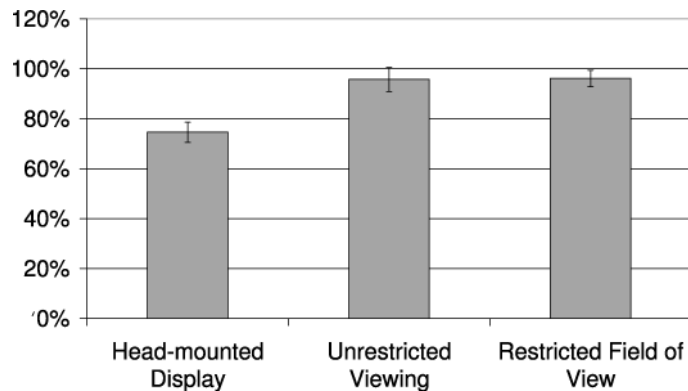


Fig. 2. Mean percentage of distance to the target walked, by viewing condition.

Subjects were told that they were performing a noncognitive 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 would be removed after they closed their eyes to prevent them walking on it.

Subjects performed six trials with targets at distances of 2 to 7 m. Though subjects were not informed of it, the first two trials, at 3 and 6 m (order varied), were treated as practice. The first two experimental trials were (in random order) 2 and 7 m; the final trials were (in random order) 4 and 5 m. Only a small number of trials were used, because we wanted to discourage the formation of task-dependent strategies that may develop with repeated testing.

2.2 Results

One set of data was removed from analysis of the HMD condition, because the subject's average ratio of perceived distance to actual distance was over 5 SDs (excluding that subject) from the other data in that condition.

As can be seen in Figure 2, distances walked in the HMD condition were much shorter than in the other two conditions. In the HMD condition, subjects walked an average of 77% of the distance viewed

(averages were computed in log space), which was significantly different from 100%, $t(5) = 5.87$, $p < 0.01$. The 95% confidence interval for the mean distance walked in the HMD condition (computed in log space) extends from 70 to 86%. The distances walked in the HMD were also reliably less than the distances walked in the rFOV condition ($M = 96\%$), $t[9] = 3.53$, $p < 0.01$, and less than the distances walked in the MONO condition, ($M = 96\%$), $t[11] = 3.83$, $p < 0.01$). Neither of these latter conditions differed reliably from each other, nor did either differ reliably from accurate walking ($p > 0.10$).

We next analyzed the data for evidence of compressiveness. Because the accommodative cues in our HMD specify a 1-m viewing distance, it was possible that the underperception of distance would increase with greater distance. To test this, we fit each subject's data to a power function. If the exponent of these power functions was less than 1, this would represent increasing compression with distance. However, in no condition did the average exponent differ reliably from 1. The average exponent in the HMD condition was 0.95. In the rFOV condition it was 0.98 and in the MONO condition it was 0.93.

There was no reliable difference between any of these values. Evidently, the underestimation of distance in the HMD was by a nearly constant factor.

2.3 Discussion

Consistent with the conclusions of Thompson et al. [2004], we have found reliable distance compression in the HMD condition, even when displaying live video. Consistent with the findings of Creem-Regehr et al. [2005] and Knapp and Loomis [2004], performance in our restricted-viewing condition did not differ from that in our unrestricted viewing condition. More importantly, however, in all three conditions, the exponents of power functions fit to the data are nearly 1. This means that distance compression in our HMD is nearly linear with distance in the range of distances tested.

If level of graphical realism and rFOV are not the cause of distance underestimation in HMDs, other aspects of the system must be responsible. One hypothesis is that the problems lie, in part, with the optics of the system and their resulting distortions in the periphery of the displayed field. Another hypothesis concerns the absence of need for changes in accommodation when scanning different parts of the scene. However, anomalies of accommodation are unlikely to be detectable for distances beyond a few meters and cannot explain the constant underperception of distance we have documented. It is likely that the resolution of our display was inadequate to represent the fine texture in the floor carpet. This would have disrupted integration of ground texture, which Wu, Ooi and He [2004] have argued is an important basis for the accurate evaluation of egocentric distance.

Despite the limitations of virtual reality for providing accurate distance information, the compression of space in our HMD seems to be roughly linear. Immersive VR therefore remains a potentially useful tool for investigating distance perception because of the ease with which a virtual environment can be manipulated. In Experiment 2, we used VR to examine the distance cue known as angular declination from the horizon. Crucially, our experimental predictions concern the changes in relative distances—the so-called *geometry of space*. We suggest that this experiment may serve as a useful model for studying distance perception using virtual environments presented in HMDs, because the near linearity of visual space seems to be preserved in the 2–7 m range we tested.

Ooi et al. [2001] showed that deflecting the direction of gaze by means of prisms produced systematic distortions in the perceived distance of objects consistent with the idea that object position in space is coded by means of angle of regard. Although Ooi et al. refer to this information as angular declination below the horizon, they did not actually manipulate the visual horizon. Rather, they considered the internal sense of “straight-ahead” to be a surrogate for the horizon. They did show that a prism-induced deflection of subjective “straight-ahead” could account for distance judgments following prism adaptation, but this is confounded with the concomitant deflection of all absolute perceived angles, not simply horizon-relative ones. Using VR, the direct manipulation of the visual horizon is quite easy

to accomplish without affecting any of the other absolute geometry in a scene. In our experiment, we will show that the predicted effects of horizon manipulations are obtained in the absence of any of the geometric distortions due to prisms, altered viewing angle, or altered viewing position. Crucially, although some expansion of visual space is predicted by the manipulation of the horizon (e.g., lowering it), the rate of expansion increases with distance. This means that effects of manipulating the horizon should turn up in our experiments in the form of exponents greater than 1.

We note that a previous study using VR manipulated the virtual height of the observer in order to examine the effects of eye level on the perceived size of objects [Dixon et al. 2000]. This research assumed that size was perceived independent of distance using eye-height scaling [Wraga 1999a, 1999b]. In contrast with their virtual eye-height manipulation (which does not move the horizon, but alters all other absolute angles in the scene), we maintain normal eye height and ground-plane geometry in near space for our observers. That is, our manipulation of the visual horizon position differs substantively from this prior work. By using an object with a fairly small vertical extent, and asking specifically for distance judgments along a richly textured ground plane, we will maintain the normal geometry appropriate to the observer in all respects, except for the location of the visual horizon.

3. EXPERIMENT 2

The angle of declination from the horizon line—even in the absence of an explicit visual horizon—is a strong distance cue. This has been shown by adjusting the perceived vertical angle of gaze, either by wearing vertically deflecting prisms or by the aftereffect of adaptation to wearing such prisms [Ooi et al. 2001]. However, these experiments were carried out in a relatively impoverished visual environment. More recently, texture gradients in ground texture have been shown to play an important role in distance perception for visually directed walking [He et al. 2004]. We sought to address the question of whether manipulating the visual horizon directly in a virtual environment would produce predicted changes in perceived distance even in the presence of a richly textured ground plane. If a textured ground plane is the most important cue to distance, what role should the horizon have?

Although we did not collect verbal estimates in Experiment 1, we further sought to compare measures based on visually directed walking with those derived from verbal estimates of distance in this experiment. The comparison of verbal and “motor” estimates of distance is of particular interest because of the common view, promoted by Goodale and Milner [1992], that there are separate visual representations available to conscious awareness and to motor action. It is sometimes argued that explicit conscious representations of distance (such as verbal estimates) show systematic distortions that are not present when motor tasks are used. However, Loomis et al. [1992] have argued that there is a single egocentric representation of space and that apparent task dissociations emerge when exocentric judgments are contrasted with egocentric ones. The advantage of egocentric distance tasks is supported by the observation that the sequential evaluation of surface textures along the ground plane will only lead to accurate performance when evaluated from the feet out to the object rather than from the object back to the feet [Wu et al. 2004].

If the verbal evaluation of distance is often made to depend primarily on exocentric judgments (comparisons of distant intervals of space—such as enforced by the measurement techniques of Gilinsky [1951], for example), then we should expect that verbal estimates would tend to show different kinds of scaling properties than do motor estimates. In particular, it might be expected that power functions fit to verbal estimates would have exponents slightly less than 1 [Teghtsoonian and Teghtsoonian 1970], whereas motor exponents are expected to be near 1 insofar as motor space is veridical. On the other hand, if both motor and verbal judgments are based on egocentric evaluations, we would expect both to have exponents that do not differ reliably from 1 (as was true for the motor task in all conditions of Experiment 1).

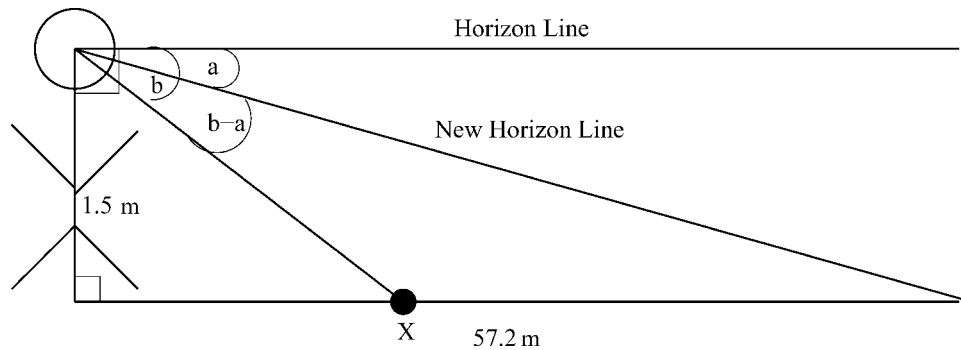


Fig. 3. This figure depicts the effects of lowering the horizon line by 1.5° .

Conversely, artificially lowering the horizon by a fixed angle can be expected to have the effect of producing accelerating power functions, with exponents greater than 1. This is because lowering the horizon line by as little as 1.5° ought to produce greater relative overestimations of distance at greater simulated distances. This is illustrated in Figure 3. In order to lower the apparent visual horizon by a fixed angle, a , assuming an observer eye height of 1.5 m, rendering of the world must stop at $1.5 \text{ m} * \tan(90^\circ - a)$ 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. Note that if the observer's eye height is 1.5 m, object X will now appear to be $\tan(90 - [b - a]) * 1.5 \text{ m}$ from the observer, based on its angular declination.

Lowering just the visual horizon in the real world is nontrivial, but in VR, it can be accomplished by having the ground plane terminate at some fixed distance from the observer. We picked a distance (57.28 m) that would produce a 1.5° drop in the apparent horizon for an observer with an eye height of 1.5 m. Because the ground texture became a fairly uniform green by about 20 m from the observer, the truncated ground plane appeared to be an infinite plane against the textureless sky. That is, the visual horizon it created with the sky, although detectably lower than appropriate for those aware of the manipulation, looked like a horizon, not a cliff. For the range of distances we used (3.0–7.5 m), we computed the expected perceived distances (see Figure 4), and calculated the exponent of the power function fit to those points. This exponent was 1.08.

If the angle of declination from this visual horizon influences the perception of distance, we would predict that distance judgments might demonstrate a similar exponent. To the extent that verbal and motor estimates of distance derive from a single egocentric representation, we further predicted that similar exponents would be found for both the verbal and the motor task. We should emphasize, however, that the task demands of walking are so different from those of producing a verbal distance estimate that we did imagine that task differences could arise. For example, verbal judgments might make subjects more likely to look at the whole scene, while a walking task might make subjects pay more attention to the ground texture between them and the target. If task strategies differed in this way, lowering the horizon might affect verbal judgments more.

3.1 Methods

3.1.1 Participants. Subjects were 20 Swarthmore College undergraduates who were naive as to the purpose of the experiment, participating 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 were female.

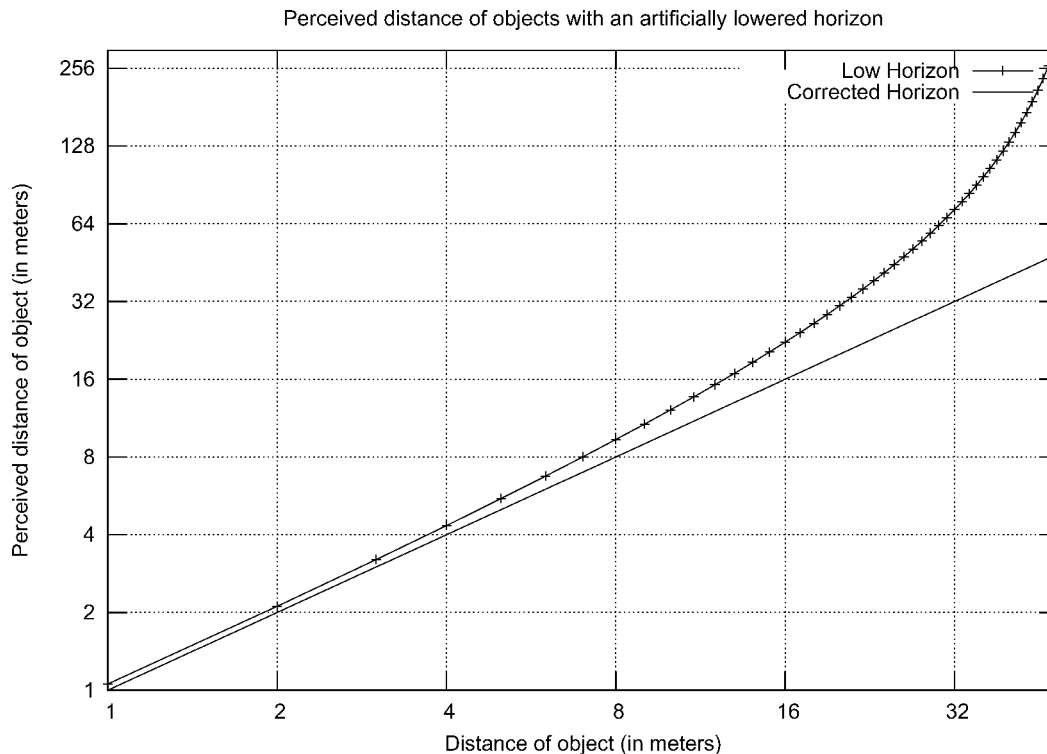


Fig. 4. The effect on distance perception of an artificially 1.5° low horizon created by cutting the rendering plane at about 57.28 m.

3.1.2 Apparatus. Subjects wore a Virtual Research Systems V8 HMD with a 3rd Tech HiBall head tracker. Subjects also wore hearing attenuators. The virtual environments were rendered in stereo at 60 Hz by two G4 computers, synchronized to within 1 ms by the arrival of head-tracker data. The display was 60 Hz with a measured lag of approximately 43 ms.

3.1.3 Displays. Subjects were shown an outdoor virtual environment with a grass field. The field texture was simulated by merging two noise patterns filtered at different spatial scales to provide multiple scales of detail corresponding roughly to individual blades of grass and to largish clumps of turf. Because this ground texture was tailored to our displays, the texture was resolvable to about 20 m from the observer. The environment included a red circle on the ground, 1 ft (0.305 m) in diameter, where subjects were supposed to stand, and an orange cone (0.25-m tall) that served as the target in both conditions. The retinal size of the target varied with its distance, providing strong geometric information about its true distance. No shadows were rendered in the scene, but the absence of this contact cue should not have been conspicuous given the diffuse lighting and distances. Moreover, because the target appeared to be a physical sport cone (rather than an abstract geometric object) physical contact with the ground was implied.

Two variations of the outdoor environment were used. In one, the visual horizon line was made to appear 1.5° too low, as in Figure 5. In the other, the horizon line was correct, as in Figure 6. To accomplish the correct horizon line, we actually added a massive green cone to the world, with an outer radius of 56 m. This cone was always centered on the observer and its upper edge was dynamically set equal to the current eye height of the observer—as is appropriate for a true horizon line. The color of this green




Fig. 5. The outdoor virtual environment with a 1.5° low horizon and a target 7 m from the observer.




Fig. 6. The outdoor virtual environment with a correct horizon and a target 7 m from the observer.

conic wall was selected so that it made no visible edge where it cut through the ground plane. As noted above, the texture merged to solid green about 20 m from the observer, so no discontinuity in texture was evident.

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 themselves (in the real hallway) between trials. This environment is shown in Figure 7.

3.1.4 Design and Procedure. Subjects were randomly preassigned to one of two response conditions, either verbal estimate, or visually directed action. In both conditions there were three practice trials, at distances of 5, 10, and 1 m. All of the practice trials had the horizon in the correct location. During the experimental trials, horizon position was varied randomly from trial to trial. (The screen went blank between trials to mask the change in horizon.) Actual distances shown were based on taking either 3, 4, 5, 6, or 7 m and then adding a random increment of up to 0.5 m in order to avoid repeating distances. These five distance categories were tested with each horizon type for a total of ten randomly ordered trials per block. Each subject completed two full blocks of trials, for a total of 20 estimates (either verbally or by walking). Separate power-function exponents were then computed for each visual horizon location and for each subject, by computing the slope of the least-squares fit to the data in log-log space.




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

3.1.5 Instructions. The instructions given to subjects in the two conditions were kept as similar as possible. In the walking conditions, subjects were told 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. They were also told that specific target distances would not be repeated (due to the added random factor). They 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. The subjects were positioned over a red virtual circle and instructed that moving away from that circle would blank the display. The 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 told 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 told not to count increments of known distances to the target. This was done in order to maintain similar instructions for the two types of task. Subjects were positioned on top of a red circle and informed accurately that it was 1 ft in diameter, although 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. They did not practice walking with closed eyes.

3.2 Results

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 female) and response (verbal or motor) 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 ($M = 0.983$), $F[1, 16] = 8.71$, $p < 0.01$. This is illustrated in Figure 8. This exponent value is quite close to that we predicted based on consideration of changes in apparent angular declination. There was no interaction with sex or of task, suggesting that the horizon manipulation affected both tasks equally.

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, although there was a trend for visually directed action to produce greater estimates of distance than verbal estimation. This is illustrated in Figure 9. Note that the average ratio of walked to actual distance

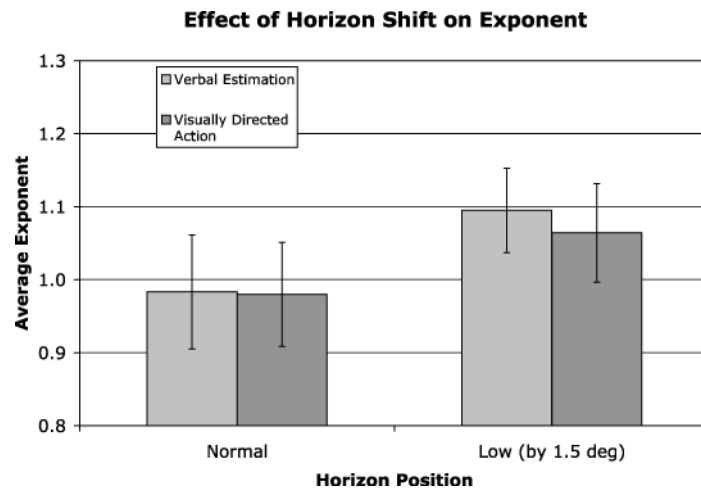


Fig. 8. Average exponents of power functions modeling subject performance in verbal estimation and visually directed action in an outdoor virtual environment with a correct or artificially low horizon line.

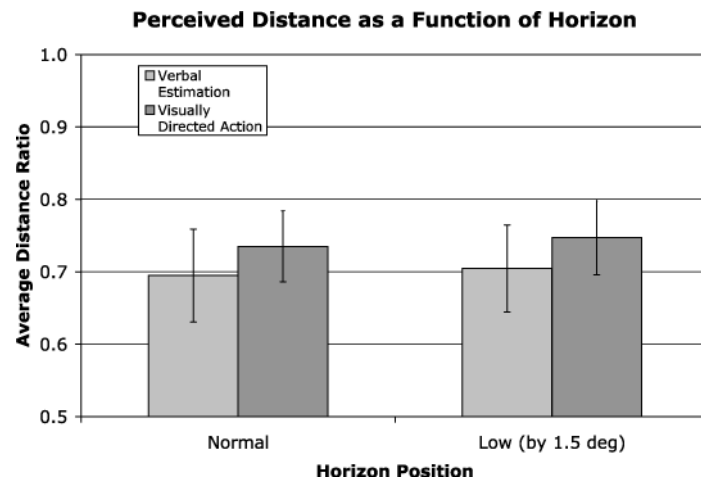


Fig. 9. Average proportion of distance walked to distance viewed in verbal estimation and visually directed action in an outdoor virtual environment with a correct or artificially low horizon line.

in the motor task with the correct horizon line (72%) was not reliably different than that in the video condition of Experiment 1 (77%), $t(14) < 1$, ns.

3.3 Discussion

As predicted, lowering the visual horizon increased the exponent of the fit power functions, consistent with the idea that distance perception in our study was strongly impacted by the apparent position of the horizon. Because we were able to manipulate the apparent visual horizon directly without changing the absolute angle of regard to the targets themselves, we are able to state that the effects we observe are truly due to angle of declination from an explicit visual horizon rather than to perceived angle of regard. Whereas Ooi et al.'s [2001] findings were interpreted in terms of the perceived location of the

horizon (in the absence of a visual horizon), ours can be interpreted with respect to the visual horizon itself. Of course, our walking results are convergent with theirs.

Unlike Dixon et al. [2000], who were studying eye-height scaling of apparent size, the transformation we imposed did not alter the absolute gaze angles in our display except for that of the horizon line (because we did not alter simulated eye height). This means that our manipulation alone distinguishes between an implicit horizon manipulation and an actual visual horizon manipulation. Moreover, our task was specifically a distance task, rather than a size task. It would be interesting to reconsider Dixon et al.'s eye-height scaling findings in the context of actually changing the visual horizon. Although it remains possible to argue that perceived distance in our experiments was derived from perceived height (it would, perhaps, have been preferable to use flat disks as targets), it would seem more efficient to base such judgments on retinal size, in which case there ought not to have been an effect of the shifted horizon. However, at least one theory of the moon illusion suggests that interactions between perceived size and perceived distance are complex [Kaufman and Rock 1962, 1989]. We note that, assuming an eye height of 1.5 m, if we had presented our (0.25 m) cone at distances over 48 m, its tip would have been above the low horizon, which ought to have forced it to appear taller than eye height. Whether this would produce conflicts with the expected size of the cone—and, therefore, changes in perceived distance—may be worth investigating in future studies.

Although magnitude estimation is currently a less popular psychophysical method, the fact that the verbal estimate data resembles the motor estimates in important ways in our experiment, suggests three comments.

1. Although verbal estimates are notoriously difficult to “get right,” the fact that they track the effects of horizon manipulation so well suggests that they were based on a fairly stable underlying representation (probably not scaled in standard units of measure). While it is premature to conclude that this is the same representation that controls action (and it seems likely, as He et al. [2004], argue, that different task strategies may sometimes produce different perceptual outcomes), the present data do argue for the utility of verbal estimates of distance as a convergent method of assessing various sources of information about distance.
2. Our technique of predicting an altered exponent does not depend on people being able to scale their verbal estimates accurately. While it is probably meaningless to say that the verbal estimates were lower than the motor estimates (since subjective units of “a foot” might simply be off by a scalar value from the objective measure—in which case a verbal “foot” is actually bigger), it appears more meaningful to argue about whether the exponents are similar and this is just the kind of argument that is normally intended by describing visual space as compressive (e.g., Gilinsky 1951). Indeed, for the purpose of computing exponents, it is possible to allow people to give verbal judgments in whatever scale they wish so long as they intend the scale to be a ratio scale (e.g., Teghtsoonian and Teghtsoonian 1969, 1970). We find no evidence that egocentric visual space is compressive in the range of distances we examined.
3. Although not conclusive, our results are entirely consistent with Loomis et al.'s [1992] idea that conscious awareness of egocentric distance may have access to the same (fairly accurate) information as that used for navigation—although the conversion of that information into standard units, such as feet and inches, seems to be off by a scalar factor.

There is one subtly surprising aspect of our results that may cut against the generality of our conclusions. Namely, in the few studies that are published in which people make estimates of distances in feet and inches (e.g., Proffitt et al. 2003), verbal estimates in real outdoor settings appear to be underscaled by a factor of about 0.70. We have observed this same ratio between verbal and motor estimates in previous work we have done in VR (e.g., Durgin et al. 2002). That is, when directed walking

tasks suggested distance was compressed by about 0.70 in VR, verbal estimates were short by a ratio of 0.5, which is roughly consistent with the multiplication of the 0.7 VR distance compression with the 0.7 verbal conversion compression, resulting in a predicted ratio of 0.49. What is striking then, is that the verbal estimates in the present study were as high as they were (about 0.7), representing nearly 95% of the motor estimates. It may turn out that the instructions we gave concerning not counting out units were crucial to producing the result that the same kind of spatial information was available in both the motor and verbal versions of our task.

The virtual environment used in this experiment was a fairly simple one, unpopulated by objects. Despite having a richly textured ground plane (visible in the ranges used, at least), the environment's paucity of landmark features, and the presence of a clearly visible horizon line may have artificially emphasized the use of the horizon line. However, the fact remains that when the horizon line was appropriate, walking and verbal judgments had exponents near 1, suggesting good calibration. Moreover, the distance compression found was similar to that found when looking at video of the real world. Our results with the manipulation of the horizon are striking not just because they show that manipulating it had an effect, but because the magnitude of the exponents were quite close to what we predicted. In a richer environment, the horizon cue might have received less weight, of course.

Although declination from the horizon is an absolute, rather than a relative depth cue, it is possible that, due to individual scaling practices, it acts more as a relative cue in practice. If it were a well-calibrated absolute cue then distance judgments in our displays should not have been as compressed as they were. It seems likely that subjects initially evaluated distances in the virtual environment using many cues, but then primarily used the declination cue as the basis for comparing distances from trial to trial. Thus, our investigation has shown evidence of the horizon line as a relative depth cue, which appropriately affected the relative distances walked or spoken by our subjects.

4. GENERAL DISCUSSION

Given the evidence that distance is underestimated in HMD VR, one goal of this paper has been to characterize the compression of distance using power functions. A further goal has been to show that power function analyses are powerful enough to support using HMD VR as a tool to investigate perception.

4.1 Realism and Distance Underestimation in HMD VR

In Experiment 1 we sought to provide a level of realism in the HMD display that would reveal whether distance errors associated with VR could be eliminated when the real world was presented through the HMD by a video feed. Although there were distinct limitations to our video—HMD apparatus, including a 433 ms lag, it did capture the geometry of the environment properly. Nonetheless, visually directed walking still revealed that distances were underperceived through the HMD (but not with a monocular FOV-restricting goggle). While this result is largely negative—a failure to identify the source of the problem—it clarifies that the issue is not realism in any cognitive sense of the term. It presently seems likely that subtle perceptuomotor failings of HMD displays are to blame. These might include optical distortion, poor absolute screen resolution, fixed accommodation, visual lag, and even headset weight.

It may be that advances in computer graphics and display resolution in conjunction with improved optics may result in a more accurate perception of space in VR. In contrast to Thompson et al. [2004], Interrante et al. [2004] have argued that compression almost completely disappears as graphics improve. Their findings, however, may be artifacts of their methodology (e.g., effects of prior knowledge of the environment or of prior adaptation to VR). On the other hand, we do not take our own work with a video feed-through as conclusive, either, because of the relatively long lag in the apparatus. In

addition, in light of the finding that mechanical aspects of the HMD may effect distance perception [Willemsen et al. 2004], the camera apparatus should probably be made lighter. In any case, we believe that, in spite of intrinsic limitations, HMDs can be used effectively to investigate principles of distance perception.

4.2 The Visual Horizon in HMD VR

Despite the limitations of our HMD as a means of perceiving space, our second experiment successfully used the HMD to investigate angular declination from the visual horizon as a cue to distance. By analyzing the data using theoretically motivated power functions, we showed that both motor and verbal estimates of distance were affected in predicted ways by a subtle manipulation of the apparent horizon line. The presentation of an artificially low visual horizon is not a manipulation that is easy to accomplish without VR. Our experiment showed that such a manipulation of the visual horizon line did indeed produce appropriate changes in the relationships among perceived distances. Moreover, by using similar instructions for verbal estimation and for visually directed walking, we found that the two measures were convergent.

Crucial to the success of our experiments has been the logarithmic analysis of changes in compression with distance (compressiveness). In summary, we have found that egocentric distance judgments in HMDs are compressed, but we did not find evidence that they are compressive either for the real world viewed with a camera or with an accurately rendered virtual environment. By simulating a truncated world, which appeared to be an infinite ground plane, but had a visual horizon about 1.5° below true straight ahead, we have, however, shown evidence that such a manipulation introduces an appropriate change in the apparent geometry of space. We believe this is a particularly useful analytic tool for the analysis of distance perception in VR.

4.3 Indoor/Outdoor and Texture Effects in HMD VR

Recently, we have used similar techniques (verbal and motor) to compare the scaling of indoor vs. outdoor environments in VR [Messing and Durgin 2004a]. Teghtsoonian and Teghtsoonian [1969, 1970] had found that verbal estimates of distances in indoor spaces were expansive (had exponents greater than 1), while those in outdoor spaces were slightly compressive (exponents slightly less than 1). We wondered if the expansive exponents indoors might be caused by the perceptual system inadvertently treating the joint between the floor and the far wall of the indoor environment as the visual horizon. However, although we were initially successful in reproducing expansive exponents for verbal judgments in an indoor virtual environment, we did not find any concomitant changes in motor actions. The exponents for the directed walking task remained near 1. The significance of this result was twofold. First, it indicated that verbal exponents could depart from motor exponents. Moreover, it seemed to undermine the idea that the effect of indoor environments on distance perception had anything to do with misperceiving the horizon, because that should have similarly affected both estimates. On the other hand, it is quite possible that the two tasks differentially encouraged attention to the far wall.

Further investigation has suggested that the factor that determined the difference in verbal exponents between our indoor and outdoor virtual environments had little to do with the structure of the space, but was instead produced by differences in the textures we had used in the two environments [Messing and Durgin 2004b]. The outdoor texture was the same as that employed in Experiment 2 of this paper, and was composed of nondiscrete, blurred elements. The “indoor” texture, on the other hand, consisted of overlapping (and unblurred) small and large circles. Because real indoor and outdoor environments frequently do have different types of visual texture, it is possible that texture effects always underlie these differences, although we are certainly not prepared to make that claim currently. We have not yet isolated the crucial texture parameters in our virtual stimuli. We would argue, however,

that VR, although typically limited in resolution, may be an excellent tool with which to begin to conduct such experiments, despite the fact that rendered spaces seem to appear linearly compressed in HMDs.

5. SUMMARY AND CONCLUSION

We had people view, through an HMD, appropriately scaled nearly-live video of the environment in which they were actually standing. Their performance at the task of walking, without visual feedback, to physical locations previewed in the HMD, suggest that they substantially underestimated the egocentric distances to those target locations in the real environment, just as they do in virtual environments. Perceived distance was compressed, but it was not compressive: All distances were similarly compressed. Compression was evident when the real world was viewed directly.

When we had people make judgments of distance to virtual targets placed in a richly textured virtual field, we measured similar levels of distance compression. When we artificially lowered the visual horizon in this virtual environment, the effect on distance judgments, captured by an analysis of power-function exponents, was precisely that expected, based on the resulting changes in the target's angle of declination below the horizon.

Although distance perception in our HMD is clearly distorted, our experiments indicate that HMD VR can nonetheless be used to successfully investigate basic principles of perception. By using a simple power function model to describe the effects of environmental changes in VR, we have successfully demonstrated the role of the visual horizon in the perception of distance.

ACKNOWLEDGMENTS

The authors would like to thank Joy Mills for helping run Experiment 1. R.M. was supported by the Hans Wallach Fellowship from the Swarthmore College Department of Psychology; This work was supported by a Swarthmore College Faculty Research Grant to F.H.D. These experiments were carried out as part of a R. M.'s senior honors thesis at Swarthmore College.

REFERENCES

- CREEM-REGHEER, S. H., WILLEMSEN, P., GOOCH, A. A., AND THOMPSON, W. B. 2005. The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual environments. *Perception* 34, 2, 191–204.
- DIXON, M., WRAGA, M., PROFFITT, D., AND WILLIAMS, G. 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., AND WALLEY, K. 2002. Perceptuomotor adaptation: More than meets the eye. In *43rd annual meeting of the Psychonomic Society* Kansas City, MO.
- FOLEY, J. M. 1980. Binocular distance perception. *Psychological Review* 87, 5, 411–434.
- GILINSKY, A. S. 1951. Perceived size and distance in visual space. *Psychological Review* 58, 460–482.
- GOODALE, M. AND MILNER, A. 1992. Separate visual pathways for perception and action. *Trends in Neurosciences* 15, 1, 20–25.
- HE, Z., WU, B., OOI, T.-L., YARBROUGH, G., AND WU, J. 2004. Judging egocentric distance on the ground: Occlusion and surface integration. *Perception* 33, 7, 789–806.
- INTERRANTE, V., ANDERSON, L., AND RIES, B. 2004. An experimental investigation of distance perception in real vs. immersive virtual environments via direct blind walking in a high-fidelity model of the same room. In *ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization* (Los Angeles, CA).
- KAUFMAN, L. AND ROCK, I. 1962. The moon illusion: I. *Science* 136, 953–961.
- KAUFMAN, L. AND ROCK, I. 1989. *The Moon Illusion Thirty Years Later*. Lawrence Erlbaum, Mahwah, NJ, 193–234.
- KNAPP, J. M. AND LOOMIS, J. M. 2004. Limited field of view of head-mounted displays is not the cause of distance underestimation in virtual environments. *Presence: Teleoperators and Virtual Environments* 13, 5, 572–577.
- LOOMIS, J. M., SILVA, J., FUJITA, N., AND FUKUSIMA, S. 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., AND 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. AND KNAPP, J. M. 2003. *Visual perception of egocentric distance in real and virtual environments*. Lawrence Erlbaum, Mahwah, NJ, Chapter 1, 21–46.
- MESSING, R. B. AND DURGIN, F. H. 2004a. Compression of distance perception in a live-video-fed head mounted display. *Journal of Vision* 4, 8, 382.
- MESSING, R. B. AND DURGIN, F. H. 2004b. Distance perception and cues to distance in virtual reality. In *ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization* (Los Angeles, CA). 176.
- OOI, T. L., WU, B., AND HE, Z. J. 2001. Distance determined by the angular declination below the horizon. *Nature* 414, 197–200.
- OWENS, D. A. 1987. *Oculomotor information and perception of three-dimensional space*. Lawrence Erlbaum, Mahwah, NJ, 215–248.
- PHILBECK, J. W. AND 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.
- PROFFITT, D. R., STEFANUCCI, J., BANTON, T., AND EPSTEIN, W. 2003. The role of effort in perceiving distance. *Psychological Science* 14, 2, 106–112.
- ROWLAND, J. H. 1999. *The effects of texture on distance estimation in synthetic environments*. Naval Postgraduate School. Department of Operations Research: Master's Thesis.
- TEGHTSOONIAN, M. AND TEGHTSOONIAN, R. 1969. Scaling apparent distance in natural indoor settings. *Psychonomic Science* 16, 6, 281–283.
- TEGHTSOONIAN, R. AND TEGHTSOONIAN, M. 1970. Scaling apparent distance in a natural outdoor setting. *Psychonomic Science* 21, 4, 215–216.
- THOMPSON, W. B., WILLEMSSEN, P., GOOCH, A. A., CREEM-REGEHR, S. H., LOOMIS, J. M., AND BEALL, A. C. 2004. Does the quality of the computer graphics matter when judging distance in visually immersive environments? *Presence: Teleoperators and Virtual Environments* 13, 5, 560–571.
- WANN, J. P., RUSHTON, S., AND MON-WILLIAMS, M. 1995. Natural problems for stereoscopic depth perception in virtual environments. *Vision Research* 35, 19, 2731–2736.
- WILLEMSSEN, P., COLTON, M., CREEM-REGEHR, S., AND THOMPSON, W. 2004. The effects of head-mounted display mechanics on distance judgments in virtual environments. In *ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization* (Los Angeles, CA), 35–38.
- WRAGA, M. 1999a. The role of eye height in perceiving affordances and object dimensions. *Perception and Psychophysics* 61, 3, 490–507.
- WRAGA, M. 1999b. Using eye height in different postures to scale the heights of objects. *Journal of Experimental Psychology: Human Perception and Performance* 25, 2, 518–530.
- WU, B., OOI, T.-L., AND HE, Z. 2004. Perceiving distance accurately by a directional process of integrating ground information. *Nature* 428, 6978, 73–77.

Received December 2004; revised April 2005; accepted May 2005