

QUANTIFYING THE RELATIVE ROLES OF SHADOWS, STEREOPSIS, AND FOCAL ACCOMMODATION IN 3D VISUALIZATION

Mike Bailey
San Diego Supercomputer Center

Thomas Rebotier
David Kirsh
Cognitive Psychology
University of California San Diego

ABSTRACT

The goal of three-dimensional visualization is to present information in such a way that the viewer suspends disbelief and uses the screen imagery the same way as he or she would use an identical, real 3D scene. To do this effectively, programmers employ a variety of 3D depth cues. Our own anecdotal experience says that shadows and stereopsis are two of the best for visualization. The nice thing is that both of these are possible to do in interactive programs. They sacrifice a certain amount of interactive speed, but they are possible.

But, there is very little information detailing *exactly* what these cues add to the perception process. The purpose of this project was to *quantify* how worthwhile using these two depth cues are, that is, is it worth losing interactivity to get them? Using a large number of student subjects, we performed a series of depth-test trials and analyzed the results. Finally, as an upper-bound control on these experiments, we also ran subject trials on physically fabricated 3D objects, viewing them through a pinhole in a controlled-lighting situation to factor out both shadows and stereopsis, leaving only focal accommodation. This paper shows the design of the experiments and the results expressed in reaction times and error rates. The results have a significant bearing on the design of 3D interactive visualization systems, particularly those that use virtual or augmented reality.

KEY WORDS

Perception, three dimensions, shadows, stereopsis, virtual reality, augmented reality.

1. INTRODUCTION

Several factors contribute to depth perception, including stereopsis, interposition, shadowing, perspective and texture perspective ([1]). Most of these factors are properties of the monocular or binocular visual images. Only two cues, vergence and accommodation, are informed by the state of ocular muscles, and, of these two, only accommodation is monocular. Few previous studies address, without methodological problems, the issue of shadows, stereopsis, and accommodation on depth perception ([2]; [3]); these studies focus on the estimation of absolute distance, and according to the results of these studies, the accuracy of distance estimate should not be good enough to help figure out differences of depth in ordinary objects held at arm distance.

This project was motivated by UCSD's Center for Visualization Prototypes project ([4], [5], [6], [7]) in which physical models are being fabricated as a form of visualization hardcopy. Anecdotally, we had observed that when scientists were presented with a physical model of something that they had been studying using 3D graphics, they noticed new features in their model even before they touched it. This was especially true with geometrically complex models, such as protein structures (Figure 1). This phenomenon suggested to us that something about the appearance of the physical models was conveying more information than even a shadowed stereographics view on a computer screen.




Figure 1: Is viewing a physical model better than viewing a shaded, shadowed, stereo image on a computer graphics screen?

2. UNDERSTANDING THE ROLES OF SHADOWS AND STEREOPSIS IN VISUALIZATION DEPTH PERCEPTION

2.1 Introduction

Shadows and stereopsis are popular depth cues for better understanding 3D scenes. Interactive stereographics has been around for some time, starting with the vertical "stretching" of half framebuffers, and then moving into the use of quad buffers. Stereographics is currently experiencing a resurgence of sorts due to the emergence of dual-projector "GeoWalls".

Shadowing was once solely the domain of software rendering, but has also experienced an interactive resurgence due to the speed and features on modern graphics accelerator cards (e.g., [8]).

But, while stereopsis and shadows are both possible on graphics hardware, each causes performance degradation. For this part of the project, we wanted to quantify what the perceptual benefits were of using stereopsis and shadows, so that a visualization developer could decide, using more than anecdotal evidence, how much perceptual benefit is worth the loss of interactive performance.

2.2 Methods

To compare the contributions of shadowing and stereoscopic vision, we built a stereoscopic viewer allowing two different 1600x1200 pixel images to be seen by the left and the right eye (Figure 2). In this viewer, we presented views of protruding plots, one of which was taller than the others, and one of which was marked by a red dot (Figure 3). Participants had to determine if the taller plot (the one protruding closer to them) was the one wearing the dot, or not. Half of the views were true stereoscopic views, with two different images, and half were showing the same image to both eyes. Orthogonally, half the views comported shadows, and half did not. Finally, the marking was borne by the tallest plot half of the time (again, orthogonally to the other two factors).




Figure 2: The experimental set-up, seen from above (distances are not to scale). The subject head was stabilized by a chin. The chamber containing the monitors was enclosed by black matte boards. The same computer controlled both monitors.




Figure 3: A stereo pair stimulus. Stereoscopic perception can be achieved by crossing one's eyes. The (x,y) positions of both plots are random within a circle encompassing 80% of the screen height, and rejecting stimuli with overlapping plots. Each image is 1600x1200 pixels, the plots have a flat top 100 pixels in diameter.

One hundred and thirteen college students, 18 to 30 years old, participated in the experiment for class credit. Each participant first took a practice with a few trials, then the main experiment which consisted of 80 trials. The stimuli were presented on two 21" hi-resolution (1600x1200) monitors in a closed chamber with matte black walls, and reflected on mirrors (normal mirrors, not first surface, but the reflection presented no visible aberration). The participant controlled the onset. All participants had normal or corrected-to-normal vision and responded with their right hand by pressing one key on a response box if the marked plot was the taller and another key if the taller plot was not the taller. Trials were randomized by the presentation software, "SuperLab". The main experiment fully crossed the following factors:

- Mono/Stereo
- Shadowed stimulus/ Bland stimulus (no shadows)
- Marking: taller plot Marked or Unmarked.

2.3 Results

We measured error rates and reaction times. We only report on Error Rates here; the RTs confirm the pattern but the statistics are more telling with Error Rates. The most sensitive result comes from the largest difference in height, $dZ=6$ pixels. (Figure 4) The effects of Stereopsis and Shadowing are comparable: an improvement of 8.5% for stereoscopic views, and of 9.4% for shadowed images. Both of these main effects are significant, for Stereopsis $F(1,26)=17.1$, $p<.001$ and for Shadowing $F(1,26)=9.9$, $p<.001$.

Surprisingly, there is also a strong bias to see the marked plot as the taller one, with a 22% difference, $F(1,26)=26.4$, $p<.001$, and this bias interacts with shadowing, so that unmarked views show a marginally significant ($F(1,26)=4.4$, $p<.05$) better improvement with shadows.

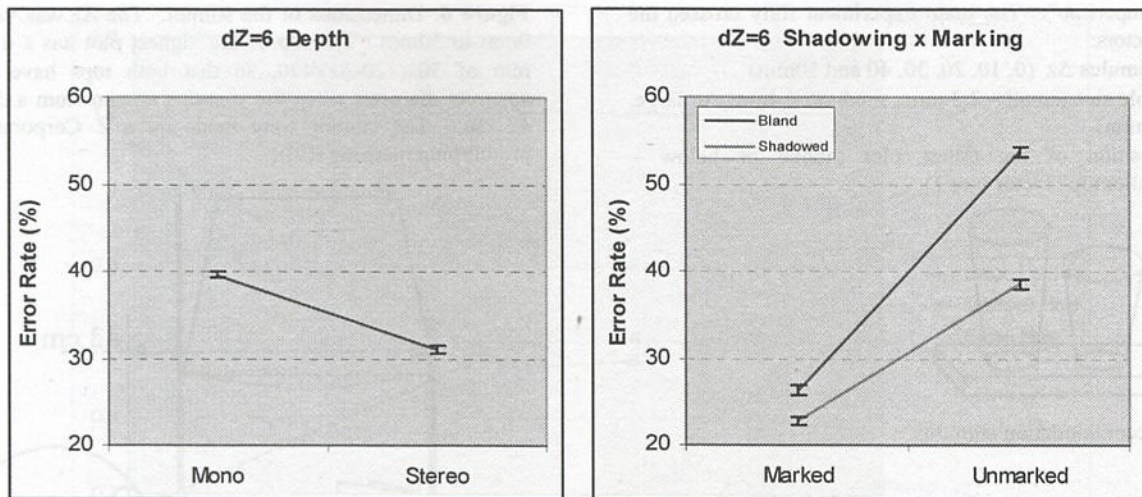


Figure 4: Error rate, in percent, for a height difference of 6 pixels. Chance level is 50%. Stereopsis, Shadowing, Marking, and to a lesser measure Shadowing X Marking are significant factors. Error bars show the 5% confidence interval on each data point.

For smaller differences the effects are reduced, and are not always significant. Table 1 summarizes all effects. Stereopsis and shadowing present the same order of magnitude.

dZ	Participants	Marking	Stereopsis	Shadowing	Sh.X.Mrk
0 (control)	34	17***	-3.4	0.6	***
2	11	23**	0.5	1.8	/
4	41	16***	6.8**	2.8	/
6	27	22***	9.4***	8.5***	***

Table 1: Effect on the Error rate, in percent, for a various height differences between the taller and the other plot(s). Chance level is 50%. *: $p < .05$, **: $p < .01$, ***: $p < .001$. For the Shadowing X Marking interaction we only indicate significance.

3. UNDERSTANDING THE ROLE OF ACCOMMODATION IN VISUALIZATION DEPTH PERCEPTION

3.1 Introduction

To study the possibility that accommodation helps understand a visualization object's shape, we built a monocular apparatus to view real three-dimensional objects. These stimuli present to the subject protruding "towers" that, seen from the top in the apparatus, present the same apparent area (see Figures 5-7). To the first order, with monocular viewing, only accommodation can help detecting which plot comes closer to the viewer's eye. Not happy with a first order, we varied the size of the hole through which the participants could look at the stimuli. By using holes smaller than the pupil, the depth-of-field is reduced, and the effect accommodation is reduced.

There were two experiment predictions. The first order effect had to be that the larger the height difference between both plots, the easier it would be for participants to detect which plot came closer to their eye. The second order effect was that as we reduced the size of the viewing hole, the first-order effect should diminish. Both predictions were validated by the results.

3.2 Methods

Forty-nine college students, 18 to 30 years old, 35 males and 14 females, participated in the experiment for class credit. Each participant first took a practice of 12 trials, using pilot models that had a larger delta-Z than the experimental models (respectively, 60mm and 80mm), then the main experiment that consisted of 36 trials. The stimuli were presented in a closed chamber with matte black walls, and lit by diffuse sideways light from two computer monitors (see Figure 7). The board behind the stimuli was also matte black, so that the sideways light could cast no visible shadow that could have clued to the height of the plots. The onset of lighting was controlled by the same computer which registered the participant's response. All participants had normal or corrected-to-normal vision, used their right eye for the experiment, and responded with their right hand by pressing one key on a response box if the taller plot was above and another key if the taller plot was below. In between trials, the participant rolled their chair back and spun around while the experimenter opened the chamber and changed the stimulus. All precautions were taken so that in between trials the participants could not see, hear, or deduce the particular stimulus used and the position of the plots. In each trial, the participant would get his head and eye in position, turn on the light by pressing any key, and respond by pressing the first or second key. Trials were randomized by the presentation

software, "SuperLab". The main experiment fully crossed the following factors:

- Stimulus Δz (0, 10, 20, 30, 40 and 50mm)
- Hole size (small - 2.3 mm, medium - 4mm, or large, 10mm)
- Position of the tallest plot (above or below - "tallontop" factor 0 or 1)




Figure 5: Accommodation stimulus





Figure 6: Dimensions of the stimuli. The Δz was varied from 0mm to 50mm. The top of the highest plot has a diameter in mm of $30 \times (420 - \Delta z) / 420$, so that both tops have the same apparent diameter when the stimulus is seen from a distance of 42 cm. The stimuli were made on a Z Corporation rapid prototyping machine ([9]).




Figure 7: The experimental set-up, seen from above. The subject head was stabilized by a chin rest and subjects were specifically instructed not to wiggle their head during viewing. The stimulus is shown in place. The chamber containing the stimulus was entirely closed by black foam/cardboard boards, preventing external light to reveal the stimulus ahead of time. The front wall was removable and adjusted to position either the large hole, the medium hole, or the small hole.

3.3 Results

We measured error rates and reaction times. Because RTs were highly variable between subject (from a median of 3 seconds to 12 seconds) we analyzed the *rank* of RT (from 1 to 36). The error rates present the predicted pattern but are very noisy (see Figure 8). In contrast, the RT ranks present much cleaner results (see Figure 8); the stimulus Δz has a clear effect on response RT, and that effect diminishes for the medium hole and entirely disappears for the small hole.

To test the significance of the results presented in Figure 8, we computed separately for each participant and each hole the regression coefficients of the RT rank to the Δz and "tallontop". (Figure 9) The 49 coefficients for Δz of the large hole (mean - 0.20) and small hole (mean +0.01) were submitted to a two-sample t-test and were shown significant: $t(96) = -3.88$ two-tailed $p < .0002$.

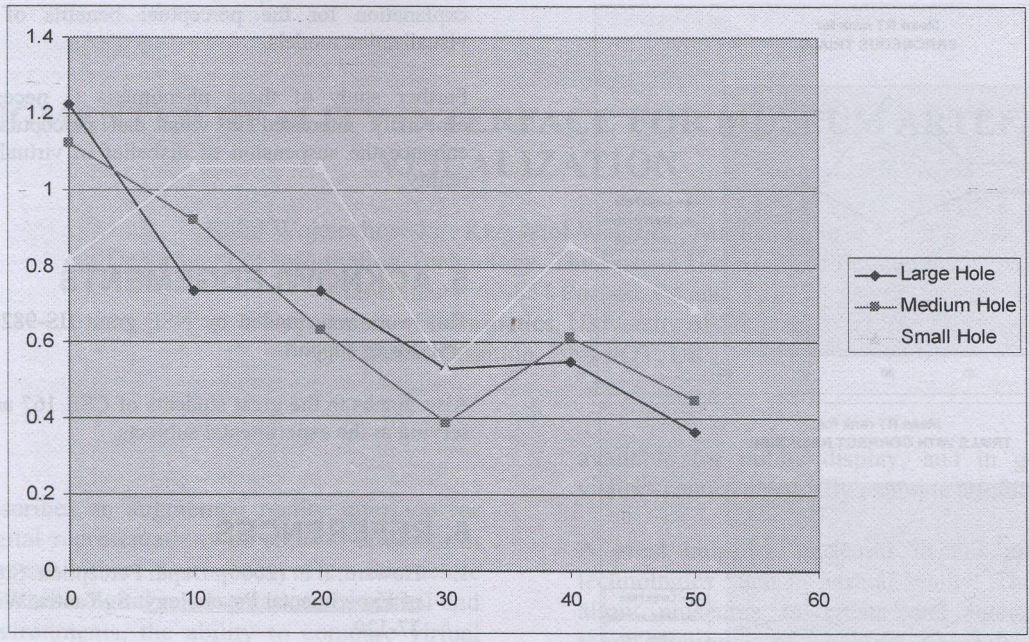
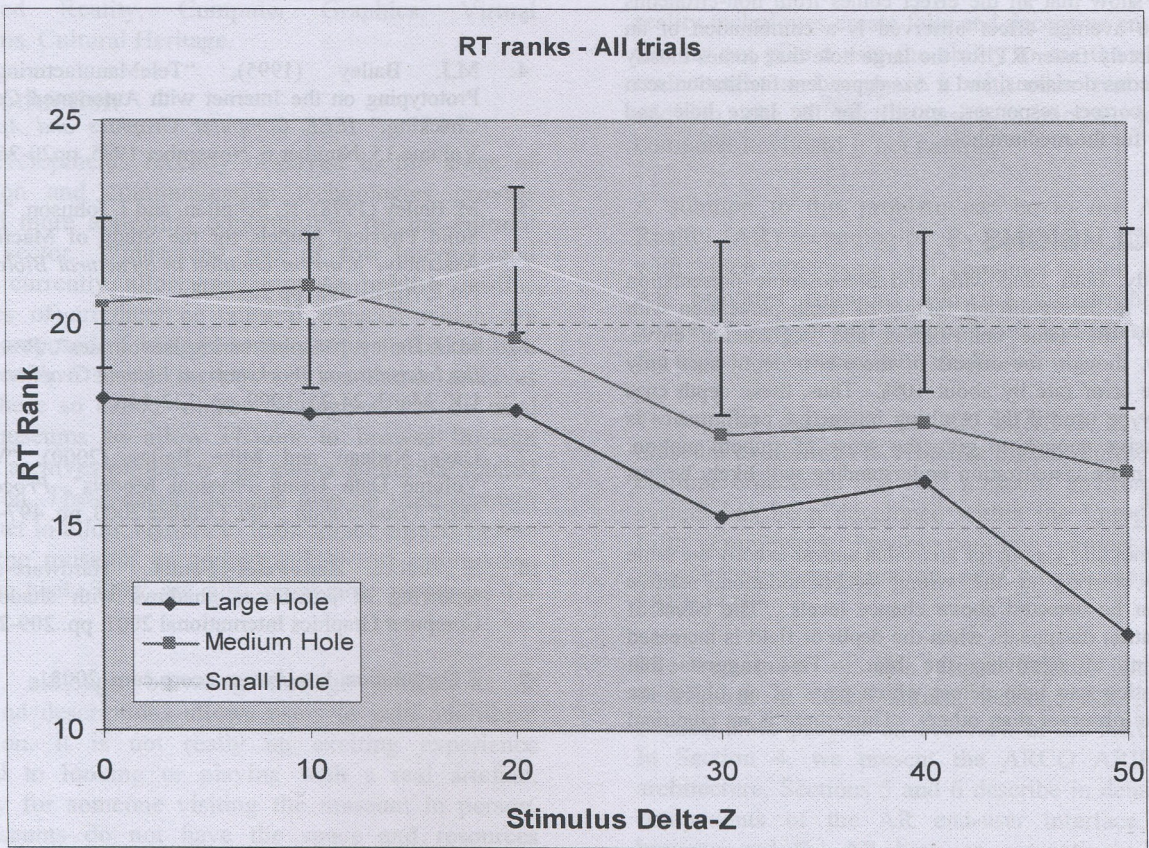


Figure 8: Error Rate results. Chance level is at 1.0. Chance decreases significantly with stimulus Δz for the large hole (mean slope -0.0144) more than for the small hole (mean slope -0.005); a t-test on the individual coefficients shows $t(96)=-3.21$, $p<0.002$. Errors for $\Delta z = 0$ are conventional: one plot of this stimulus was conventionally defined as being taller than the other, in order to provide a baseline.



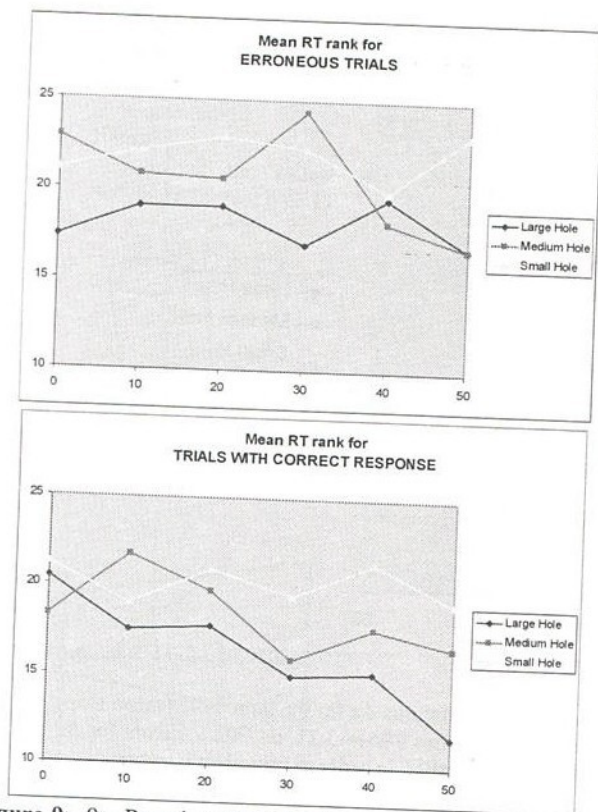


Figure 9: 9a. Reaction Time Ranks, as a function of stimulus. 9b and 9c show that all the effect comes from non-erroneous trials. The average effect observed is a combination of an indiscriminately faster RT for the large hole that comes mostly from erroneous decisions, and a Δz -dependent facilitation seen only with correct responses, mostly for the large hole and moderately for the medium hole.

4. CONCLUSIONS

In this study, both Shadowing and Stereoscopic presentation were found to be significant factors of depth perception with substantially the same contribution and degradation curve. Surprisingly, though, the effects of these two depth cues only reduced the error rate by about 10%. Thus, these depth cues should only be used if the resulting interactive performance is still responsive enough to give the sense of smooth motion. Otherwise, more visualization understanding will likely be lost than gained.

Even in monocular vision, for areas that appear exactly the same except for blurring at the wrong accommodation, relative position can be detected above chance levels. The effect of accommodation disappears when the depth of field is increased to the point of removing the blur. This suggests that accommodation can help detect which parts of an object are closer to the observer than others. Thus, there is an empirical

explanation for the perceptual benefits of using physical visualization models.

Further study of these phenomena is necessary. We are especially interested in what can perceptually be done to enhance the suspension of disbelief in virtual and augmented reality.

5. ACKNOWLEDGEMENTS

This work was funded by NSF grant IIS-9820594. We thank NSF for its support.

Also thanks to the great students of CSE 167 and MAE 152 for serving as the experimental subjects.

6. REFERENCES

- Howard, I. P. (2000). Depth Perception. *Steven's Handbook of Experimental Psychology*. S. Yantis, Wiley and Sons. 1: 77-120.
- Fisher, S. K. and K. J. Ciuffreda (1988). "Accommodation and apparent distance." *Perception* 17: 609-621.
- Mon-Williams, M. and J. R. Tresilian (1999). "Some recent studies on the extraretinal contribution to distance perception." *Perception* 28: 167-181.
- M.J. Bailey (1995). "TeleManufacturing: Rapid Prototyping on the Internet with Automated Consistency-Checking." *IEEE Computer Graphics and Applications*, Volume 15, Number 6, November 1995, pp20-26.
- M. Bailey (1998), K. Schulten, and J. Johnson. "The Use of Solid Physical Models for the Study of Macromolecular Assembly." *Current Opinion in Structural Biology*, Vol 8, No 2, April 1998, pp 202-208.
- M. J. Bailey, "Manufacturing Isovolumes", *Proceedings of the International Workshop on Volume Graphics*, Swansea, UK, March 24-25, 1999, pp. 133-146.
- Dave Nadeau and Mike Bailey (2000), "Visualizing Volume Data Using Physical Models", *Proceedings of IEEE Visualization 2000*, October 2000, pp. 497-500.
- S. Brabec, Hans-Peter Seidel., "Hardware-accelerated rendering of antialiased shadows with shadow maps", *Computer Graphics International 2001*, pp. 209-214, 2001
- Z Corporation, <http://www.zcorp.com>, 2003.