

Evaluating Space Perception in NPR Immersive Environments

Amy Ashurst Gooch

Peter Willemsen

University of Utah

www.cs.utah.edu

Abstract

Non-photorealistic rendering (NPR) methods allow us to emphasize or omit detail in order to communicate information more effectively. An important issue to consider is how can we evaluate and validate the communication adequacy of these techniques? In this paper we present a methodology for quantitatively evaluating spatial perception in a functional, NPR, stereo, immersive environment. The methods described in this paper borrow from both the art and perceptual psychology communities and will provide NPR researchers with a methodology for conducting user studies to perceptually validate new rendering approaches within immersive environments.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—3D Graphics

Keywords: perception, human factors, immersive environments, virtual environments, virtual reality, egocentric, non-photorealistic rendering.

1 Introduction

Many NPR techniques have evolved by cultivating well-developed artistic processes originally intended to enrich the sense of realism, presence, and information portrayed in an artistic medium. In the past NPR researchers have looked to the artistic community for inspiration and now must also focus on psychology and what is known about human cognition and perception to better understand and evaluate how NPR might be used to convey a sense of space in immersive environments. In understanding human perception we may more aptly convey information through NPR methods.

Functional realism, as described by Ferwerda [2001], is differentiated from physical and photorealistic realism in that functional realism provides knowledge about the properties and relations of objects allowing a user to make reliable visual judgments and perform useful visual tasks. For example, many paintings, by artists such as Pollock, Monet, or Picasso, appeal to individual viewers in different ways. In contrast, technical drawings, such as those that might be found in an automobile manual, attempt to communicate the same information to each viewer. This latter work is considered functional because the communicated information is intended to be invariant across viewers. This is an important distinction and one that NPR is well-suited to address. Non-photorealistic rendering,




Figure 1: A screen shot of the NPR hallway, displaying only feature edges, with a scaled inset of a screen shot looking down at the target.

by definition, represents a form of functional realism. We are interested in how the functional realism of NPR can impart a sense of space to people viewing an immersive environment.

In this paper, we describe a quantitative study that evaluates one aspect of space perception in NPR immersive environments, emphasizing the use of an experimental methodology well-suited for understanding spatial perception. Note that this paper does not attempt to justify the functional use of NPR methods for conveying spatial information in immersive environments. The chief contribution of this paper is to emphasize the importance of evaluating NPR methods in the context of related psychology research. The data presented in this paper focuses on the ability of humans to perceive depth when asked to perform directed walking tasks in an environment viewed through a head-mounted display (HMD) system and rendered only with feature edges [Gooch and Gooch 2001]. An example of this type of environment is shown in Figure 1. Our hypothesis is that participants will be able to perceive distances in this simple NPR immersive environment, although perhaps not as well as reported by some studies that used traditionally rendered, immersive environments.

In Section 2 we review research in space perception, including work from psychology, as well as work which has been conducted in immersive environments. Sections 3 and 4 describe the experimental methodology, equipment, and results obtained from our study. The remainder of the paper discusses our conclusions and the role that interdisciplinary collaboration plays in evaluating the effectiveness of NPR in immersive environments.



Figure 2: A photograph of the real hallway, with a scaled inset of a screen shot looking down at the target

2 Related Work

Several computer graphics researchers have been motivated by perceptual psychology, including those who presented a wide range of topics at the ACM-sponsored Campfire on “Perceptually Adaptive Graphics” [2001]. In addition, several research projects in the area of NPR have been motivated by perceptual psychology, including some of those presented at the SIGGRAPH 2001 Course on NPR scientific visualization [Healey et al. 2001]. However, few of these computer graphics researchers have conducted formal user studies. An exception is Interrante [1996] who conducted a user study to evaluate her artistically driven volume renderings. She evaluated how judgments about shape and relative depth of a layered transparent surface may be affected by the addition of a opaque texture.

In order to experimentally evaluate NPR techniques in immersive environments, we examined research by perceptual psychologists and computer graphics researchers. Depth perception is particularly important in immersive environments and is controlled by several physiological and pictorial cues, including accommodation, convergence, stereopsis, familiar size, relative size, eye height, aerial perspective, texture, shading, and motion parallax [Rogers 1995; Cutting and Vishton 1995; Palmer 1999; M.W. Dixon and Williams 2000]. In this paper, our experiment focuses on the perception of absolute egocentric distances in immersive environments. Absolute egocentric distance is defined to be the actual metric distance between an observer and an external object.

Measuring the accuracy of depth perception via verbal reporting is not ideal since subjects are often unaware of systematic bias in what they say they see compared to what they actually perceive. In addition, subjects may believe they are inaccurate and change what they report to reflect these beliefs. These subjects can contaminate the results of the study by modifying their verbal reports [Loomis et al. 1992]. Several perceptual psychologists have shown that directed action tasks such as blind walking to a previously viewed target are good probes for investigating how physical space maps to visually perceived space [Rieser et al. 1990; Loomis et al. 1992; Philbeck and Loomis 1997; Philbeck et al. 1997]. Results from these studies, conducted in real world buildings and outdoor spaces, under full cue conditions, show that subjects are accurate at judging distances between 2 and 25 meters.

Visually directed action tasks have been used as probes for gathering data on how humans perceive distances in virtual environments [Witmer and Kline 1998; Knapp 1999; Loomis and Knapp



Figure 3: A picture illustrating the equipment a subject wears while participating in the immersive environment.

in press]. Results from these and other distance estimate studies [Lampton et al. 1995] have typically shown a large underestimation in perceived distance in the virtual environment that varied from about 50% to 70% of the intended distances. In our previous research, we examined perceived egocentric distances in the real world and in immersive environments using visually directed walking tasks [Willemsen and Gooch 2002]. We found that subjects viewing a photographic panorama immersive environment judged distances to be 87% of the intended distance while subjects judged distances to be 82.5% in a traditionally rendered OpenGL immersive environment. Results from this study are consistent with previous work showing that there is a significant difference in distance judgments between real and immersive environments (e.g., [Knapp 1999]). Although, we were unable to explain why our previous work found underestimation values which were out of the typical range gathered by other researchers, the observed decrease in underestimation may have been due to the display or image quality which differed from previous experiments. However, the differences between distance judgments in photographic environments and environments modeled by traditional computer graphics are small, suggesting that the display device effects distance judgments in immersive environments.

Our work combines the efforts of perceptual psychologists with current NPR techniques to provide the NPR research community with a methodology for evaluating the fidelity of information provided by our imagery.

3 Experiment Methodology

We conducted our experiments with two environmental conditions. In the first condition, shown in Figure 2, subjects performed the direct walking task in the real (physical) hallway. In the second condition, shown in Figure 1, subjects viewed a scale model of the hallway through an HMD in which boundaries, creases, and silhouettes were drawn in black on a pale background.

In order to probe subjects’ perceptions of egocentric distances in our NPR immersive environment and in the real world, we used a direct walking task in which subjects walked without vision to a previously viewed target. The target was placed on the ground at distances of 2m, 3.5m, and 5m from the subject. These distances were chosen due to the limits of our tracked lab space. All subjects were college age volunteers, were given a stereogram eye test, and had normal or corrected to normal vision with the use of con-




Figure 4: Average intended versus perceived distance in feature-edge immersive environment shown in Figure 1. Error bars represent ± 1 SEM (Standard Error of the Mean). The dashed line represents ideal performance.

tact lenses. Our subject pool consisted of 6 females and 6 males. Both environmental conditions were presented binocularly. Three training trials were conducted for each condition followed by a random ordering of three trials at each of the three distances, resulting in each subject experiencing 12 trials per condition. The order in which the environment conditions were presented was evenly distributed over the subject pool as well. No feedback was provided during any phase of the experiment.

Subjects wore white-noise generating headphones to help remove local sounds. Subjects also wore a collar which physically and visually restricted the view of their feet since feet are not rendered in the virtual environment. The collar, shown on a person in Figure 3, was constructed of heavy duty poster-board cut with an exterior diameter of 54.5 cm and an interior diameter of 18.5 cm, with a slice through one side of the ring such that it could be twisted in order to be placed around the neck. On the bottom side of the collar were strips of velcro to which two foam blocks (25.5 x 15.5 x 15.5 cm) could be affixed to keep the collar in place, allowing the subject to move freely.

Prior to each experimental condition, subjects practiced blind walking for at least 5 minutes. The training is helpful in building trust between the experimenter and the subject, but more importantly accustoms the subject to walking blind with the headphones and collar. Subjects were trained to respond to verbal commands to start walking and stop walking. After the training session for the immersive environment condition, subjects were led, still blindfolded, into the immersive environment laboratory so that their movement during the experiment would not be inhibited by *a priori* knowledge of the location of the walls in the lab. While keeping their eyes closed, the HMD was placed upon their head. Once on, subjects were allowed to open their eyes and adjust the fit and focus of the HMD. After which we sampled the orientation of the HMD on the subject's head to calibrate the visual rendering frustum, aligning it with the natural resting position of the HMD on the subject. The calibration ensures that the horizon line at in-




Figure 5: Average intended versus perceived distance in the real hallway, as shown in Figure 2. Error bars represent ± 1 SEM. The dashed line represents ideal performance.

finitly is correctly situated in the middle of the display screen as the subject looks straight ahead.

For each trial, subjects were shown a target located on the floor. In the real world the target was a disk of red foam-core board, about 37cm in diameter. In the NPR environment, the target was of the same dimensions, represented only by its feature edges. Subjects were instructed to obtain a good image of the target and their local surroundings. Each subject was told that a "good image" is obtained if when the subject closed their eyes, he or she would still be able to "see" the environment, and most importantly, the target. Cognitive maps such as these have been extensively explored in the perception literature, as well as some immersive environment studies [Billinghurst and Weghorst 1995]. Subjects were allowed to rotate their head, but were instructed not to move their head from side to side or back and forth, in order to reduce depth cues resulting from motion parallax. Subjects were instructed not to bend over to view their feet. Subjects viewed the environment until they felt they had obtained a good image of the target and the scene. Once a good image was achieved, subjects were instructed to close their eyes. In the real world environment the subjects then blindfolded themselves, where as in the immersive environment condition the HMD screen was cleared to black. Subjects then started walking toward the target location. All subjects were directed to walk purposefully and decisively to the target location and stop walking when they felt their feet were where the target had been observed. Instructing subjects to walk purposefully and decisively helped to remove their tendencies to second guess themselves. The same instructions were provided for each environment.

3.1 Equipment and Software

The equipment used in this experiment consisted of an n-Vision Dativisor HiRes head-mounted display with interlaced 1280x1024 resolution, full color, and a 52 degree diagonal field of view. The angular resolution of the HMD is on the order of 2 arc minutes per pixel. The display was configured with 100% stereo overlap




Figure 6: *Intended versus perceived distance. Comparison of results from both environments. The dotted line represents ideal performance.*

between the two eyes.

The computer generated environments were rendered on an SGI Onyx2 R12000 with two IR2 rendering pipelines. One rendering pipeline was used for each eye to provide stereopsis. In keeping with our previous experimental configuration, we used a fixed interpupillary distance of 6.5 cm for the software generating the immersive environment. We recorded the interpupillary distance for each subject and the average interpupillary distance over our subject pool was 6.22 cm, with lower and upper bounds of 5.25 cm and 7.0 cm respectively. However, it is important to note that precise control of effective inter-pupular distance in HMDs is difficult [Wann et al. 1995]. We lack precise control over image position and accommodation, and these adjustments cannot be determined by polling the HMD.

The NPR immersive environment ran at no less than 25 frames per second. The polygon count in this model was approximately 8,000. Our software uses an OpenGL silhouette drawing algorithm [Raskar and Cohen 1999], as well as Sense8's WorldToolKit libraries, and Nate Robins' OBJ model reader [Robins 1997]. The model was created in Alias|Wavefront's Maya Complete.

During the portion of the experiment conducted in the real world hallway, the subjects' positions and the target locations were recorded by one of the experimenters with a tape measure. In the portion of the experiment in which subjects wore the HMD, subject positions were recorded using an IS600-Mark2 Intersense tracker. Subjects were afforded head rotation, but translations did not update the rendering the subjects viewed.

4 Results

We analyzed the data using a 2 (environment) x 3 (distance) x 2 (sex) repeated measures ANOVA (analysis of variance), with environment and distance as within-subjects variables and sex as a between-subjects variable using statistics software package SPSS [SPSS 2002]. We found that there is a significant statistical difference ($p < 0.001$) in performance between the real world condition and the immersive virtual environment condition. A statisti-

cal significance of 0.001 means that there is a 0.1% chance that the effects we are seeing are due to random events. It is worth noting that most psychologists agree that such relative data comparisons are statistically significant if $p \leq .05$.

Figures 4, 5, and 6 graph the results of the perceived versus intended distance analysis of our direct walking task experiments to be roughly 66.0% of the intended distance in the feature edge immersive environment, and about 97.0% in the real world. The data for each distance is averaged across the participants. In Figures 4 and 5, error bars represent ± 1 standard error of the mean (SEM). The plot in Figure 6 provides a direct comparison of the data.

These results are expected, especially given the previous body of work in perception. Although 66% seems to be a low number, it does correspond well with previous findings in traditionally rendered immersive environments in which perceived distances were found to be 50% to 75% of the intended distances. It may have been difficult for subjects to obtain an accurate sense of depth because the target may not have appeared to be connected to the ground plane in the feature edge rendered environment. However, as shown in the error bars, subjects were quite consistent in judging distances and not merely guessing at the intended distances. The ability for subjects to perceive any depth in this environment was chiefly afforded by the linear perspective cues provided at the boundaries between walls, the floor, and the ceiling, as well as by the various features (e.g. doors) along the hallway itself. We would expect line drawings to do a good job in conveying relative distance, since linear perspective is such a powerful relative distance cue. The surprise is that they can convey absolute distance as well.

Studies are currently underway which investigate additions that could be made to the current feature edge rendering to improve spatial perception. In order to test these hypotheses, we conducted a pilot experiment using the feature edge rendered environment, augmented with a line texture on the ground plane. Figure 7 shows a screen capture of the modified immersive environment. The addition of these lines enhance occlusion cues by the way that the target occludes the ground texture, possibly tying the target to the ground plane. Additionally, this enhancement increases texture gradient and linear perspective cues.

Figure 8 shows the preliminary data we gathered on the average perceived versus intended distance in this task with only two distances (3.1m and 4.7m) over a much smaller subject pool. The same underlying experimental methodology was used to conduct the pilot study. The six subjects (3 female, 3 male) that participated in this pilot study did not participate in the first study.

The original study showed that subjects judged distances to be roughly 66% of the intended distance when the ground plane was not explicitly rendered. In this pilot study, subjects perceived distances to be on average 62% of the intended distance. However it should be stressed that this is a pilot study and the data is only preliminary. Whether or not including ground plane surface detail actually helps in distance judgments remains an open question until a more complete study is conducted.

Creating pilot studies is an important part of conducting user studies. We have found that we often spend more time conducting small pilot studies in which you can change the conditions, instructions, and basic variables of the experiment. In these pilot studies the subject numbers are kept small in order to reserve the subject pool for full experiments. The most important issue to note is that data gathered from these types of experiments can only indicate a possible future direction for experimentation and data gathering, and does not provide conclusive evidence for or against the hypothesis of the pilot study.




Figure 7: A screen shot of the NPR hallway, displaying feature edges and a line textured ground plane, with a scaled inset of a screen shot looking down at the target.

5 Conclusion

Our paper is the first to examine space perception in a functional, non-photorealistically rendered immersive environment. We present a study that quantitatively evaluates a functional, non-photorealistic rendering algorithm in a stereoscopic, immersive environment.

To date, very little perceptual research has been directed at NPR images. We are unaware of any studies indicating the degree to which NPR images are capable of conveying a veridical sense of spatial organization. It is desirable to objectively quantify measurements of task performance in these environments, but it should be noted that performing these experiments correctly is difficult and care should be taken to understand the factors involved. Casual experiments may not produce data which is meaningful. In conducting our research, extensive cross-collaboration with psychologists and computer vision researchers has been extremely important. Their suggestions provided insight into the factors involved in designing and running experiments, such as the number of subjects, number of trials per subject, and the extreme importance of correctly worded and consistent experimental instruction.

As we move away from photorealistic rendering, it is important to be sure that in a functional context the information we do provide faithfully communicates the desired information. In the past, NPR researchers have embraced interdisciplinary collaboration involving the artistic community, and now, NPR researchers must expand their interdisciplinary efforts by focusing on and understanding research into human perception. Through these efforts, we have a means to understand the effectiveness of NPR, especially in immersive environments.

6 Acknowledgments

The authors are grateful to William Thompson for supporting our research. We are also grateful for our subjects, Bruce Gooch, Bill Martin, and the members of the University of Utah computer graphics groups. Thanks to Natalie Sargent for helping us conduct the experiments. Special acknowledgment goes to Sarah Creem-Regehr, John Rieser, Herb Pick, and Jack Loomis for attempting to teach us about perceptual psychology. Thanks to Alias|Wavefront for their




Figure 8: Average intended versus perceived distance in feature-edge immersive environment with the addition of a line textured ground plane shown in Figure 7. Error bars represent ± 1 SEM. The dashed line represents ideal performance.

generous donation of Maya Complete, which we use to create the geometric models for our experiments. This material is based upon work supported by the National Science Foundation under Grants No. CDA-9623614 and No. 58500700. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsors.

References

- BILLINGHURST, M., AND WEGHORST, S. 1995. The use of sketch maps to measure cognitive maps of virtual environments. *Virtual Reality Annual International Symposium*, 40–47.
- CUTTING, J. E., AND VISHTON, P. M. 1995. Perceiving layout and knowing distance: The integration, relative potency and contextual use of different information about depth. In *Perceptual of Space and Motion*, W. Epstein and S. Rogers, Eds. Academic Press, New York, 69–117.
- FERWERDA, J. 2001. Hi-fi rendering. In *Perceptually Adaptive Graphic, Preconference Proceedings*, A. McNamara and C. O’Sullivan, Eds. ACM Campfire. <http://isg.cs.tcd.ie/gap/>.
- GOOCH, B., AND GOOCH, A. 2001. *Non-Photorealistic Rendering*. AK Peters Ltd, July.
- HEALEY, C. G., INTERRANTE, V., KREMERS, D., LAIDLAW, D., AND RHEINGANS, P. 2001. Nonphotorealistic rendering in scientific visualization. In *SIGGRAPH 2001 Course Notes Nonphotorealistic Rendering in Scientific Visualization*, C. G. Healey, Ed. ACM.
- INTERRANTE, V. 1996. *Illustrating Transparency: communicating the 3D shape of layered transparent surfaces via texture*. PhD thesis, University of North Carolina at Chapel Hill.

- KNAPP, J. M. 1999. *The visual perception of egocentric distance in virtual environments*. PhD thesis, University of California Santa Barbara.
- LAMPTON, D. R., McDONALD, D. P., SINGER, M., AND BLISS, J. P. 1995. Distance estimation in virtual environments. In *Human Factor and Ergonomics Society*, 1268–1272.
- LOOMIS, J. M., AND KNAPP, J. M. in press. Visual perception of egocentric distance in real and virtual environments. In *Virtual and Adaptive Environments*, L. J. Hettinger and M. W. Haas, Eds. Erlbaum, Hillsdale, NJ.
- LOOMIS, J. M., DA SILVA, J. A., FUJITA, N., AND FUKUSIMA, S. S. 1992. Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance* 18, 906–921.
- M.W. DIXON, M. WRAGA, D. P., 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.
- PALMER, S. E. 1999. *Vision Science: Photons to Phenomenology*. The MIT Press.
- 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.
- PHILBECK, J. W., LOOMIS, J. M., AND BEALL, A. C. 1997. Visually perceived location is an invariant in the control of action. *Perception and Psychophysics* 59, 4, 601–612.
- RASKAR, R., AND COHEN, M. 1999. Image Precision Silhouette Edges. In *Proc. 1999 ACM Symposium on Interactive 3D Graphics*.
- RIESER, J. J., ASHMEAD, D. H., TAYLOR, C. R., AND YOUNGQUIST, G. A. 1990. Visual perception and the guidance of locomotion without vision to previously seen targets. *Perception* 19, 675–689.
- ROBINS, N., 1997. Smooth Normal Generation with Preservation of Edges. www.pobox.com/~nate/smooth.html.
- ROGERS, S. 1995. Perceiving pictorial space. In *Perception of Space and Motion*. Academic Press, Inc, ch. 4.
- SPSS, 2002. <http://www.spss.com/>.
- 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.
- WILLEMSEN, P., AND GOOCH, A. 2002. Perceived distance in virtual environments: An experimental comparison between real, image-based, and traditional virtual environments. In *To appear in Proc. IEEE Virtual Reality*.
- WITMER, B. G., AND KLINE, P. B. 1998. Judging perceived and traversed distance in virtual environments. *Presence: Teleoperators and Virtual Environments* 7, 2 (April), 144–167.