

Running head: STRATEGIES FOR USING VIRTUAL REALITY

Controlled interaction: Strategies for using virtual reality to study perception

Frank H. Durgin and Zhi Li

Swarthmore College

Frank H Durgin

Swarthmore College, Department of Psychology

500 College Avenue

Swarthmore, PA 19081

fdurgin1@swarthmore.edu

610-328-8678

fax: 610-328-7814

Abstract

Immersive virtual reality systems employing head-mounted displays offer great promise for the investigation of perception and action, but there are well-documented limitations to most virtual reality systems. We suggest strategies for studying perception/action interactions that try to depend on scale invariant metrics (like power function exponents) and careful consideration of the requirements of the interactions under investigation. New data concerning the effect of pincushion distortion on the perception of surface orientation are presented as well as data documenting the perception of dynamic distortions associated with head movements with uncorrected optics. A review of several successful uses of virtual reality to study the interaction of perception and action emphasizes scale-free analysis strategies that can achieve theoretical goals while minimizing assumptions about the accuracy of virtual simulations.

There are two main reasons that the idea of virtual reality (VR) has so deeply appealed to students of perception. The first is the promise of being able to present any large-scale environment desired. To study basic space perception, for example, researchers can change the horizon line while leaving all other scene geometry correct (Messing & Durgin, 2005) and they can make large-scale hills of any surface orientation (Creem-Regehr, Gooch, Sahm & Thompson, 2004; Li & Durgin, 2009; Proffitt, Bhalla, Gossweiler & Mldgett, 1995). The second reason VR has such appeal is because it simulates, and thus allows us to study, the interaction between behavior and perception (e.g., Warren, Kay, Zosh, Dushon & Sahuc, 2001). Such interactivity is fundamental and the control it allows is revolutionary (Loomis, Blascovich & Beall, 1999).

Powers (1973) famously proposed that behavior is best construed as the “control of perception.” By this he meant to argue that perceptual feedback is more than a means of guiding action. It is often the proximal goal of action. At the very least, action involves perceptual prediction (Wolpert & Flanagan, 2001). Immersive virtual environments are intended to capture this interactivity: I turn my head to bring an object of interest into view and it comes into view. I walk forward to experience a new region of visual space and there I am. VR systems can be used to alter the perceptual gain of actions, such as walking or rotating, thus re-specifying the quantitative relationship between perception and action (e.g., Mohler, Thompson, Creem-Regehr, Willemson, Pick & Rieser, 2007).

With regard to creating any environment desired, VR systems have proven disappointing in practice because of their apparent failure to generate the perceptions we expect. Interactivity provides a visceral sense of immersion in a virtual environment that can be quite compelling even when the computer graphics are clearly artificial. However, the virtual spaces created nearly always appear much smaller than they are simulated to be. This phenomenon is well-documented but still not well understood

(Loomis & Knapp, 2003; Thompson, Willemsen, Gooch, Creem-Regehr, Loomis & Beall, 2004). If one is studying qualitative spatial cognition, these quantitative errors of scale may not be critical. For example, to study spatial re-orientation in a small virtual room, we simply simulated a much larger room to achieve the sense of spatial scale that we intended (Grodner, Pietsch & Durgin, 2009). But this is not ideal. There is some evidence that interacting with a virtual environment that exactly matches the real environment one has just been immersed in helps scale the virtual space (Akagi & Durgin, 2007; Interrante, Anderson & Reis, 2006; Interrante, Reis, Lindquist & Anderson, 2007), but requiring prior exposure to a real version of the virtual space seems to undermine one of the chief virtues of VR: the ability to present any environment desired.

Because the perceived scale of virtual environments can generally be assumed to be different from (smaller than) the simulated scale, studies of space perception in VR are faced with a conundrum: It is possible to manipulate the environmental information to try to investigate the effectiveness of certain kinds of visual information, but the absolute scale is known to be wrong. What counts as accurate and what counts as inaccurate? It is worth considering whether the second-order scaling of VR (compression of scale with distance) is more or less uniform. If it is, then we can ask questions about the rate of compression such as have been traditionally been asked in studies of distance perception using scale-free methods (e.g., Teghtsoonian & Teghtsoonian, 1970). The scale-free method of using the exponents of power functions to represent perceptual experience has proven very useful in a variety of studies using VR in our lab. Here we review examples of how we have previously applied the scale free method in VR to distance perception, self-motion perception and geographical slant perception. Finally we will present a novel experiment using the method to demonstrate the perceptual importance of a subtle optical correction in virtual environments.

The perceived distance to an object along the ground may be calculated in terms of its visual declination below the horizontal (Ooi, Wu & He, 2001). That is, assuming a level ground plane and a standing observer, there is direct relationship between how far away an object is from the observer and the visual angle below horizontal where its base appears to meet the ground. Messing and Durgin (2005) sought to test whether subtly lowering a visible horizon (by 1.5 deg) would affect distance estimates in near space (1-7 m). Knowing that distance would be underestimated in virtual reality, and that the absolute shifts in perceived distance would be difficult to detect because of inter-subject variability, Messing and Durgin expressed their predictions of horizon effects in terms of differences in the exponents of power functions fit to distance estimates with a true (normal) horizon and an artificially low horizon. Although a power function only approximates the expected perceptual change, lowering the horizon line by 1.5 deg predicts an exponent of about 1.15 for this range of distances. Thus, while each participant made ten distance estimates with each horizon line, the statistics used for analysis simply examined the individual exponents of power functions for each horizon condition by participant. Both for verbal estimates and for action measures (walking to the previewed target location without visual feedback), it was found that although distances were underestimated in all conditions, exponents were normal (near 1) with the correct horizon. This suggests the compression of perceptual space in VR was uniform. Moreover, the shifted horizon reliably produced the predicted change in exponent for both measures. This illustrates the advantage of looking for higher order (scale free) predictions when studying perception in virtual environments – even when a scaled quantity (i.e., perceived distance) is under investigation.

The ability to create interaction between action and perception is not unique to immersive virtual reality systems, but the power and generality of the technology makes the study of that interaction much more versatile. Consider that Wallach and Flaherty

(1975) built a mechanical device that allowed them to present an expanding flow-field (based on a rotating spiral image) that was yoked to the mechanical motion of an observer on a swing, thus allowing them to demonstrate that motion aftereffects were reduced when the motion stimulus (the expanding flow-field) was consistent with the forward physical motion of an observer (no such reduction was found if the expanding flow field was presented on the back swing instead). This important finding was later replicated using a computer-controlled motion stimulus that was yoked electronically to the motion of a cart carrying an observer (Harris, Morgan & Stone, 1982). With the advent of immersive VR, such contingencies have become the bread and butter of a variety of studies. For example, using immersive VR, we have quantified the reduction in perceived visual speed during self-motion and shown that it can be modeled as a subtractive function proportional to the speed of self-motion under a variety of conditions (Durgin, Gigone & Scott, 2005).

Once again, our initial strategy was to use scale free representations (exponents of power functions fit to verbal speed estimates) to create a basis for comparing two conditions (static vs. moving observers) rather than trying to directly measure a change in any one particular visual speed. The versatility of immersive virtual reality meant that we could easily compare the amount of speed reduction caused by biomechanical activity (walking on a treadmill) with that from physical self-motion (being carried on a cart) and from the two in combination (walking on solid ground in a wide area immersive virtual environment). The comparison could be made using the same visual information (virtual environment) and measurement strategies across all these conditions. The amount of speed reduction in the case of normal locomotion was roughly the sum of the amounts found separately for biomechanical activity of walking and for (passive) physical translation. Although the scaling of perceived speed based on VR might be distorted in much the way that visual space is distorted in VR, the theoretical inferences

of these studies is not limited by this fact because they focus on second order (scale free) relationships.

Similarly, evidence that the perceptual discrimination function for visual speed is altered during self-motion, such that speed sensitivity is tuned to expected visual speeds (Durgin & Gigone, 2007) is not undermined by the knowledge that the spatial scaling of VR is known to be problematic. Discrimination is discrimination. Similarly, evidence that object motion may not be susceptible to this same kind of speed subtraction (Durgin, Reed & Tigue, 2007) and that perceived biomechanical speed of self-motion is related to stride frequency (Durgin et al., 2007) are findings that were facilitated by being able to control the interaction between perception and action in ways that VR makes easier to achieve (though we did also develop a specialized robotic treadmill for this latter work). Thus, although biases in the perception of distance in VR have not been eliminated, it is possible to study variables such as perceived visual speed by modeling higher-order relationships that are not disturbed by scalar errors in the perception of distance.

Graphics systems today are much faster than those of a decade ago, and their use has gotten more sophisticated. Still, most head mounted displays (HMDs) have evident optical pincushion distortions (the optically-correct Sensics HMDs are a notable exception). Whereas correcting for these distortions has not been a standard practice until recently, correction is now relatively easy to achieve in commercial software packages. The need for such correction is evident when observers are encouraged to scan a scene. In an initial study of demand-characteristics in the perception of slope (Russell & Durgin, 2008; discussed in Durgin, Baird, Greenberg, Russell, Shaughnessy & Waymouth, 2009), we noticed that the perceived surface orientation of our virtual hills appeared to dynamically distort when we scanned up and down. Although a study using these uncorrected displays produced the anticipated differences in judgments of slope based on demand characteristics of wearing a backpack (Russell & Durgin, 2008), we

have noted that many prior studies of slope perception in VR did not use optical correction. For example, Proffitt et al. (1995) report what they regard as “normative” data that includes data from very primitive VR (e.g., uncorrected optics, refresh rates of about 10 Hz, and lags of several hundred ms). In retrospect, it seems likely that their participants depended on a pictorial cue (path edge convergence) that turned their slope estimation task into a cognitive problem-solving task. Other limitations of their design and measurement techniques have recently come to light (Bridgeman & Hoover, 2008; Durgin, Hajnal, Li, Tonge & Stigliani, 2010).

As an alternative model of using VR to study slope perception, Li and Durgin (2009) studied the effect of direction of gaze on downhill slope perception. They used both virtual hills and real (smaller) surfaces to collect slope judgments when observers were positioned at different horizontal distances from the surfaces so that their view of the target surface was either more parallel with the surface (as when standing back from the edge of the top of a hill) or more frontal to it (as when standing nearer the edge and looking more directly down onto the surface). Again, rather than directly interpreting the verbal judgments, Li and Durgin modeled the relationship between optical slant (the orientation of the surface relative to the line of gaze), gaze orientation, and perceived slope. Using a simple geometric model, they found that the functions relating true/simulated optical slant to (inferred) perceived optical slant became appropriately aligned when the model assumed that perceived direction of gaze was misperceived by a constant gain factor. The analysis, again, depends entirely on comparing differences in representations of perceived optical slant without assuming that the scaling in VR corresponds to that in the real world. The overestimation of perceived head and gaze orientation was then measured separately in two further experiments and found to fit the predictions from the VR data.

The model developed by Li and Durgin (2009) showed that the perception of optical slant was exaggerated in both VR and the real world, but that the perception of head orientation was also (and similarly) exaggerated. As a consequence, for horizontal surfaces, the two biases should cancel each other out, but for uphill, sloped surfaces, the two errors would tend to combine to exaggerate the perceived slope, thus providing a new account of errors in uphill slope perception. For downhill surfaces the two biases partly compensate for each other, resulting in a shallower perception of downhill slopes when standing at the edge than when standing further back. Indeed, whereas the mathematical function relating optical slant to perceived optical slant was different in VR than in the real world (a scaling difference, which was not surprising given the myriad differences in visual information available in the two cases), the functions relating perceived head and gaze orientation to perceived slope generalized across both contexts. What is particularly important about this example is that, rather than attempting to reproduce a known pattern of data in VR, Li and Durgin discovered the geometric relationship using VR and then confirmed that it extended to the real world using small-scale wooden surfaces.

It is important to note that Li and Durgin (2009) used a software shader in Virtools, a commercial rendering system, to correct for the pincushion distortion produced by the optics of their nVis HMD. Pincushion distortion results from imperfectly corrected optics in the HMD, and its effects are illustrated in Figure 1. Kuhl, Thompson and Creem-Regehr (2009) have recently described their technique for correcting and calibrating their nVis HMD and they arrived at very similar values to ours (42 x 34 deg) for their true display size, though our optics appear to have produced greater pincushioning than theirs prior to correction. We recommend consulting their strategy for correcting and calibrating an HMD.

-----Figure 1 about here-----

Kuhl et al. (2009) reported that the presence or absence of pincushion correction (at the levels they tested) had no effect on judgments of distance. However, our informal observations regarding the perception of slanted surfaces with and without correction suggest that there might be measurable effects on slope perception. Moreover, rather than approaching the problem in the classic way of looking for absolute differences in slope estimates, we again sought to detect higher-order differences in perception that might emerge in log space (exponents of power functions).

Most VR displays use stereoscopic disparity as a source of information about distance. Pincushion distortion would cause disparity information in the outer corners of the display to be exaggerated relative to the middle of the display. In theory, this could tend to make various parts of the lower and upper parts of the display appear nearer and farther than appropriate and thus produce evidence that the slope is both shallower and steeper than intended. Similar effects could arise from the 2D distortion of texture gradients in the image. The expected outcome should be an increase in uncertainty and a concomitant reduction in sensitivity. This could be reflected as a decrease in the exponent of the power function – a compression of the perceptual scaling of slant perception.

Moreover, because there is a great deal of variability in the manner with which people may provide numeric estimates of slope, the advantage of summarizing estimates in terms of power function exponents is that the exponent of the power function can retain information about the rate with which slope information (for example) is used while ignoring the absolute scale (which might otherwise add a great deal of between–subject variance).

We therefore expected that the effect of pincushion distortion would be to reduce the exponent of a power function fit to the data by muting sensitivity to differences between slopes. To confirm our subjective impressions that pincushion correction caused perceptual instability during head motions, we also collected ratings of stability. We expected that direct ratings of the stability of scene geometry during head motions would be higher for a corrected display than for an uncorrected display.

Methods

Participants. Twenty-four Swarthmore College undergraduate students who were naïve to the hypotheses were reimbursed \$5 for their participation.

Apparatus. Displays were rendered to an nVis HMD, with nominal XGA resolution updated at 60Hz using Virtools 4.1 rendering software. Head orientation was monitored by a Vicon optical tracking system with minimal lag. Participant sat in a comfortable chair while looking at displays that were either corrected using an HLSL shader or presented uncorrected. Our correction constant was .3 (see Kuhl et al., 2009, for an explanation of this parameter).

Virtual environment. The virtual environment was limited to a single rough surface composed of a flat sand-like texture and a plethora of jagged 3D “rocks” that protruded from the surface at irregular intervals. In addition, a small white ball on the surface of the hill served as a fixation point. The ball was 2 m from the observer (either directly ahead at eye level or along a line 30 deg declined from horizontal). For slope estimation, participants were required to fixate the ball and discouraged from moving their head. For “distortion” judgments, they were required to move their head vertically between the straight-ahead position and the ball in its lower location. Note that in the forward gaze condition, blue sky was visible beyond the end of the surface for the lowest

3 orientations. In the downward gaze condition, the edge of the surface was not visible. Each view of the 12-deg surface is shown in Figure 2.

-----Figure 2 about here-----

Design and procedure. The experiment had two parts (slope estimation and stability judgments). The slope estimates were always completed first in two blocks. In one block gaze was forward. In the other block, gaze was declined. The order of these blocks alternated between subjects. In each block of trials 5 different slopes were tested in each optical condition (10 trials total per block). The simulated slopes were 6, 12, 18, 24 and 30 deg relative to the horizontal. Order of presentation was randomized within each block. A black screen appeared for 1 second between trials so that changes in correction and surface orientation would be masked. Slope judgments were made verbally in deg ranging from 0 for horizontal to 90 for vertical. Following the initial two blocks of verbal slope estimates, the second part of the experiment presented three slopes (6, 18, and 30 deg) in each of the two optic conditions (6 trials in random order). Participants rated the amount of distortion experienced during head movement on a scale of 1 (completely stable) to 10 (very distorted). It was explained to participants that as they moved their head on each trial they should try to notice whether the slanted surface appeared to rotate and/or distort, or whether it appeared fully rigid and stable.

Analysis. Angle estimates and slopes were converted to logarithms (after adding 1 to eliminate zeroes) and the slopes of least-squares fits (i.e., power function exponents) were computed for each participant for each head orientation and optical correction state. Note that the use of the slope in log-log space to provide a scale free summary of a set of data does not require that the data have the form of a power function. Rather this use of a power-function exponent is a one-parameter model that

captures spread while neglecting scale. Scale is contained in the second parameter of a power-function – the intercept in log-log space, which we ignore in this case. Raw rating scale measures were compared according to optical correction condition.

Results

Slope estimates. We predicted greater sensitivity (higher exponents) to slope differences when optical correction was applied. A 2 (Gaze: Up, Down) x 2 (Optical Correction: Corrected, Uncorrected) repeated-measures ANOVA revealed only a marginally reliable effect of Optical Correction, $F(1,69) = 2.70$, $p = .0528$ (one-tailed) in the predicted direction. However, the analysis revealed a highly reliable effect of Gaze, $F(1,69) = 26.4$, $p < .0001$, likely indicating a strong influence of the presence of a horizon line formed by the top of the hill in the Gaze Up condition (see Figure 2), where the average exponents (1.13) were much higher than in the Gaze Down Condition (0.52). Because our predictions regarding pincushion correction were based on expected distortions of optical slant information (from texture and stereopsis), whereas the horizon line could have acted as a cognitive reference, the apparently large effect of the hill-top horizon line was treated as a confounding influence and data from the Gaze Down condition, in which the horizon was never visible, were analyzed separately.

Median estimation data are shown in Figure 3, in log-log space, as a function of Optical Correction for the Gaze Down condition. When a repeated measures ANOVA was conducted on the Gaze Down data separately, the average exponents in the Uncorrected pincushion condition (0.37, $se = 0.17$) were found to be reliably lower than those in the optically Corrected condition (0.68, $se = 0.21$), $F(1, 23) = 8.06$, $p = .0093$. Thus, as expected, pincushion distortion seems to have degraded the scaling of the optical slant information required to estimate optical slant in the gaze down-condition. Note that an analysis of least-squares fits in linear space, because it is sensitive to

arbitrary scale selection, did not reliably detect even the large difference in the Gaze Down condition, $t(23) = 1.47$, $p = .155$, indicated by the analysis in log-log space. Our argument is not that there is no effect in linear space, but rather that such analyses are likely to be less sensitive to theoretically important differences because they are scale-sensitive even when the theoretical question under investigation are not.

-----Figure 3 about here-----

Ratings of dynamic distortion. Pincushion distortion is not obvious during static viewing of irregular textures, but produces dynamic distortions of geometry when the head is moved. Average ratings of dynamic distortion are shown in Figure 4. For our corrected displays, the average rating of dynamic distortion was only 3.4 (on a 1 to 10 scale), whereas the average rating in the pincushion condition was 6.0, which was reliably greater, $t(23) = 4.45$, $p = .0002$.

-----Figure 4 about here-----

Summary of results. Uncorrected pincushion distortion in our nVis HMD had measureable consequences for the extraction of spatial parameters, like surface orientation. Pincushion distortion also produced noticeable dynamic distortion when the head was re-oriented. Note that in our study of perceived surface orientation, we did not allow head movements because we sought to isolate static effects of pincushion distortion. It is likely that dynamic aspects of pincushion distortion could add additional noise to the perception of surface layout. Although human perception is highly adaptive, studies that intend to investigate active perceptual exploration of spatial parameters using HMDs for immersive VR should seek to correct their optics.

Discussion

VR is a powerful tool for studying the interaction between perception and action. Whereas VR systems remain imperfect in their scaling of space, we have recommended using scale-free strategies for harnessing the power of VR.

One strategy, the use of power function exponents, is both a theoretical strategy and a methodological one: Whereas the variability in our participants' slope judgments meant that standard analyses of means would be unable to detect effects of pincushion distortion on slope perception, the use of scale-free, one-parameter representations of data (such as power function exponents) have, once again, provided a way of extracting the essential pattern of data we expected: When slope estimates must be based on optical slant, pincushion distortion adds uncertainty to estimates. This can be measured as a proportionally compressed set of estimates of slope (a lower exponent) for the distorted case than the corrected case, as we have shown. Verbal judgments are prone to biases, but considered in aggregate and transformed into higher-order (scale-free) statistics, they seem to be quite sensitive to subtle perceptual effects. Messing and Durgin (2005) depended on the theoretical value of higher-order statistics to detect the effect of a shifted horizon. Here we have emphasized that scale-free analyses are better suited to detect effects of pincushion distortion on the static perception of spatial variables like surface orientation.

Our second scale-free strategy is to take advantage of perceptuo-motor interaction, but to do so in a way that demands as few assumptions about the stimulus presented as possible. When Li and Durgin (2009) discovered that the proprioception of head orientation seemed to play a role in the misperception of downhill slope, they investigated this by creating contexts in which the interaction of the observer with the virtual world included head motions. (They did not, for example, as we have done here,

allow participants to keep their head declined for a block of trials.) Their modeling, however, could remain neutral in its assumptions about the verisimilitude of VR because it did not depend on the form of the relationship between perceived and simulated optical slant, but only assumed that there was a stable relationship. That is, it did not depend on the scale being accurate. Nonetheless, because metric space perception itself was under investigation (i.e. slope perception), and head movements were required, the elimination of dynamic distortions due to pincushion distortion seems important to their use of immersive VR.

Whereas immersive VR remains an imperfect medium for the study of perception, both the use of scale-invariant metrics (like power function exponents) and the carefully controlled engagement of perception and action as a principled research strategy hold great promise while we wait for future improvements.

References

- Akagi, M., & Durgin, F. (2007). Accurate perception of visual space from live-video in a head-mounted display [Abstract]. *Journal of Vision*, *7*, 283a. doi:10.1167/7.9.283
- Bridgeman, B., & Hoover, M. (2008). Processing spatial layout by perception and sensorimotor interaction. *Quarterly Journal of Experimental Psychology*, *61*, 851–859.
- Creem-Regehr, S. H., Gooch, A. A., Sahm, C. S., & Thompson, W. B. (2004). Perceiving virtual geographical slant: Action influences perception. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 811-821. doi: 10.1037/0096-1523.30.5.811
- Durgin, F. H., Baird, J. A., Greenburg, M., Russell, R., Shaughnessy, K., & Waymouth, S. (2009). Who is being deceived? The experimental demands of wearing a backpack. *Psychonomic Bulletin & Review*, *16*, 964-969. doi: 10.3758/PBR.16.5.964
- Durgin, F. H., & Gigone, K. (2007). Enhanced optic flow speed discrimination while walking: Contextual tuning of visual coding. *Perception*, *36*, 1465-1475.
- Durgin, F. H., Gigone, K., & Scott, R. (2005). The perception of visual speed while moving. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 339-353.
- Durgin, F. H., Hajnal, A., Li, Z., Tonge, N., & Stigliani, A. (2010). Palm boards are not action measures: An alternative to the two-systems theory of geographical slant perception. *Acta Psychologica*. doi: 10.1016/j.actpsy.2010.01.009
- Durgin, F. H., Reed, C., Tigue, C. (2007). Step frequency and perceived self-motion. *ACM: Transactions in Applied Perception*, *4:5*, 1-27.

- Grodner, D., Pietsch, C., & Durgin, F. H. (2009). The use of categorical features in adult spatial reorientation. *Proceedings of the 2009 Meeting of the Cognitive Science Society*. Amsterdam: The Netherlands.
- Harris, L. R., Morgan, M. J., & Still, A. W. (1981, September 10). Moving and the motion after-effect. *Nature*, 293, 139–141.
- Interrante, V., Anderson, L., & Ries, B. (2006). Distance perception in immersive virtual environments, revisited, *IEEE Virtual Reality 2006*, pp. 3-10.
- Interrante, V., Ries, B., Lindquist, J., & Anderson, L. (2007). Elucidating factors that can facilitate veridical spatial perception in immersive virtual environments, *IEEE Virtual Reality 2007*, pp. 11-18.
- Kuhl, S. A., Thompson, W. B., & Creem-Regehr, S. H. (2009). HMD calibration and its effects on distance judgments. *ACM Transactions on Applied Perception*, 6:19.
- Li, Z., & Durgin, F. H. (2009). Downhill slopes look shallower from the edge. *Journal of Vision*, 9(11):6, 1-15. doi: 10.1167/9.11.6
- Loomis, J. M., Blascovich, J. J., & Beall, A. C. (1999). Immersive virtual environment technology as a basic research tool in psychology. *Behavioral Research Methods, Instruments, & Computers*, 31, 557-564.
- Loomis, J. M., & Knapp, J. M. (2003). Visual perception of egocentric distance in real and virtual environments. In L. J. Hettinger and M. W. Haas (Eds.) *Virtual and Adaptive Environments* (pp 21-45), Lawrence Erlbaum, Mahwah, NJ.
- Messing, R., & Durgin, F. H. (2005). Distance perception and the visual horizon in head-mounted displays. *ACM Transactions on Applied Perception*, 2, 234-250. doi: 10.1145/1077399.1077403
- Mohler, B. J., Thompson, W. B., Creem-Regehr, S. H., Willemsen, P., Pick, H. L., & Rieser, J. J. (2007). Calibration of locomotion resulting from visual motion in a

- treadmill-based virtual environment. *ACM Transactions on Applied Perception*, 4, 1-15. doi: 10.1145/1227134.1227138
- Ooi, T. L., Wu, B., & He, Z. (2001). Distance determined by the angular declination below the horizon. *Nature*, 414, 197–200.
- Powers, W. T. (1973). *Behavior: The control of perception*. Chicago: Aldine.
- Proffitt, D. R., Bhalla, M., Gossweiler, R., & Midgett, J. (1995). Perceiving geographical slant. *Psychonomic Bulletin & Review*, 2, 409–428.
- Russell, R., & Durgin, F. H. (2008). Demand characteristics, not effort: The role of backpacks in judging distance [Abstract]. *Journal of Vision*, 8(6), 755a.
- Teghtsoonian, R., & Teghtsoonian, M. (1970) Scaling apparent distance in a natural outdoor setting. *Psychonomic Science*, 21, 215-216.
- Thompson, W. B., Willemson, P., Gooch, A. A., Creem-Regehr, S. H., Loomis, J. M., & 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, 560–571.
- Wallach, H., & Flaherty, E. W. (1975). A compensation for field expansion caused by moving forward. *Perception & Psychophysics*, 17, 445–449.
- Warren, W. H., Kay, B. A., Zosh, W. D., Duchon, A. P., & Sahuc, S. (1991). Optic flow is used to control human walking. *Nature Neuroscience*, 4, 213-216.
- Wolpert, D. M., & Flanagan, J. R. (2001). Motor prediction. *Current Biology*, 11, R729–R732.

Acknowledgements

Supported by a Swarthmore College faculty research grant.

Figure Captions

Figure 1. Pincushion distortion (left) means that a rectangular image (middle) will end up getting displayed to the eye appearing compressed (left). Most of the geometrical distortion is in the periphery, though the center of the display is also minified. Pincushion correction algorithms spherically pre-distort the image in the opposite direction (right) so that the optical pincushion distortion ends up producing the originally intended rectangular image.

Figure 2. Sample frames from the experimental displays. Each represents a 12-degree slope. The one on the left is depicted with gaze declined by 30 degrees from horizontal. The image on the right was viewed with horizontal forward gaze toward the white ball.

Figure 3. Median estimates of geographical slant with and without pincushion correction are shown for the Gaze Down condition in log-log space. These medians suggest that in the presence of pincushion distortion, surface orientations were less well differentiated from each other. Power function exponents represent the best-fitting slope in log-log space. Statistics based on analyses of power function exponents revealed that this pattern of compressed responding was statistically reliable.

Figure 4. Ratings of dynamic scene distortion during head motion as a function of whether pincushion distortion is present or corrected. Standard errors of the means are shown.




Figure 1





Figure 2

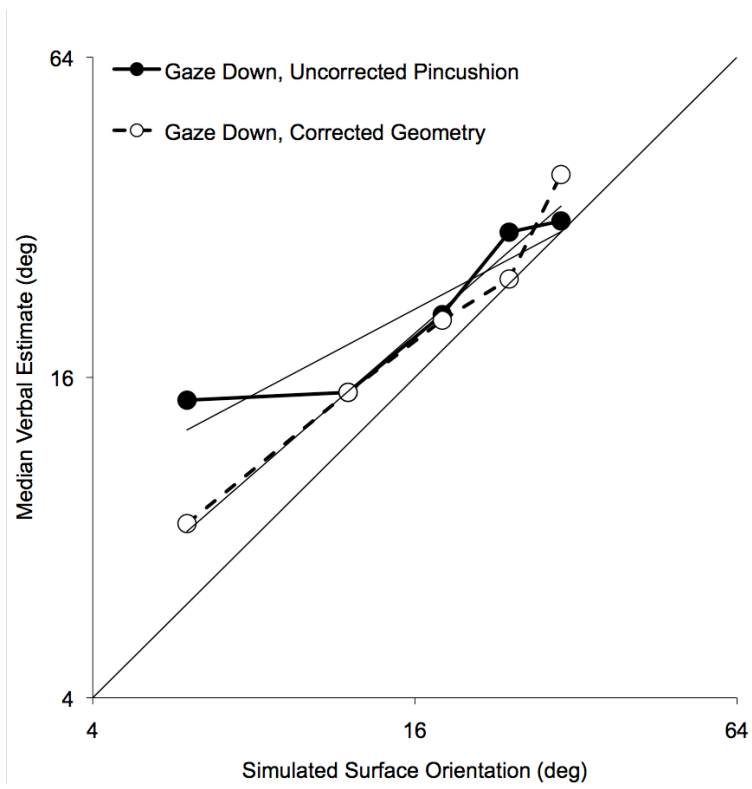


Figure 3

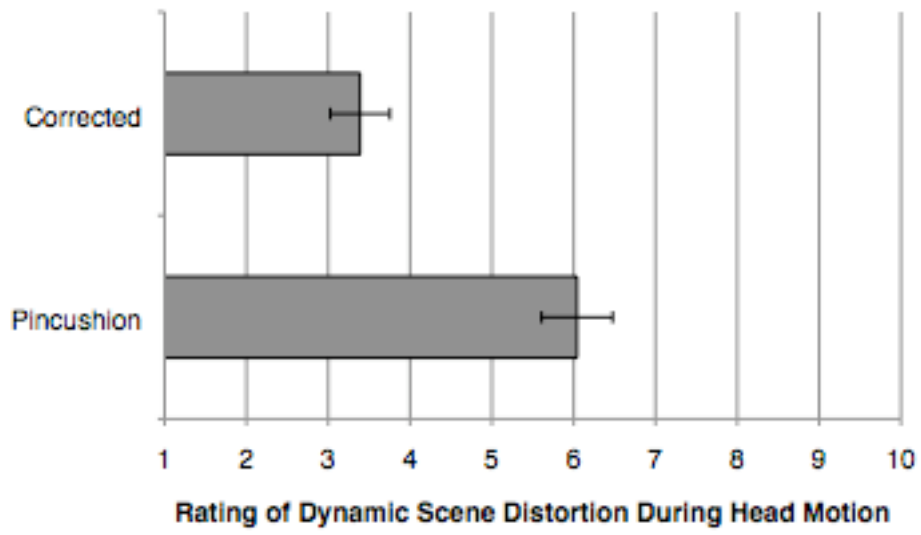


Figure 4