Effects of Rendering on Shape Perception in Automobile Design

James A. Ferwerda*, Stephen H. Westin*, Randall C. Smith†, Richard Pawlicki†
*Program of Computer Graphics, Cornell University
†General Motors R&D

Abstract

The goal of this project was to determine if advanced rendering methods such as global illumination allow more accurate discrimination of shape differences than standard rendering methods such as OpenGL.

To address these questions, we conducted two psychophysical experiments to measure observers' sensitivity to shape differences between a physical model and rendered images of the model. Two results stand out:

- The rendering method used has a significant effect on the ability to discriminate shape. In particular, under the conditions tested, global illumination rendering improves sensitivity to shape differences.
- Further, viewpoint appears to have an effect on the ability to discriminate shape. In most of the cases studied, sensitivity to small shape variations was poorer when the rendering and model viewpoints were different.

The results of this work have important implications for our understanding of human shape perception and for the development of rendering tools for computer-aided design.

CR Categories: I.3.7 [Three-Dimensional Graphics and Realism]

Keywords: rendering, perception, shape, fidelity

1 Introduction

The goal of design visualization is prediction: to show the design exactly as it will appear when built. Existing CAD visualization tools have proven their value, but have also shown some important limitations. Sometimes the visualizations produce "surprises", where the final physical object looks different than it did in the computer rendering.

Many commercially available modeling and rendering tools have been developed for the entertainment industry, where artistic skill is applied to convince viewers that the image on the screen is real. These tools typically simplify the image rendering process by using ad hoc models of local light reflection and ignoring global illumination effects within the environment.

For many years, advanced rendering methods that more accurately simulate light reflection and transport have been available, but these methods are only useful to the design community if they improve the fidelity of the visualizations with respect to the final physical artifacts. For the design of automobiles and other complex objects, shape perception is an important issue, and a critical question is whether advanced rendering methods allow more accurate perception of object shape than standard methods.

To address this question we conducted a pair of psychophysical experiments in which observers compared rendered images of an object to its real physical counterpart, and we tested whether they were able to make finer discriminations of shape with advanced renderings than with standard renderings. Our primary goal was to determine if advanced rendering methods allow designers to make more accurate visual evaluations of their designs.

The paper is organized into the following sections. Section 2 surveys related work in this area. Section 3 describes the overall design of the experiments and the methods used for generating the visual stimuli. Sections 4 and 5 describe the experiments themselves, and the results are analyzed in section 6. Section 7 summarizes the conclusions that can be drawn from the experiments and Section 8 points toward future work.

2 Related Work

Perceiving the shapes of objects is one of the central functions of vision (see [Palmer99] for an introductory review and [Interrante98] for a more comprehensive survey). The visual system uses many sources of information to accomplish this task including occlusion, perspective, shading, shadows, texture, motion, and disparity. Of these, surface shading by itself is thought to be a relatively weak source of metric information about surface shape [Todd83, Erens93]. However it has also been suggested that the non-metric cues provided by shading, in combination with other cues can provide powerful information for shape perception and discrimination [Todd89].

The pattern of shading observed on a curved surface depends not only on its shape, but also upon its material properties, the surrounding environment, and the observer's viewpoint. The specularities produced by shiny materials appear to provide useful information for shape perception [Norman95, Blake91], though under some restricted conditions shape and material properties may be confounded [Ramachandran88]. The visual system also appears to be able to take advantage of the shading patterns provided by both diffuse and glossy interreflections in solving the shape perception problem [Christou96, Norman04].

The pictorial representation of shape from shading is an important topic that has been studied for centuries [Kubovy86, Miller98], and much of the modern research in the area of shape perception has employed shaded computer graphics images. Therefore, it is surprising that relatively little attention has been paid to the issue of how the shading methods in rendering affect shape perception ([Rodger00] is an exception), and only recently have researchers started to use global illumination algorithms rather than standard local shading methods in their studies [Pellacini00, Madison01, Fleming03]. While this is a promising trend, to our knowledge no direct comparisons of global and local algorithms with respect to shape perception have been conducted.

Comparison of the appearance of real and rendered objects is also a topic that has a long history [Hagen80, Ellis91], and recently computer graphics researchers have started conducting studies to validate the fidelity of renderings with respect to the real world [Meyer86, Rushmeier95, McNamara00], however the focus on these studies has been reflectance and illumination perception per se, rather than how these shading factors affect shape appearance. Therefore we believe our current studies make a unique contribution to this literature.

3 Design of the Experiments

Our experiments required two kinds of visual stimuli: a physical test object and rendered images of it. There were two driving requirements for these stimuli: the images had to be faithful representations of the physical model, and the shape of the rendered model had to be systematically variable. In the following sections we will first describe how we selected the physical test object, then we discuss the processes used to generate the rendered images.

3.1 Test object

Since the driving application for this work is automobile design, the most obvious choice would have been to use a real car as a test object. Unfortunately this was impossible, not only because of lab space limitations, but also because we did not have access to the geometric data.

Another possibility would have been to use a simple generic shape such as a perturbed sphere or superquadric. However, such a shape is geometrically abstract, and we were concerned that the results might not generalize well to real objects. In addition, there would have been the problem of fabricating a physical model to match the geometric model.

We chose a compromise: a generic car-like shape called a *frog*. The frog is a glass-fiber shell about 2 feet long used in the auto industry as a reference shape for evaluating different paint finishes. While it is much simpler in shape than a real automobile and manageable in size, it embodies some of the complexity that is found on a real car, including convex, concave, and saddle curvatures. We based the mathematical model for rendering on a 3D scan of a physical frog, constructing surfaces in Alias® AutoStudio™ to match.

3.2 Shape Variations

Just as important as deciding on a base shape was defining the shape variations that would be used in the experiment. It was important to hold the object silhouette constant so that differences in shading would be the only cue subjects could use to distinguish the variations. We chose an area on the side of the body that displayed subtle, but visible, curvature. We created six variations of the base model, moving the surface further outward in each.

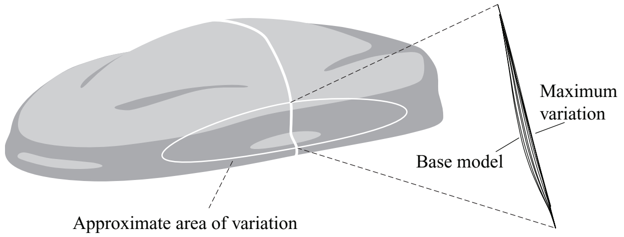

Figure 1: range of shape variations used in the experiments.

Figure 1 indicates the area where variation was concentrated, and shows cross-sections through the seven different surfaces. Figure 2 shows renderings of the base model (shape 0), the midpoint of the variation range (shape 3), and the maximum variation (shape 6).

3.3 View Variations

Since objects are often observed from different viewpoints, we also wanted to determine if viewpoint differences affect shape discrimination. To do this we rendered each shape variation from three different viewpoints by rotating the model. Figure 3 shows renderings from the three viewpoints used: left, center, and right. The center view is nominally a projective match to the view of the physical model used in the experiments.

3.4 Viewing Environment

In order to make a direct comparison between the physical model and the renderings, we needed to match not only the test object itself, but also its surrounding environment, because this determines surface illumination and reflections. The environment had to be realizable both physically and as a mathematical model. Since cars are typically viewed outdoors, we constructed a simple outdoor scene. Figure 4 shows the physical setup.

We built a closed box (approximately 5’ on each side) with a painted blue interior to simulate a clear sky. We used high-resolution texture images of tiles and bricks printed on a large-format inkjet printer to give natural reflections from the floors and walls. The environment was intentionally rather simple, since every detail had to be modeled in both the physical world for the renderings, but it did contain enough detail to be plausible as a real environment.

The most severe restriction on the environment came from our need to match luminances between the physical environment and the computer display. A sunlit outdoor scene can have a dynamic

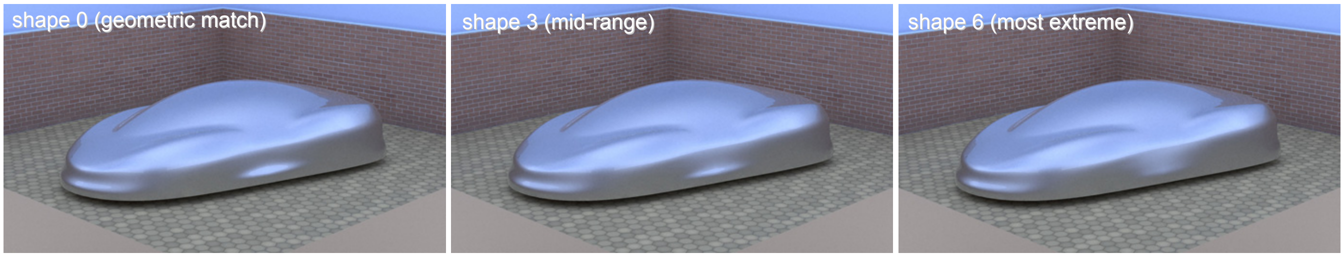

Figure 2: Range of shape variation used in the experiments.

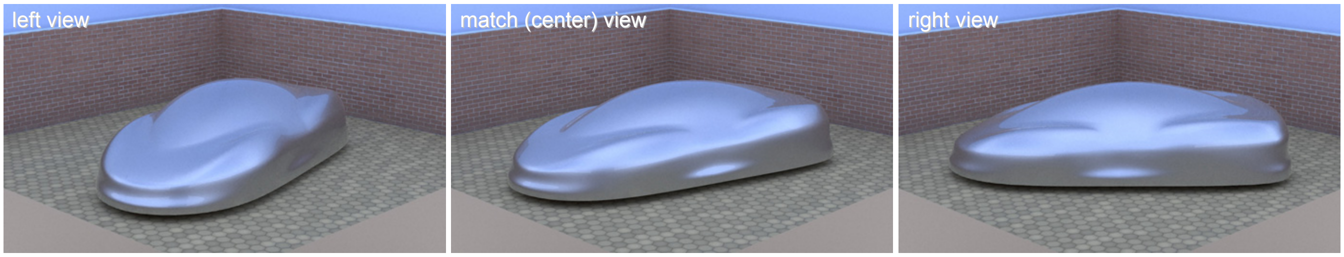

Figure 3: Viewpoints used in the experiments.

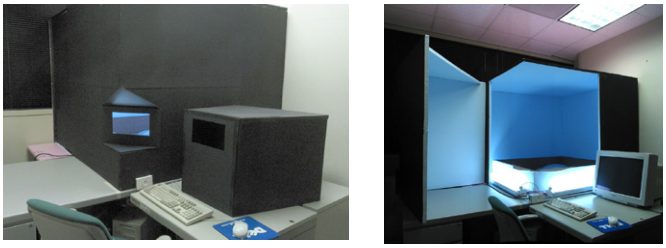

Figure 4: Experimental setup (L). With apertures removed (R).

range exceeding 10,000:1, but conventional displays can only reproduce about a 100:1 range. To resolve this mismatch, we used completely indirect lighting that limited the range of luminances to the range that could be reproduced by our display (a Sony GDM-W900 CRT). As a bonus, this also simulated the "desert before sunrise" look favored by automotive photographers, illustrators, and designers.

To match colors between the physical environment and the rendered images, we used two standard fluorescent strip lights with high color rendering index and a 5000K correlated color temperature, and switched our monitor to a 5000K white point. The fluorescents were dimmable so that we could visually match brightness of the physical environment to that of the monitor.

We wanted the subjects to view the physical model and the renderings at the same size and from a similar angle, so we placed the environment and the monitor on a tabletop and had the subjects view the physical model and the computer display through two identical apertures. Figure 5 shows both the principle involved and the view seen by the subjects. The apparent tilt of the images is a projective distortion caused by the wide angle of the photograph. Both apertures were physically upright.

3.5 Rendering methods

To make the global illumination renderings, we used our in-house research software based on the Metropolis method [Veach97]. Reflectance was simulated with the Ward light reflection model [Ward92], fitting it to reflectance measurements of a sample of the same paint used on the physical model of the frog. Each image was rendered at 1600x900 pixels and 24 bits.

We made corresponding OpenGL renderings for comparison, using lighting and shading parameters that matched the environment and surface finish as well as possible. Shadows were computed using a depth-map shadow method, and an environment map was used for reflections.

3.6 Resulting Images

Figures 6a and 6b show the two types of computer rendering used in this experiment. Figure 6a is an OpenGL rendering, representing the quality of rendering available in many design packages. Figure 6b shows the same geometry rendered using global illumination and an accurate reflectance model. Figure 6c shows a photograph of the physical model.

4 Experiments

The question we seek to answer through our experiments is whether advanced computer graphics rendering techniques such as global illumination allow for more accurate discrimination of shape differences than standard rendering methods such as those available through OpenGL. We are also interested in how an additional factor, viewpoint, affects our ability to discriminate shape. The experiments and statistical analyses presented in the following sections were designed to address these questions.

4.1 Stimuli

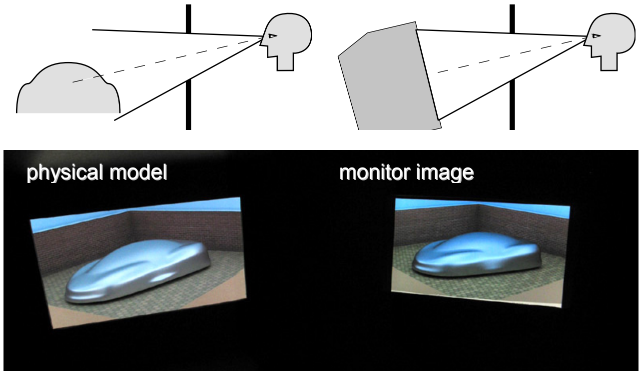

Figure 5: Viewing geometry: physical model, monitor image.

The stimuli for the experiments consisted of six sets of computer-generated images of the object shown in Figure 1. The sets were defined by all combinations of rendering method (Open GL, global illumination) and viewpoint (left, match, right). Each rendering/view set consisted of a sequence of seven images showing a progressive change in the shape of the physical frog model. Figure 2 shows the range of shape variation in each sequence, where shape 0 was a geometric match to the physical model, while shape 6 represented the greatest geometric difference. Figure 3 shows the three views (left, match, right) used. The match view was a projective match to the subject's view of the physical model in the experimental setup.

5 Ranking Study

Prior to the main experiment, we conducted a ranking study. Our goals were twofold. First we wanted to familiarize the subjects with the character and range of the images we were using. Second, we wanted to confirm that we had chosen an appropriate range of shape variations for study.

5.1 Procedure

Subjects sat in front of the table shown in Figure 7. On each trial they were presented with one of the six rendering/view image sets. At the top of the table was a "standard" image of the model (shape 0, global illumination rendering, match view).

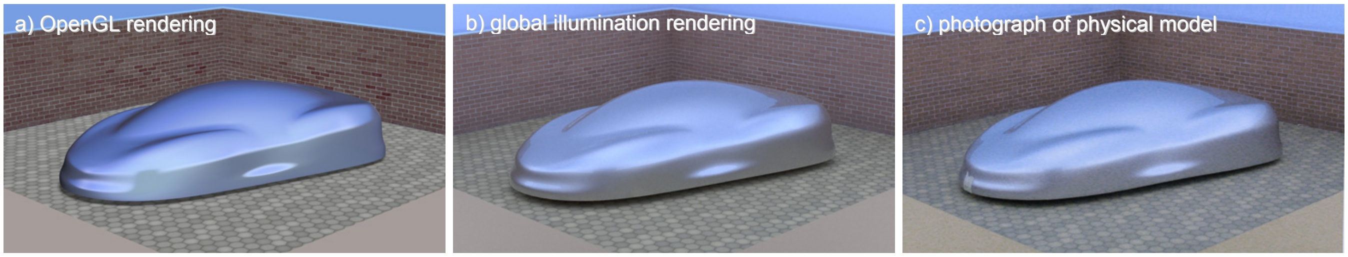

Figure 6: Renderings used in the experiments (a,b) with a photograph of the physical model (c) for reference.

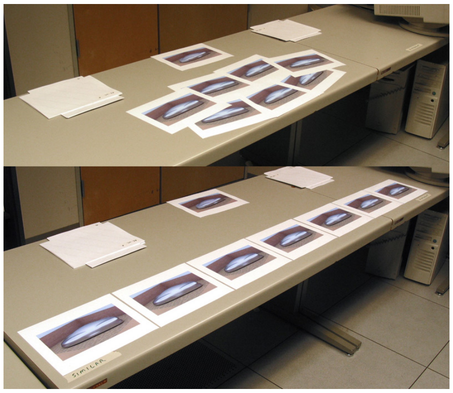

Figure 7: Ranking study procedure (top) before, (bottom) after.

Initially the images were presented in a random pile as shown in the top part of Figure 7. The subject's task was to place the images in order of similarity with respect to the standard (shown in the bottom of Figure 7). The same procedure was used for each image set, and the presentation order of the sets was randomized from subject to subject.

Ten subjects participated in the study. The subjects were staff and graduate students in the Cornell Program of Computer Graphics, were all naïve to the purpose and design of the experiment. However they all had significant experience viewing images and discriminating image differences, so their performance is probably most representative of a professional population.

5.2 Results and Analysis

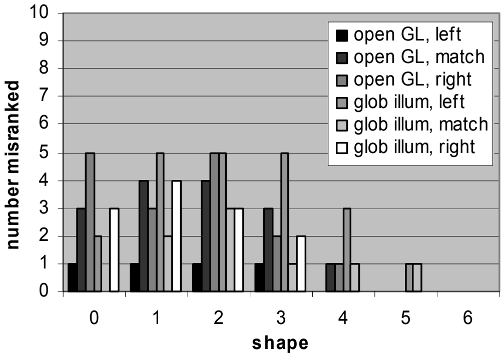

Figure 8: Ranking study results.

Figure 8 summarizes the results of the ranking study. Each set of colored bars corresponds to one of the six rendering/view combinations presented. The horizontal axis in each graph indicates the shape variation tested (with 0 being the standard), and the vertical axis indicates the number of times each of these shapes was misranked (for example, a bar for shape 0 shows the number of times it was placed in the 1st, 2nd, etc. position in the subjects' rankings). There are several things to notice:

- First, the rankings are not perfect. Subjects made errors in ordering the shapes with respect to their true geometric variations.
- Second, most of the errors occurred for the smaller variations in shape (shapes 0-4). Larger shape differences were typically ranked in the correct order.
- Finally, there were differences in the patterns of misrankings among the different rendering/view combinations indicating that these factors made the task more or less difficult.

There are several insights that can be drawn from this study.

- First, the study confirms that the range of shape variation we are using is sufficient and that the shape differences are indeed visible.
- Second, the fact that there are generally more misrankings for the smaller shape variations indicates that these differences are near the threshold for shape discrimination, which is likely the region of interest in the design process.
- Finally, the variance in ranking performance across different rendering/view combinations show that these parameters do in fact affect the visual information available for judging shape, and therefore these are appropriate variables to be studying.

While the results of this study are only preliminary, they give us confidence that the stimuli and methods we are using are appropriate for more quantitative study.

6 Rating Experiment

As stated earlier, our goal was to determine whether advanced rendering methods such as global illumination allow for more accurate discrimination of shape differences than standard rendering methods such as OpenGL. We are also interested in how an additional factor: viewpoint, affects our ability to discriminate shape. To accomplish this we ran a second experiment in which we compared renderings and a real physical model. The setup used in this experiment was previously described in Section 3.

6.1 Procedure

The experiment used a graphical rating procedure. Subjects saw a series of images, and for each image they were asked to rate how different the shape shown in the image was with respect to the shape of the physical model. For each image, subjects indicated their responses using a mouse to adjust an onscreen slider. The movement of the slider was continuous and the ends of the range were labeled "identical" and "most different". Slider position was coded into a numerical rating on a 0 to 100 scale. The images were the same ones used in the ranking study and consisted of a total of 42 images representing all combinations of the 7 shape variations, 2 rendering methods, and 3 views. The images were presented two times each in random order. Trials were under subject control through the use of an onscreen "next" button. The same 10 subjects that participated in the ranking study participated in this experiment.

6.2 Results: mean ratings

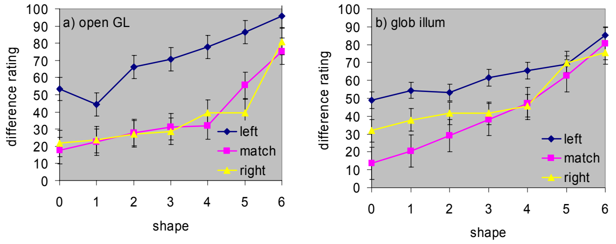

Figure 9: Mean rating curves.

Figure 9 summarizes the results of the rating experiment. The curves in the two graphs show the mean ratings given to each of the seven shape variations as a function of rendering method and view. Figure 9a shows results for the OpenGL renderings and Figure 9b shows results for the global illumination renderings.

Although the trends shown are based on descriptive statistics (means), and need to be confirmed with inferential methods, several interesting observations can be made.

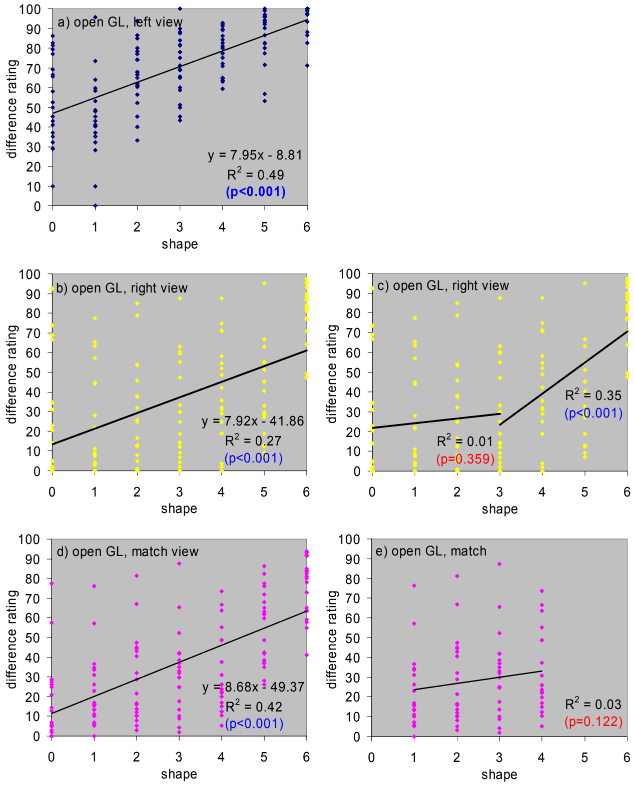

Figure 10: Full and sub-range regressions, OpenGL renderings

- First, overall, all the curves have positive slopes. This suggests that subjects are in fact able to do the task with some ability, and means that larger physical differences receive systematically higher ratings.
- Second, there are absolute differences in the ranges of ratings given to different views, with the range for the left view being highest, the match view lowest, and the right view somewhere in between. This suggests that there may be overall differences in how similar the shapes rendered in the different views appeared with respect to the physical model
- Finally, it appears that there are differences in the slopes of the curves across both view and rendering method. If there are slope differences this is interesting, because the slopes of the curves are an indicator of sensitivity to shape variations. A steeper slope means that the different shapes are given a greater range of ratings, which suggests that the subjects are able to see greater differences between them.

6.3 Analysis: regressions

Although the trends seen in Figure 9 are interesting, before we can draw any conclusions, we need to examine the results more thoroughly using inferential statistics. The methods that we will use are linear regression and analysis of covariance. We will examine each of the rendering/view combinations in turn.

6.3.1 Open GL, left view

We start with a regression analysis of the data from the OpenGL sequences beginning with the left view. Figure 10a shows a scatterplot of the ratings given by different subjects for each of the shape variations. The first thing to notice is that there is substantial variance in the ratings given by the subjects for each shape. While studying individual differences in the perception of shape is a worthwhile goal, it is beyond the scope of this project. Therefore, to allow group analysis of the ratings we normalized each subject's ratings by their overall minimum and maximum responses. This effectively compensated for differences in use of the slider range without destroying the internal consistency of each subject's ratings. Also, recall that each subject rated each shape twice. One option would have been to average these ratings before analysis, however given the variance in the data this seemed inappropriate, so we included both ratings directly in the regressions. If anything, this should make it more difficult to achieve significance in the statistical tests. This also explains why there are 20 data points for each shape in the scatterplots.

Figure 10a also shows a regression line fit to the scatterplot data for the OpenGL/left view condition and the inset summarizes: 1) the equation of this line; 2) R^2, the coefficient of determination; and 3) the p-value indicating the statistical significance of a test for a non-zero slope to the regression line (in the figures that follow, blue and red text will be used to indicate significant and non-significant results of statistical tests). There are several useful pieces of information that can be gleaned from this analysis.

- First, the small p-value ($p<0.001$) for the regression indicates that there is a statistically significant positive trend in the data. This means that under these conditions (OpenGL rendering, left view) larger geometric differences are reliably rated as being more visibly different from the physical model, and that these renderings do in fact convey useful information for shape discrimination.
- Second however, the 0.49 R^2 value indicates that only half of the variation in ratings can be explained by this model and that there is clearly a large amount of variation that is due to other factors. Familiarity/learning, and/or inter-subject differences might be two possibilities.

6.3.2 Open GL, right view

Figure 10b shows a regression analysis of the data from the OpenGL, right view condition. As before, the small p-value indicates that there is a statistically significant positive trend in the overall range of the ratings, but the low R^2 and the kinked appearance of the mean rating curve (Figure 9a), suggests that a simple linear function might not be an appropriate model for this data.

Figure 10c shows regressions on the lower and upper halves of the shape variation ranges. Two things can be seen from these analyses.

- First, the slope of the regression line fit to the upper half of the range is very steep and statistically significant ($p<0.001$).This means that over this range, subjects are able to make clear discriminations among the different shapes.
- Second and more importantly however, the large p-value ($p=0.359$) for the regression line fit to the lower half of the shape range means that the slope of the line over this range is not significantly different from zero, which indicates that the subjects are not reliably able to see differences between these small shape variations under these conditions.

What we are seeing in this data is a threshold effect. Below a certain magnitude of shape variation it is very hard to see differences, while above this level the differences are clearly visible.

6.3.3 Open GL, match view

Finally, Figure 10d shows a regression analysis of the data from the OpenGL/match view condition. The overall fit is shown on the left. Here both the p-value and the R^2 indicate a statistically significant positive trend in the ratings, but again the kinked shape of the corresponding mean rating curve in Figure 9a suggests the

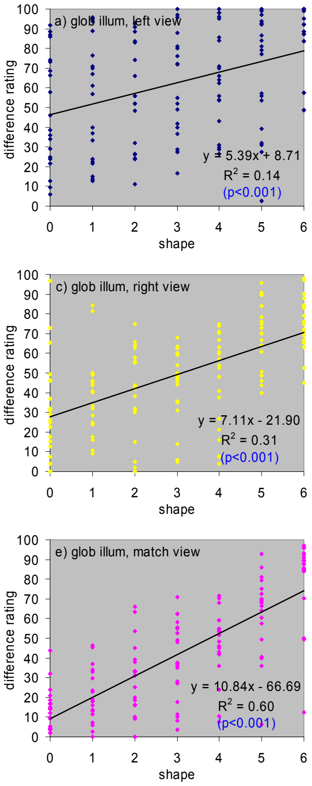

Figure 11: Regressions, global illumination renderings

need for further investigation. Figure 10e shows a sub-range regression on this dataset. Here the non-significant slope of the regression line fit to the lower-middle shape variation range means that subjects are not able to reliably discriminate these shapes under these conditions, and that the significance of the regression line fit to the full dataset is only produced by the extremes of the range (0 and 4-6). Similar to previous results for the right view, the analysis suggests that discrimination is below threshold for the lower part of the range, except for shape 7 where the shape of the rendered model is most similar to the physical model.

6.3.4 Global illumination, left and right views

We now turn our attention to analysis of the ratings for the global illumination renderings. The results of regression analyses for the left and right view conditions are summarized in Figures 11a,b and 11c,d respectively. The patterns of results here are similar to the threshold pattern we already saw in the analyses of the OpenGL renderings. For small shape variations (range 0-3) the slopes of the regression lines fit to the data (left panels) are not significantly different than zero, indicating poor discrimination, while for larger variations (range 3-6, right panels) the slopes are positive and significant, indicating a reliable ability to accurately discriminate shape differences.

Overall in this experiment, three of the four mismatched view conditions have shown this threshold-like pattern, and the fourth (OpenGL, left view), while showing a statistically significant positive trend, has a relatively shallow slope which indicates relatively low sensitivity to shape variation.

Taken together these results suggest that mismatches between the object and image viewpoints may make it difficult to judge small to moderate shape differences.

6.3.5 Global illumination, match view

To complete our analyses of the data from the rating study we now look at the global illumination/match view condition. The results of the regression analysis are shown in Figure 11e. There are several interesting results of this analysis.

- First, the small p-value ($p<0.001$) indicates that the positive trend is statistically significant. Of particular importance is the fact that unlike the other curves that could be broken down into two segments with different slopes, here analyses of shape subranges did not yield better fits to the data, suggesting that a simple linear model is indeed appropriate.
- Second, the R^2 (0.60) for this fit is the largest value seen in this experiment, suggesting that the unexplained variance in ratings is smaller under these conditions possibly because the discriminations are easier and therefore subjects' ratings are more orderly.
- Finally, the slope of the regression line (10.84) is the highest seen in the experiment, which means that under these conditions subjects are using a greater portion of the response range in making their ratings. Again this suggests that these conditions make it easier to see the differences in shape, and this is reflected by the expansion of the range of ratings. In simpler words, under these conditions, the most similar rendered shape (0) looks more similar to the physical model, and the most different rendered shape (6) looks more different.

A tantalizing hypothesis is that these particular rendering and view conditions (global illumination, match view) maximize apparent shape contrast making it easier to accurately discriminate shape differences. This hypothesis is directly related to the original purpose of this project: to evaluate whether advanced rendering techniques improve shape discrimination. The analysis presented in the following section will examine this hypothesis in greater detail.

6.4 Quantifying the effects of rendering

The previous analyses have shown that some of the positive trends seen in the mean rating data in Figures 9 are statistically significant, which indicates that these rendering/view combinations are providing reliable visual information for the discrimination of shape. But a question more to the point is whether these trends are statistically different from each other. In other words, are subjects more sensitive to shape differences when some combinations are used than others? To determine this we can test to see if the slopes of the regression lines fit to the different rendering/view data sets are significantly different from one another. The statistical procedure used is called an analysis of covariance (ANCOVA).

In principle, we could do this analysis on any of the rendering/view combinations, but we have already seen that there is a negative effect of viewpoint differences on sensitivity so we decided to eliminate cross-viewpoint comparisons from the analysis and focus on the match view conditions. The basic question we are trying to answer is whether there is a difference in how well the OpenGL and global illumination renderings allow us to discriminate shape. If so, this will be reflected by a significant difference in the slopes of the regression lines fit to each data set.

6.4.1 Analysis of covariance

The results of the analysis of covariance on the OpenGL and global illumination match view conditions are shown in Figure 12. Figures 12a and 12b show the mean rating data for these two conditions, and 12c shows confidence interval plots produced by

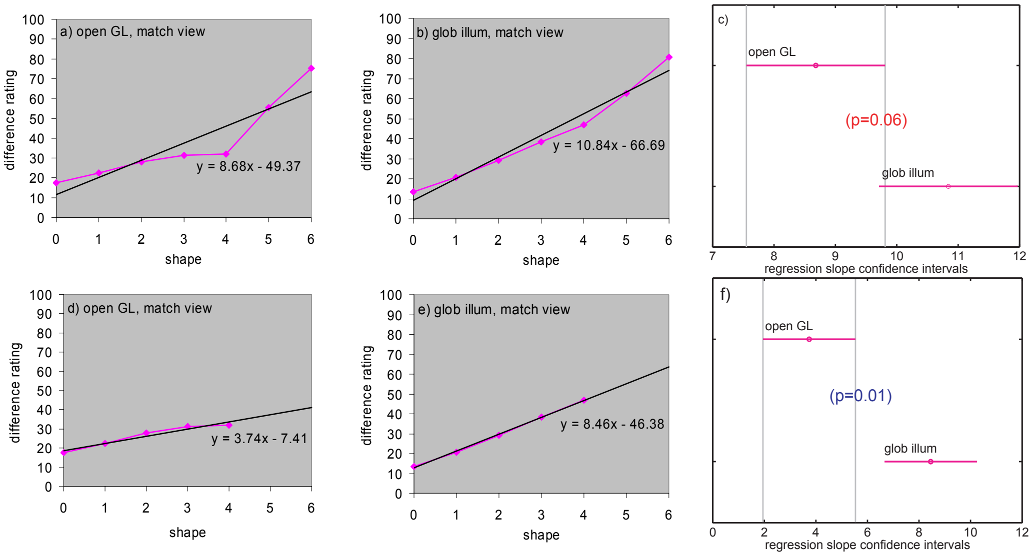

Figure 12: Analysis of covariance. Full (top) and sub-range (bottom) regressions.

the analysis of covariance, that indicate whether the slopes of the regression lines fit to the underlying datasets are significantly different from one another. The overlapping ranges of the confidence intervals indicate that the slopes of the regressions fit to the full range (0-6) of the datasets are not significantly different, (p=0.06) suggesting no advantage of one rendering method over the other.

However we already know from the analysis done in the previous Section that a linear model is a poor fit to the full range of data from the OpenGL/match view combination. So to take this into account we will redo the analysis, restricting the input to the lower range of shape variations (0-4). This makes sense, since our primary interest is in determining whether these two rendering methods differ in their ability to convey the small differences in shape represented by this range.

Over this range the slope of the regression line on the OpenGL dataset (3.74) is itself statistically different than zero (p<0.01). Recall however, that the question the analysis of covariance is designed to answer is whether this slope is significantly different than the slope of the regression line for the global illumination dataset (8.46).

The results of this range-restricted analysis are shown in Figure 12d-f. The graphs on the left side of the figure show the mean rating curves for the OpenGL and global illumination renderings. The confidence interval plots for the slopes of the regression lines are shown on the right. Now the non-overlap in the confidence intervals show that the slopes of the regression lines are significantly different from one another (p=0.01).

The conclusion that can be drawn from this analysis is that under the conditions tested, for small variations in shape, global illumination renderings provide better visual information for shape differences, and allow better discrimination of shape. This is reflected in the steeper slope of the regression line fit to the data and the greater range of ratings given to the different shapes.

7 Discussion/Conclusions

The original question we sought to answer through this project is whether advanced rendering techniques such as global illumination allow more accurate discrimination of shape differences than standard rendering methods such as OpenGL. Further we were interested in how an additional factor -- viewpoint -- affects our ability to discriminate shape differences. Following the experiments and analyses described in the previous section we can reach the following conclusions. Under the conditions studied:

- Viewpoint appears to have an effect on our ability to discriminate shape. Under most of the conditions tested, discrimination, especially for the smallest variations, was worse when the viewpoints of the physical and rendered models were different. Performance was generally better under matched view conditions, and this is reflected in the increased dynamic range of ratings given to the different shape variations.
- Finally, to address the main question of this project, rendering method does have an effect on the ability to discriminate shape differences. Even under the best case match view conditions, there was still a significant improvement in subjects' ability to discriminate shape differences when global illumination rendering methods were used. If we take ratio of the slopes of the regression lines as a measure of the magnitude of the effect, then using global illumination methods more than doubled sensitivity (x2.26).

8 Future Work

There is still much more work to be done. First, it is clear that surface reflections are an important source of information for shape and that these reflections are affected by material properties so exploring the relations between surface finish and shape seems like an important next step. Testing different kinds of shape variation would also be worthwhile, and in fact we already have a second set of variations on the frog model that we were unable to test due to time limitations. Finally, one of the major differences between looking at the physical and rendered models is that we view the physical model with true binocular vision while when we look at the renderings we have a binocular view on a single point perspective image. In these studies we did not regard this as a drawback since our premise was that we wanted the subjects to have the best possible understanding of the shape of the physical model when judging differences with respect to the rendered models. However this does leave open the question of whether shape discrimination performance would be even better if we used stereoscopic renderings. This also opens the question of the relative magnitudes of the rendering and stereo factors in shape discrimination. Any or all of these topics would be worthwhile subjects for future study.

9 Acknowledgments

Thanks to the students and staff of the Cornell Program of Computer Graphics for participating in the experiments described in this document. Thanks also to GM Design staff who provided and painted the frog models and paint samples, and contributed to formative discussions of this project. This work was supported by

General Motors R&D and the Cornell Program of Computer Graphics.

10 References

BLAKE, A. AND BÜLTHOFF, H.H. (1991). Shape from specularities: computation and psychophysics. *Philosophical Transactions of the Royal Society (London) Series B 331,* 237-252.

CHRISTOU, C., KOENDERINK, J.J, AND VANDOORN A.J. (1996). Surface gradients, contours, and the perception of surface attitude in images of complex scenes. *Perception, 25,* 701-713.

ELLIS, S. R. (ED.) (1991). Pictorial communication in virtual and real environments. *Taylor and Francis, London.*

ERENS, R.G.F., KAPPERS, A.M.L., AND KOENDERINK, J.J. (1993). Perception of local shape from shading. *Perception and Pyshcophysics, 54(2),* 145-156.

FLEMING, R. W., DROR, R. O. AND ADELSON, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision, 3(5),* 347-368.

HAGEN, M.A. (1980). The perception of pictures. *Academic Press, New York.*

INTERRANTE, V. (1998). Perceiving and representing shape and depth. *In. V. Interrante (Ed.) Applications of visual perception in computer graphics, Course 32, Proceedings of SIGGRAPH 1998*, 1-27.

KUBOVY, M. (1986). The psychology of perspective and Renaissance art. *Cambridge University Press, New York.*

MADISON, C., THOMPSON, W., KERSTEN, D., SHIRLEY, P. AND SMITS, B. (2001). Use of interreflection and shadow for surface contact. *Perception and Psychophysics, 63*, 187-194.

MCNAMARA, A., CHALMERS, A., TROSCIANKO, T., AND GILCHRIST, I., (2000). Comparing real and synthetic scenes using human judgements of lightness. *Proceedings of the 11th Eurographics Workshop on Rendering*, 207-218.

MEYER, G.W., RUSHMEIER, H.E., COHEN, M.F., GREENBERG, D.P. AND TORRANCE,K.E., (1986). An experimental evaluation of computer graphics imagery. *ACM Transactions on Graphics 5(1),* 30-50.

MILLER, J. (1998). On reflection. *Yale University Press, New Haven, Conn.*

NORMAN, J.F., TODD, J.T., AND PHILLIPS, F. (1995). The perception of surface orientation from multiple sources of information. *Perception and Psychophysics, 57(5),* 629-636.

NORMAN, J. F., ORBAN, G., AND TODD, J. T. (2004, In press). Perception of 3-D shape from specular highlights and deformations of shading. *Psychological Science.*

PALMER, S.E. (1999). Vision science: photons to phenomenology. *MIT Press, Cambridge, Mass.*

PELLACINI, F., FERWERDA, J.A., AND GREENBERG.D.P. (2000) Toward a psychophysically-based light reflection model for image synthesis. *Proceedings SIGGRAPH '00*, 55-64.

RAMACHANDRAN, V.S. (1988). Perception of shape from shading. *Nature, 331,* 163-166.

RODGER, J.C., AND BROWSE, R.A. (2000). Choosing rendering parameters for the effective communication of 3D shape. *IEEE Computer Graphics and Applications, 20(2),* 20-28.

RUSHMEIER, H.E., WARD, G., PIATKO, C., SANDERS, P., AND RUST, B., (1995). Comparing real and synthetic images: some ideas about metrics. *Proceedings of the 6th Eurographics Workshop on Rendering*, 82-91.

TODD, J.T. AND MINGOLLA E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology Human Perception and Performance, 9(4),* 583-595.

TODD, J.T. AND REICHEL, F.D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review, 96(4),* 643-657.

VEACH, E., AND GUIBAS, L.J. (1997). Metropolis light transport. *Proceedings of SIGGRAPH 1997,* 65-76.

WARD-LARSON, G.J. (1992). Measuring and modeling anisotropic reflection. *Proceedings of SIGGRAPH 1992,* 265-272.