Using surface markings to enhance accuracy and stability of object perception in graphic displays

Roger A. Browse[a,b], James C. Rodger[a], and Robert A. Adderley[a]

[a]Department of Computing and Information Science,
[b]Department of Psychology, Queen's University, Kingston, Ontario, Canada

ABSTRACT

Effective computer graphic applications should accurately convey three dimensional shape. Previously, we investigated the contributions of shading and contour, specular highlights, and light source direction to three dimensional shape perception. Our experiments use displays of convex solid objects based on the superquadric parameterization, permitting continuous variation in their cross-sectional shapes. Our present work concerns the impact of surface markings. Rotating wireframe or uniformly shaded objects may produce perceptually distorting shapes. We investigate the idea that such distortions interfere with shape judgements, and that surface markings may either enhance perceptual accuracy by encouraging stability, or impair it by interfering with global shading patterns. Our displays include rotating objects with no surface markings, stripes, latitudinal or longitudinal stripes, each with two different scene illuminations. Observers view pairs of objects, a target shape and a second object whose shape they adjust, using mouse clicks, to match that of the target. Our principal result is that these surface patterns do not enhance performance, even though the chosen stripe intensities minimise interference with global shading, and the stripe patterns may actually encode surface curvature. We are now investigating alternatives for applying surface patterns to modelled objects, including hardware supported texture mapping. Our long term goal remains the identification of a comprehensive set of conditions for optimising shape understanding of graphic objects.

Keywords: shape perception, surface patterns, 3D graphics

1. INTRODUCTION

The effectiveness of computer graphics applications such as scientific visualization, telerobotic interfaces, and computer aided design depends on accurately portraying the three dimensional shape of objects. Generally, shape understanding can be enhanced through the use of graphic display techniques that render the surfaces more accurately and realistically. However, our previous research showed that certain of the graphic rendering variables are more important than others when observers attempt to judge the shape of displayed curved-surface objects. We investigated the relative contributions of shading and occluding contour, the lack of contribution of specular highlights, and the effect of light source direction on the accuracy of perceiving three-dimensional shape.[1,2]

Our techniques for examining shape perception use the display of three-dimensional solid shapes based on superquadric parameterization[3]. Variations in the superquadric parameters produce objects varying continuously in the shape of their orthogonal cross-sections. We use the range from a sphere to a cube to present a wide variety of convex objects. Each has a simple numerical representation of its shape that we can use to measure performance in shape discrimination and matching tasks. The displayed objects rotate about two axes to encourage observers to apprehend the complete shape rather than forming judgements on the basis of occluding contours. In our present research, we use the same approach to address the issue of the effects of surface markings on the perception of object shape.

2. SURFACE MARKINGS IN THE PERCEPTION OF SHAPE

When the display of a rotating solid has the shading due to lighting removed so that only the occluding contour is presented, observers often report the perception of a stationary but distorting object. In our previous work with shaded, rotating objects of uniform reflectance, observers still sometimes reported noticing distortion of the rotating object. We believe that these conscious experiences of distorting shape contribute to errors in shape judgement, and further, that subliminal distortions may interfere with the judgement of shape. It is a reasonable conjecture that the presence of surface markings on a rotating object will encourage the perception of rigidity and thereby enhance shape perception accuracy. Another way to view this issue is to consider the idea that accurate perception of the shape of a rotating object may rely on an understanding of the movement of the individual locations on the surface of the object. If true, then it seems reasonable to expect that the application of surface markings would enhance the observer's stable identification of surface locations, and hence enhance shape perception.

On the other hand, surface markings may reduce accuracy of shape perception. In our previous studies we have shown that the presence of shading across the surface of an object is a significant source of information in shape understanding[1]. The presence of surface markings may impede the observer's integration of these shading patterns across the surface of the object resulting in less accurate shape perception. Certainly there is a convincing argument that in the extreme, as in the case of camouflage, surface marking may hinder and distort perception.

There is a wide range of possible surface markings that could be tested for their influence on the perception of object shape. As a first step in understanding the role of surface markings in shape perception, we have used a simple form of striping applied to the superquadric shapes. The objects used in the unstriped condition have a brightness that is the average of the two brightnesses that alternate in the striped objects, thereby controlling for overall brightness. The brightness levels used in the striping were selected to have a low contrast in order to minimally interfere with the shading across the surface. We used two different directions of illumination.

3. EXPERIMENTAL PROCEDURE

We generated the displays for the experiment using a Silicon Graphics O2 system with a 19 inch monitor. The interface for the experiment, illustrated in Figure 1, was a straightforward extension of that employed in our previous research[m]. In that research, observers adjusted a two-dimensional contour to attempt to match the cross-sectional shape of a rotating superquadric object displayed alongside it. In the current experiment, the observers viewed a pair of superquadric objects in the Animation Window, and adjusted the right hand (adjustment) object to match the shape of the left hand (target) object. They made the adjustments by clicking with the mouse on either the circle or square in the Adjustment Window, located immediately below the adjustment object.

We varied conditions relating to both the objects and the rendering parameters. Since our primary interest concerned possible contributions of surface patterning to shape perception, the pairs of objects in each display could be in one of three surface patterning conditions. These were either latitudinal stripes, longitudinal stripes, or uniform surface (no striping). We produced the stripes by varying the material properties of strips of the quadrilateral primitives used to tessellate the surfaces. While the stripes for the two objects in a pair were in the same direction, both the width of the stripes and the alternation of light and dark differed, preventing observers from simply matching stripes to achieve a global shape match. Figure 2 depicts the three display conditions applied to different superquadric shapes.

The superquadric parameterization of shape allows for the definition of cross-sections that range from square (at $\varepsilon = 0.0$) to circular (at $\varepsilon = 1.0$). Each object is defined by two cross section parameters (ε_1 and ε_2), and the object's extension along each of the local Cartesian axes. The shapes that we used all had an aspect ratio of 1:1 in the horizontal plane, and had an aspect ratio of 1:1.4 for each of those axes relative to the vertical axis. The cross-sectional shape of the target object ranged from 0.2 to 0.8 in eight equal increments of the superquadric shape parameters, with $\varepsilon_1 = \varepsilon_2$. The shape of the adjustment object always differed initially, offset by either 0.135 or 0.18 from that of the target object, and equally often towards either the square or round end of the shape parameter scale.

When we found, in preliminary data, that observers were exhibiting a bias in their shape matches, tending to make the squarer objects rounder than the exact match, we decided to add another condition to the experiment. We conducted the initial testing using a light source direction 45 degrees to the left and elevated 70 degrees. For the full experiment, we added a second light direction condition, with the light source direction coincident with the viewing direction (that is, from directly in front of the objects, with no elevation).

These variations produced 108 distinct combinations of conditions, in addition to the two light source directions. Each observer participated in one light direction condition, and viewed two replications of each of the other combinations of variables, for a total of 216 trials. We recruited our participants by E-mail postings to the graduate and undergraduate students in the Department of Computing and Information Science at Queen's University. There were 12 observers for each of the light directions, for a total of 24 participants in all. We used a different random ordering of trials for each observer.

During a trial, both of the objects rotated continuously with a compound motion. This combined complete rotation about the vertical axis with swaying back and forth through an arc of 60 degrees about a horizontal axis. The starting points for the swaying motion differed between the two objects, so that, at any given time, the observer did not see matching views of them. The observer clicked repeatedly on the shapes in the Adjustment Window, until a satisfactory match was achieved, and then initiated the next trial by clicking on the Test Window. There was an upper limit of 30 seconds per trial, so that if the observer had not clicked on the Test Window, the system automatically initiated the next trial.

For each trial the testing software recorded the final shape of the Adjustment object, the elapsed time and the number of clicks the observer used to arrive at the match. We analysed these data using the SYSTAT for Windows statistical package.

4. EXPERIMENTAL RESULTS AND DISCUSSION

Figure 3 shows the mean error across all 24 observers for each of the three striping conditions. An analysis of variance shows that this a reliable difference in error for the different surface patterns. This favors the interpretation that the striping has interfered with, or camouflaged the shading across the surface and reduced the informative aspect of that cue to a greater degree than the presence of striping has increased the shading through stabilization of the objects. Though not statistically significant, the mean error for the longitudinal striping is greater than that of the latitudinal striping. If the striping were stabilizing the object under rotation, it is reasonable to expect the opposite effect. The rotation of the object is about the vertical axis, and so a perceived distortion of the horizontal cross-section would be difficult when the object is striped longitudinally, because to perceive such a distortion would require an alteration in the perceived width of the striping. The latitudinal striping would not introduce that requirement.

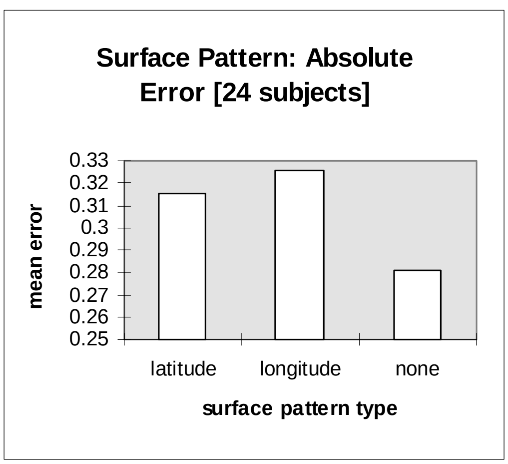

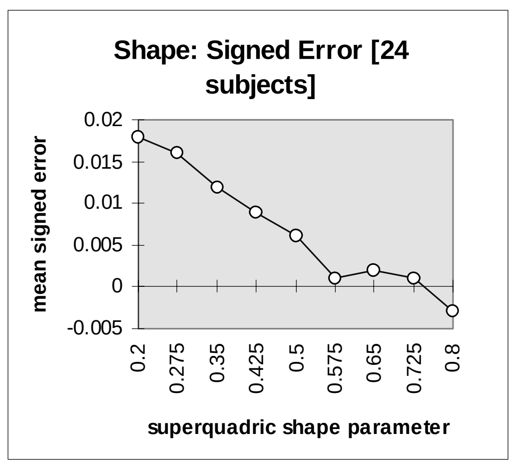

Figure 3. Mean error for each surface pattern.

Figure 4. Signed error for each shape value.

The main conclusion of our experiment is that forms of striping which one would expect to enhance performance, in fact degrade the accuracy of shape perception. This is true in some instances for striping carefully chosen with low contrast to minimally interfere with the visibility of the surface shading due to lighting. We found in some cases, that the accuracy of perception was degraded with striping which, in the spacing of the stripes, actually encoded the surface curvature. Thus there may be a surprisingly small range of surface striping that will enhance the accuracy of shape perception for arbitrarily rotating and illuminated objects. For rendering and rotation configurations that induce object distortions, it may be possible to devise striping perpendicular to the direction of the distortion, thus offsetting the distortion with bands that cannot be seen to change in width.

A second main result is the perplexing fact that observers tend to systematically overestimate the cross-sectional superquadric parameter. Figure 4 shows the mean signed error across all subjects and conditions plotted against the target superquadric shape parameter. We expected there to be comparable error levels for underestimation and overestimation of the parameter because the shape that was being adjusted began equally often with offsets on either side. A significant component of the error resulted from situations in which the adjustment object was initially more square than the target, with the observer adjusting past the point where the two shapes were identical, and beyond to the point of responding with an overestimation of the roundness of the cross section of the object. The graph in Figure 4 also indicates the statistically significant result that this bias towards overestimation increases as targets become more square in cross section

We first noticed the bias towards overestimation when observers were carry out the task with a scene lighting direction that was always above and off to one side. In order to find out if this lighting was contributing to the bias, we had a second set of 12 observers carry out the task with scene lighting that originated down the viewing axis. The bias was also present for this second set of subjects, ruling out lighting as the cause. In fact, light source direction had no significant effect on the estimation of the superquadric parameter, and so the results shown in Figures 3 and 4 represent both of the lighting conditions combined.

In the course of this ongoing research program we have developed a variety of techniques to measure human performance in the judgement of shape. These techniques include forced choice similarity, same different judgements, cross sectional matching, and now shape matching. We have used these techniques to examine a wide variety of parameters that control the presentation and rendering of graphic objects, of which local surface reflectance properties is our latest concern. Our long term goal is the identification of a complete set of presentation and rendering conditions that optimize the accuracy of human judgement of the shape of graphic objects.

ACKNOWLEDGEMENTS

The research reported in this paper was conducted with financial support from the Institute for Robotics and Intelligent Systems and the Natural Science and Engineering Research Council of Canada.

REFERENCES

1. R.A. Browse and J.C. Rodger, "Investigations of three-dimensional shape perception for telepresence using superquadric primitives", *Proceedings of the SPIE, Volume 2179, Human Vision and Electronic Imaging, I* pp.235-246, 1994.
2. R.A. Browse, J.C. Rodger, and R.A. Adderley, "The effect of lighting direction on the perception of shape in graphic displays", *Proceedings of the SPIE, Volume 2179, Human Vision and Electronic Imaging II*, pp.353-359, 1997.
3. A.H. Barr, "Superquadrics and angle-preserving transformations", IEEE Computer Graphics & Applications, 1(1), pp.11-23, 1981.