

Visual Cues for Imminent Object Contact in Realistic Virtual Environments

Helen H. Hu
Computer Science Department
University of Utah

Amy A. Gooch
Computer Science Department
University of Utah

William B. Thompson
Computer Science Department
University of Utah

Brian E. Smits
Computer Science Department
University of Utah

John J. Rieser
Dept. of Psychology and Human Development
Vanderbilt University

Peter Shirley
Computer Science Department
University of Utah

Abstract

Distance judgments are difficult in current virtual environments, impacting their effectiveness in conveying spatial information. A situation in which this problem is apparent is when contact occurs while a user is manipulating objects. In particular, the computer graphics used to support current generation immersive interfaces does a poor job of providing the visual cues necessary to perceive that contact between objects is about to occur. This perception of imminent contact is very important in human motor control. Its absence prevents a sense of naturalness in interactive displays which allows for object manipulation. This paper reports results from an experiment evaluating the effectiveness of binocular stereo, cast shadows, and diffuse interreflections in signaling imminent contact in a manipulation task.

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

Keywords: Virtual Reality, Head Mounted Displays, Human Visual Perception

1 Introduction

An effective visualization system must accurately communicate shape and other spatial information to the user. While immersive systems can aid in this process, it is important that they do so in a natural and intuitive manner. The accurate and natural presentation of spatial information is particularly important in interfaces that allow a user to grasp and manipulate simulated objects. Grasping and placing of one object in contact with another involves fine motor control that is dependent on precise perceptual judgments. Such situations arise wherever objects are manipulated in a virtual space, as is done in virtual surgery and any visualization application on an immersive workbench that includes a strong “ground plane.” Only limited studies exist on the perception of quantitative spatial information in immersive interfaces. Almost nothing is yet known about the effectiveness of immersive interfaces in supporting a natural sense of grasping and contact.

Manually bringing one object in contact with another involves two distinct stages. A person intending to bring one object in contact with another is able to sense when contact is imminent and adjust his or her control of the manipulated object appropriately. The perceptual cues driving this action are necessarily purely visual. Actual contact is signaled perceptually by a combination of visual, haptic, and acoustic information. To see the importance of the visual cues for imminent contact, try placing an item such as a coffee mug on a table with your eyes closed. Even though the approximate location of the table is known, the lack of a visual indica-



Figure 1: Mechanism for allowing a subject to manipulate the virtual world. The task was to lower a virtual cylinder as close to a virtual table as possible without contacting the virtual table.

tion of approaching contact *prior* to the contact actually occurring makes the haptic sensation of actual contact more of a surprise than when the same task is performed with your eyes open.

This paper explores some of the visual cues likely to be effective in signaling imminent contact in immersive displays. To concentrate on visual effects, we have limited our investigation to the approach phases of contact and leave to other studies the even more difficult evaluation of the interaction between vision, haptics, and audition involved in perceiving actual contact. Perceptual judgments of spatial properties appear to be task-dependent. In particular, estimates of distance can vary between appearance-based judgments and action-based judgments. As a result, in this study we measure the effectiveness of visual cues for imminent contact using manual actions typical of those likely to occur when a user is manipulating objects.

In the physical world, grasping and object manipulation are typically graceful and accurate [17, 18, 23], whereas in existing virtual worlds they are imprecise and often awkward [7]. This imprecision is due to a lack of spatial information in the virtual displays. A major cause of the awkwardness is the difficulty people have in making precise distance judgments based on computer-generated imagery. Perspective, occlusion, shading, motion, and binocular stereo all contribute to a strong sense of spatial organization in rendered images [2, 9]. Despite the visually compelling nature of modern computer graphics, viewers still find it difficult to precisely judge ob-




Figure 2: Rendering geometry used in the experiment (with the light shown as a star). The cylinder had a radius of 3.9 cm and a height of 20 cm. It was placed 38.1 cm away from the subject in the immersive environment. Subjects were instructed to stand approximately the same distance away from the physical cylinder.

ject sizes, exocentric distances between different scene points, and egocentric distances from the observation point to points in the scene [14]. As an aid in understanding why this is so, this paper evaluates the effectiveness of binocular stereo, shadow, and diffuse interreflection cues on the precision with which people can manipulate a virtual object in a virtual environment.

Binocular stereo is part of most immersive displays and is a primary depth cue for the range of distances involved in object manipulation [19]. Shadows are known to be a cue for the closeness of an object to an extended surface [25], an effect that is heavily used in computer animation and video games. Interreflection is the process of light bouncing from one object to another and includes all illumination not directly from a light source. Recently, evidence is emerging that interreflection may also aid in visually judging contact [15, 16, 21]. In the remainder of this paper, we present the results of an experiment aimed at evaluating the effectiveness of binocular stereo, cast shadows, and interreflections in signaling imminent contact between an object and an extended surface towards which it is moving (an exocentric task).

2 Background

There is an extensive literature on depth perception (see Cutting and Vishton [2] for a current overview and Foley [4] for a more in-depth discussion of binocular stereo). Whereas numerous studies [8, 19, 29] have shown stereopsis to be a strong cue for distance judgments, relatively few studies have been performed on the role that shadows play in space perception [11, 25, 30]. Also there is almost no published research on how interreflection affects distance judgments, except work by Madison *et al.*. They found interreflections convey similarly strong information about depth information for object contact as shadows despite being considerably less visually prominent in scenes [15, 16, 21].

Only recently have computer scientists started to look at depth perception when viewing renderings of 3D environments (*e.g.*, [9, 12, 26]). Wanger *et al.* compared the effectiveness of six different depth cues for three simple tasks, finding that for each task a different combination of depth cues were important [26]. Kjell Dahl studied a different set of depth cues and found illumination and object placement were the most effective cues for conveying relative depth information [12]. However, both of these studies did not compare their depth cues to stereopsis.

Servos *et al.* examined only stereopsis and found binocular vision to be significantly important for distance judgments for a grasping task [19]. Hubona *et al.* compared the effectiveness of stereo, cast shadows, and background images for distance judgments for 3D inferencing [8]. However, these studies examined

depth cues for tasks other than the object manipulation task studied in this paper. The set of depth cues that are important for each task varies depending on the type of task involved [26].

Most importantly, it is unclear if any of these previous findings will hold in virtual environments, because depth perception in immersive environments is complicated by the limitations of rendering and display systems [27] and by the complexity of the tasks often expected of users. For this reason, efforts are underway by a number of groups to study depth perception within the context of virtual environments [3, 5, 14, 20, 24, 28] and the related problem of achieving a compelling sense of presence in such environments [22]. This work is in its infancy, and relatively little is yet known about how specific visual cues, rendered using immersive environment technology, contribute to performance in the sorts of tasks for which immersive environments are likely to be useful.

3 Methodology

Measuring the accuracy of depth perception via verbal reporting is not ideal. Subjects are often unaware of systematic bias in what they say they see compared to what they actually perceive. In addition, subjects often come to believe that they are inaccurate and change what they report to reflect beliefs. These subjects may contaminate the results of the study by modifying their verbal reports [6, 13]. As a result, we instead had subjects perform a task involving depth perception and measured the accuracy and precision of their task performance, rather than asking subjects to report numerical estimates of distances.

Participants stood at a physical table and manipulated a physical cylinder while viewing a virtual cylinder which moved in a manner coupled to the real cylinder. Movement of the real cylinder was constrained along a vertical path of fixed location (see Figure 1 and video). The height of the cylinder was sensed using a SensAble Technologies PHANToM 1.5, which has a nominal spatial resolution of 0.03mm. The PHANToM was used purely as a position sensing device. No force feedback was presented to subjects if they placed the virtual cylinder in contact with the virtual table. The height of the physical cylinder was used to adjust the height of the rendered cylinder viewed through a head-mounted display (HMD), such that that the virtual cylinder was displayed in the immersive environment at the same position in space as the physical cylinder. Positional drift in the PHANToM was not a problem for this experiment because the virtual table, the virtual cylinder, and the distance between the two were all determined by the coordinate data from the PHANToM. (In other words, positional drift affected both the virtual table and the virtual cylinder equally, so the distance between the two remained constant despite drift.)




Figure 3: Sample images for the cases enumerated in Table 1. Interreflections are visible in (a) and (c) at the base of the cylinder on the table. See the accompanying video for more example images.

Geometry and shading on the cylinder was rendered interactively using an SGI Onyx2 R12000 with two IR2 rendering pipelines. The illumination on the table top was computed using a standard Monte Carlo path tracer [10]. (See Figure 2 for the rendering geometry.) The illumination was stored in a 3D texture map, where u, v corresponded to the width and depth of the table, and w was a function of the height of the cylinder above the table. We used a non-linear mapping for the height in order to have better resolution when the cylinder is near the table. At display time, the illumination maps were blended with the surface color to achieve a correctly lit table top. The table top was textured with a checkerboard to strengthen the binocular stereo cue. (To prevent the checkerboard from providing a 2D positional visual cue to the user, it was translated randomly on the table for each trial.)

Images were presented using an n-vision Datavisor HiRes head-mounted display with 1280 by 1024 full color resolution, a field of view of 52 degrees, and an angular resolution on the order of 2 arc minutes per pixel. The display was configured with 100% stereo overlap between the two eyes. The n-vision HMD uses CRT technology, which avoids a number of visual artifacts in LCD-based displays that distract from visual realism. Since we wanted to concentrate on a limited number of visual cues, head tracking was not used so as to eliminate depth information obtained from motion parallax [1] or viewing strategies using user-chosen viewpoints. Thus the displayed scene did not change with subject movement, only with cylinder movement.

Subjects were presented with multiple trials in which they were instructed to lower the virtual cylinder as close to the virtual table as possible without actually making contact. Individual trials began with the cylinder high in the visual field and well above the appar-

shadows and interreflection, in stereo
shadows and interreflection, without stereo
shadows only, in stereo
shadows only, without stereo
interreflection only, in stereo
interreflection only, without stereo
neither shadows nor interreflection, in stereo
neither shadows nor interreflection, without stereo

Table 1: Different combinations of visual cues examined in the experiment.

ent height of the table. The subjects were expected to bring the cylinder down to the table and back up again within three seconds. If they did not complete the task within the allotted time, an alarm sounded and that trial was recorded as incomplete. (Subjects were not informed that the data from these trials were ignored.) This timing restriction was added to encourage subjects to bring the cylinder down to the table and back up to its original height in one fluid and natural motion. This methodology also avoided any ambiguity if subjects had been unable to stabilize the cylinder at the closest position to the table.

The goal of the experiment was to compare the precision with which subjects adjusted the cylinder, across variations in the stimulus information that specified the cylinder and table top in the virtual world. Eight stimulus conditions resulted from crossing the presence and absence of binocular stereo cues, shadow cues, and

Visual Cue Combination	Average Accuracy	Standard Deviation
stereo, shadow, interreflections	4.1mm	3.5
stereo, shadows	4.4mm	3.6
stereo, interreflections	5.6mm	4.9
stereo	6.2mm	5.5
shadows, interreflections	10.4mm	9.3
shadows	12.6mm	11.8
interreflections	13.3mm	12.1
no cues	11.6mm	10.1

Table 2: Mean stopping distances and standard deviations above the table, pooled for all subjects. These values were especially susceptible to subject bias.

interreflection cues (see Table 1 and Figure 3). Each of the participants completed 80 repeated trials in each of the eight stimulus conditions for a per subject total of 640 trials. The experiment took about 40 minutes per subject, with short pauses every 40 trials and one long break after 320 trials. After each pause, there were five additional trials that we ignored, to allow the subjects to regain their rhythm. The trials were randomly ordered with respect to the four conditions involving presence or absence of shadows and presence or absence of interreflections. Every four trials we also alternated between presentations utilizing binocular disparity and presentations where the binocular disparity was set to zero, although from post-experiment questioning, we found subjects were not aware of any absence of binocular disparity in the trials. The weight of the cylinder and the positioning of the apparatus with respect to the subjects were chosen to minimize muscle fatigue effects.

To control for motor memory (and a variety of other possibly confounding effects), the height of the virtual table was varied randomly between trials, and the checkerboard pattern on the virtual table moved randomly in two dimensions from trial to trial. To prevent subjects from viewing the table movements, the virtual environment was replaced by a single sphere floating in an empty world between each trial during the table movement. The sphere allowed for continuous stereoscopic vision. Subjects were instructed to fixate on the bottom of the cylinder as they lowered it towards the table to make sure they attended to the visual cues of interest to the study.

4 Experimental results

The main goal of the study was to compare how precisely participants could adjust the cylinder’s position as a function of the presence and absence of the different cues. The results presented here came from four subjects, one of whom is a co-author of this paper. The other three were unfamiliar with the goals and methods of the study. All four subjects were naive users of immersive environments. Both pooled data and individual subject performance were examined.

We had data from a total of 320 trials for each of the eight cue combinations, for a grand total of 2560 trials. 2.9% of these trials were marked invalid because the subject did not finish the trial within the allotted time frame. Exceeding the time limit happened three times as frequently when binocular disparity was absent, a result consistent with reports that binocular disparity reduces the time needed to perform grasping tasks [19].

As indicated in the introduction, the visual perception of imminent contact is quite different from the visual/haptic/acoustic perception of actual contact. For the purposes of this experiment, we limited our analysis to those situations in which subjects were suc-




Figure 4: Variable error, pooled for all subjects, shown with standard error bars. Error is lowest when binocular stereo information is available. Even without stereo, shadows and possibly interreflection help in performing the task.

cessful in placing the cylinder close to the table but not making contact. This occurred in 28.2% of the trials in which the time limit was not exceeded. We are not sure if this relatively low percentage is the result of subject bias caused by attempts to be as accurate as possible or some systematic effect. While the experimental setup provided no haptic feedback, visual cues for the interpenetration of the cylinder and table were present and included the virtual cylinder shrinking and the shadow (if present) ceasing to change position or shape. These effects limit the utility of the data associated with trials in which the cylinder was lowered too far.

For above the table results (Table 2), subjects stopped closer to the table when binocular disparity was present than without. However it is unclear how important accuracy is because it is so susceptible to bias by the subject. For example, if a subject is more conservative about not contacting the table for a given visual cue, that subject will stop the cylinder higher on average for that cue, and the corresponding accuracy error will be worse on average.

Therefore, to better compare the effectiveness of binocular disparity, shadows, and interreflections at signaling contact between the cylinder and the table, we examined subject precision for each cue combination. By precision, we mean the amount of variability from the mean. We used two different measures to evaluate the precision of their adjustments. One measure consisted of the variable error scores, which were determined by computing the standard deviation of the errors for each trial set for each subject in each condition (Figure 4). These scores directly reflect how consistently subjects placed the cylinder relative to the virtual table top surface across their many repeated trials in a condition. The variable errors reflect how precisely the subjects can control their actions relative




Figure 5: Correlation between actual table height and minimal height of cylinder, given for each subject. (Each circle or square corresponds to one subject's data.) Correlations were higher when binocular stereo cues were present than when they were absent for all four subjects. Correlations were also higher when shadows and/or interreflections were present for some subjects.

to a visual situation. The variable error scores themselves varied as a function of the stimulus conditions. Since the motor control side of the cylinder placement task were the same across all of the stimulus conditions, the effects of stimulus condition on task could only reflect variations in how precisely the subjects were able to use the different visual cues that were available in the different conditions.

The second measure was the simple correlation of the variations in the heights of the virtual table with variations in their adjustments. For this measure, we evaluated each subject's data separately to find the difference between how our subjects used the visual cues. Correlations were computed across the repeated trials in each condition for each subject (Figure 5). High correlations indicate high precision, but they do not indicate the absolute distances involved.

The experiment was designed with 80 repeated trials per condition so that each subject's data could be individually analyzed. In each condition, variable error scores were computed for each consecutive set of five trials. This resulted in 16 variable errors per condition for each of the four subjects. To find out whether the precision of the subjects adjustments varied as a function of the presence and absence of the binocular stereo, shadow, and interreflection cues, the variable error scores were submitted to a $2 \times 2 \times 2 \times 4$ (stereo by shadows by interreflections by subject) analysis of variance with repeated measures on the first three factors.

Figure 5 shows the correlations which were computed across the variations in the virtual table height and the adjustments for each subject's 80 trials in each stimulus condition. As can be seen in the figure, the correlations were higher when binocular stereo cues

were present than when they were absent for all four subjects, and there is more modest variation in the correlation across the variations in shadows and interreflections. Note that the correlations were somewhat above zero even in the condition where there was no stereo, no shadows, and no interreflections. The reason for this is there are additional cues available to subjects when the virtual cylinder goes below the surface of the virtual table, as previously explained.

The mean variable error scores, averaged across the scores of all four subjects, appear in Figure 4. (The standard error bars appear with the means.) The results of the analysis of variance showed significant main effects of stereo ($p < .001$), of shadows ($p = .001$), and of interreflections ($p = .003$), where $(1-p)$ represents the probability that the subjects were responding to the different types of stimulus cue. Thus, overall the subjects were able to use the presence of all three cues to improve the precision of their adjustments of the cylinder in the virtual world.

The data for individual subjects were also analyzed with $2 \times 2 \times 2$ (stereo by shadow by interreflection) analysis of variance with repeated measures on all three factors. These analyses showed that all four participants used the stereo cues to improve their precision, but they differed in the degree to which they used the shadows and the interreflections. One subject clearly benefited from the shadows, whereas the others did not to an effect that was statistically significant. A different subject showed a modest benefit from the interreflections.

5 Discussion

Without haptic feedback, contact between a manipulated object and other objects or surfaces in current immersive environments is unnatural. Even with haptic feedback, contact often comes as a unexpected surprise. It is therefore of significant importance to explore ways of improving the visual cues for imminent contact in such systems. We have shown the importance of binocular stereo in one sort of manipulation task common to many virtual environment applications. Adding shadows further helps some individuals in the task. Interreflections may provide some additional benefit but this is still an open question which warrants more investigation. Without binocular stereo, shadows and interreflections provide little information for judging that contact is about to occur, at least with this configuration of object geometry and lighting.

It is important to avoid over-generalizing from these results. These results hold for a particular configuration of object shapes, surface markings, and lighting. Different conditions may change the effectiveness of different visual cues. The results are also likely to be dependent on display conditions. In our case, we utilized a head-mounted display with the current best commercially-available image quality and stereo resolution. Other display devices may do less well in presenting either stereo or shading information. On the other hand, the resolution of the display is still far less than that associated with observing physical objects under natural viewing conditions.

Clearly, much additional work is required in performing controlled, quantitative examinations of perceptual accuracy in virtual environment displays. We plan on examining other object geometries and other lighting arrangements. We believe by conducting more experiments across the design space, we will be able to determine whether subjects are performing this placement task by measuring the distance between themselves and the table or by measuring the distance between the table and the cylinder: two very different mechanisms.

6 Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. CDA-96-23614. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Scene models were created using Maya, generously donated by Alias|Wavefront.

We thank Herb Pick for his advice in designing the experiment, and Greg Coombe, Bruce Gooch, Bill Martin, Simon Premoze, and David Weinstein for contributing suggestions to the pilot experiment.

References

- [1] A. C. Beall, J. M. Loomis, J. M. Philbeck, and T. J. Fikes. Absolute motion parallax weakly determines visual scale in real and virtual environments. In *Proc. Conference on Human Vision, Visual Processing, and Digital Displays*, pages 288–297, Bellingham, WA, 1995. SPIE.
- [2] James E. Cutting and Peter M. Viston. Perceiving layout and knowing distance: The integration, relative potency, and contextual use of different information about depth. In William Epstein and Sheena Rogers, editors, *Perceptual of Space and Motion*, pages 69–117. Academic Press, New York, 1995.
- [3] Stephen R. Ellis and Brian M. Menges. Judgments of the distance to nearby virtual objects: Interaction of viewing conditions and accommodative demand. *Presence*, 6(4):452–460, August 1997.
- [4] John M. Foley. Binocular distance perception. *Psychological Review*, 87(5):411–434, 1980.
- [5] Sergio S. Fukusima, Jack M. Loomis, and José A. DaSilva. Visual perception of egocentric distance as assessed by triangulation. *Journal of Experimental Psychology: Human Perception and Performance*, 23(1):86–100, 1997.
- [6] Melvyn A. Goodale and Angela Haffenden. Frames of reference for perception and action in the human visual system. *Neuroscience & Biobehavioral Reviews*, 22(2):161–172, 1998.
- [7] J. M. Hollerbach, W. B. Thompson, and P. Shirley. The convergence of robotics, vision, and computer graphics for user interaction. *International Journal of Robotics Research*, 18:1088–1100, November 1999.
- [8] Geoffrey S. Hubona, Philip N. Wheeler, Gregory W. Shirah, and Matthew Brandt. The relative contributions of stereo, lighting, and background scenes in promoting 3D depth visualization. *ACM Transactions on Computer-Human Interaction*, 6(3):214–242, September 1999.
- [9] Victoria Interrante. *Perceiving and Representing Shape and Depth*. PhD thesis, UNC-Chapel Hill, Department of Computer Science, 1996.
- [10] J. T. Kajiya. The rendering equation. In David C. Evans and Russell J. Athay, editors, *Computer Graphics (SIGGRAPH '86 Proceedings)*, volume 20, pages 143–150, August 1986.
- [11] D.K. Kersten, P. Mamassian, and D.C. Knill. Moving cast shadows induce apparent motion in depth. *Perception*, 26(2):171–192, 1997.
- [12] Lars Kjelldahl and Martin Prime. A study on how depth perception is affected by different presentation methods of 3D objects on a 2D display. *Computers & Graphics*, 19(2):199–202, March 1995.
- [13] Jack M. Loomis, José A. DaSilva, Naofumi Fujita, and Sergio S. Fukusima. Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance*, 18(4):906–921, 1992.
- [14] Jack M. Loomis and Joshua M. Knapp. Visual perception of egocentric distance in real and virtual environments. In L.J. Hettlinger and M.W. Haas, editors, *Virtual and Adaptive Environments*. Erlbaum, Hillsdale, NJ, in press.
- [15] Cindee Madison and Daniel Kersten. Use of interreflection and shadow for surface contact (poster). In *Annual Meeting of the Association for Research in Vision and Ophthalmology (ARVO)*, 1999.
- [16] Cindee Madison, Daniel J. Kersten, William B. Thompson, Peter Shirley, and Brian Smits. The use of subtle illumination cues for human judgement of spatial layout. Technical Report UUCS-99-001, Computer Science Department, University of Utah, January 1999. <http://www.cs.utah.edu/vissim/bibliography/papers/madison-99.pdf>.
- [17] D. Rosenbaum. *Human Motor Control*. Academic Press, San Diego, 1991.
- [18] D. Rosenbaum, L. Loukopoulos, R. Meulenbroek, and J. Vaughan. Planning reaches by evaluating stored postures. *Psychological Review*, 102:28–67, 1995.
- [19] Philip Servos, Melvyn A. Goodale, and Lorna S. Jakobson. The role of binocular vision in prehension: A kinematic analysis. *Vision Research*, 32(8):1513–1521, 1992.
- [20] R.T. Surdick, E.T. Davis, R.A. King, and L.F. Hodges. The perception of distance in simulated visual displays: A comparison of the effectiveness and accuracy of multiple depth cues across viewing distances. *Presence: Teleoperators and Virtual Environments*, 6(5):513–531, 1997.
- [21] William B. Thompson, Peter Shirley, Brian Smits, Daniel J. Kersten, and Cindee Madison. Visual glue. Technical Report UUCS-98-007, Computer Science Department, University of Utah, March 1998. <http://www.cs.utah.edu/vissim/papers/glue/glue.html>.
- [22] Martin Usoh, Kevin Arthur, Mary C. Whitton, Rui Bastos, Anthony Steed, Mel Slater, and Jr. Frederick P. Brooks. Walking > walking-in-place > flying in virtual environments. In *Proc. SIGGRAPH 99*, pages 359–364, August 1999.
- [23] C. von Hofsten, H. L. Pick, Jr., and G. Neeley. The role of binocular information in ball catching. *Journal of Motor Behavior*, 24, 1992.
- [24] Thomas E. von Wiegand, David W. Schloerb, and W.L. Sachtler. Virtual workbench: Near-field virtual environment system with applications. *Presence: Teleoperators and Virtual Environments*, 8(5):492–519, October 1999.
- [25] Leonard Wanger. The effect of shadow quality on the perception of spatial relationships in computer generated imagery. *Symposium on Interactive 3D Graphics*, 25(2):39–42, March 1992.
- [26] Leonard R. Wanger, James A. Ferwerda, and Donald P. Greenberg. Perceiving spatial relationship in computer-generated images. *IEEE Computer Graphics & Applications*, pages 44–58, May 1992.
- [27] John P. Wann, Simon Rushton, and Mark Mon-Williams. Natural problems for stereoscopic depth perception in virtual environments. *Vision Research*, 35(19):2731–2736, 1995.
- [28] Bob G. Witmer and Paul B. Kline. Judging perceived and traversed distance in virtual environments. *Presence*, 7(2):144–167, April 1998.
- [29] Y. Yeh. Visual and perceptual issues in stereoscopic color displays. In *Stereo Computer Graphics: and Other True 3D Technologies*, pages 50–70. Princeton University Press, 1993.
- [30] Albert Yonas, Lynn T. Goldsmith, and Janet L. Hallstrom. Development of sensitivity to information provided by cast shadows in pictures. *Perception*, 7:333–341, 1978.