Biol. Cybern. 53, 137–151 (1986)

Biological Cybernetics

Perception of Solid Shape from Shading

E. Mingolla[1,*] and J. T. Todd[2]

[1] Center for Adaptive Systems, Boston University, 111 Cummington Street, Boston, MA 02215, USA
[2] Brandeis University, Waltham, MA 02154, USA

Abstract. Observers judged the slants and tilts of numerous regions within shaded images of ellipsoid surfaces that varied in shape, orientation, surface reflectance, and direction of illumination. The perceived three-dimensional structure of each surface was calculated from these judgments. Much of the error in observers' responses resulted from a tendency to perceive surfaces whose axes were aligned with the display screen. The presence of specular highlights or cast shadows, in contrast, had no effect on performance. The results of the experiment indicate that several assumptions of certain formal models for perception of shape from shading are not psychologically valid. The most notable of these assumptions are that the visual system initially assumes that all surfaces have Lambertian reflectance and that illuminant direction must be known before shape detection can proceed. These assumptions are often accompanied by a third assumption that surface orientation is detected locally, and global shape determined by smoothing over local surface orientation estimates. The present experiment indicates that an alternative approach offered by Koenderink and van Doorn may be more psychologically accurate, as it avoids all three assumptions.

Introduction

While much of our environment consists of solid objects with extended and often smoothly curved surfaces, surprisingly little research has been done on the perception of solid shape. Two formidable obstacles have slowed the progress of such research. First, an experimenter must generate displays that depict different shapes using precisely controlled optical variables. Second, subjects' impressions of shape must be accurately recorded. The computer graphics technology needed to execute such procedures has only recently become available and affordable for a typical perceptual laboratory.

At present psychologists do not even have a consensual attitude toward what solid shape with respect to vision *is*. While classical differential geometry offers elegant and powerful descriptions of curved solid shapes themselves, psychologists seem until recently to have considered such descriptions to be cognitive or indeed intellectual achievements, having little to do with the sensory and perceptual processes of vision. Differential geometry and perceptual theories employing it are considered at some length later in this paper, but it is worthwhile to consider briefly the approach to solid shape perception which is tacit in many other visual theories.

Certain aspects of solid shape perception are often analytically dissected and studied piecemeal. A search for studies of perception of shape in three dimensions tends to find near misses under the entries "form perception" and "depth perception". In fact the perception of form, or shape in *two* dimensions in a picture plane, has a rich tradition (Rock 1973; Zusne 1970)[1]. On the other hand, the study of depth perception has largely concentrated on how the third dimension in three dimensional structure is seen at all, with little reference to differences in shape in three dimensions. A typical study in depth perception might investigate the relation of textural compression or motion parallax to perceived slant in depth, but the object slanting in

* Supported in part by the Air Force Office of Scientific Research (AFOSR 82-0148). The research reported in this paper was performed while the first author was a graduate student at the University of Connecticut at Storrs

1 Researchers often distinguish two dimensional "form" from "shape", taking shape to describe the arrangement of all points on an object, while form refers to an intrinsic structure or essence common to perhaps many related shapes

depth remains a planar segment throughout. Tacit in this research tradition is the implication that apprehension of solid shape is not rooted in optical information at all, but occurs only after some object recognition process matches stored three dimensional representations to the two dimensional form and depth information available perceptually.

Common to most studies in both the depth and form percpetion has been a reliance on experimental apparatus with which continuous variation in displayed solid shape, such as from a sphere to an ellipsoid, could not easily be achieved. While some form perception studies have employed curved line contours in the picture plane, most depth perception studies have employed depictions of planar or polyhedral objects. Only recently have the manipulation of smoothly curved objects and the analysis of visual information for curvature attracted attention (Braunstein et al. 1982; Cutting and Millard 1984; Lappin et al. 1980; Todd and Mingolla 1983). Meanwhile increasingly sophisticated formal procedures for determining solid shape from optical information have been proposed (Horn 1975, 1977; Koenderink and van Doorn 1979, 1980, 1982a, b; Pentland 1982b), but techniques for evaluating the perceptual validity of such procedures have yet to be developed. Accordingly, the present paper presents an experiment which introduces new methods for investigating the perception of curved solid shapes.

The displays employed in the present experiment depict curved solid objects solely through variations in shading. We perceive the shapes of objects in part through variations in shading resulting from various positions of illuminants, object surfaces and our eyes. These variations of shading often result from smooth variations in the local orientation of surfaces of curved objects, making shading a natural source of information to manipulate for the study of solid shape perception.

The problems posed in the study of perception of solid shape require certain new concepts and techniques for conducting perceptual experiments. Accordingly, the first section of this paper considers issues involved in the very description of solid-shapes and the connection of shape descriptions to theoretical analyses and perceptual experiments.

The next section outlines how methods for describing solid shape are related to information available in shaded images. Most types of visual information, such as motion parallax, texture compression, or binocular disparity are typically taken to be functions of eye position and visible surface layout alone, but intensity values in a shaded image are clearly affected by the illuminant direction and surface reflectance characteristics, as well as surface and eye positions. Any analysis of perceptual information available through shading must therefore consider how the effects of surface reflectance and illuminant direction can be untangled from those of solid object shape and position.

Since not all information available in visual displays is actually employed by the visual system, perceptual experiments that measure differences in performance as information is manipulated must be performed. In the case of perception of shape through shading, however, paradigms for evaluating perceptual performance are needed. What does it mean to say that the shape of, say, a coffee cup is more accurately seen in one condition or another? If a subject does not report experiencing quite the shape depicted by the experimenter, just what *does* the subject experience? In the experiment reported below the depiction of solid shapes through shading is manipulated in such a way as to begin to answer these questions. The experimental methods employed enable experimenters to compare computed solid shapes of displayed objects with reconstructions of subjects' perceived shapes, as well as to assess the impact of display variables, such as highlights or cast shadows, on perceptual performance.

What Is Solid Shape?

Describing Solid Shape

Studying the perception of three dimensional shape poses problems in both the domains of analysis of visual information and the measurement of perceptual performance. Almost all of these difficulties can be traced to the nature of descriptions of objects' solid shapes, which seem to be either short but imprecise or precise but cumbersome. A few important shapes have names, such as "sphere", "cylinder", or "cone". Even describing an object as being shaped "like a fish" or "like a mushroom" can be informative and efficient in some contexts, but clearly poetic gains are accompanied by losses in precision in such descriptions. If on the other hand one possessed a complete list of all the three dimensional Euclidean coordinates of the points on the surface of an object up to some arbitrarily fine resolution, such a description could be said to embody all the information for the object's solid shape. The information for shape in such a description would be buried, however, and it is difficult to even imagine how a perceiver could extract functionally meaningful aspects of shape from the coordinates. Even from a purely mathematical standpoint, moreover, a mere list of points does not seem to be a description *of* solid shape, but more adequate modes of mathematical description exist.

Classical differential geometry offers general and precise descriptions which are in principle applicable to the surfaces of all physically realizeable solid shapes. A number of recent theoretical analyses of the perception of solid shape have in fact employed concepts of differential geometry, and a close examination of the assumptions involved in relating these concepts to human perceptual performance is warranted (Koenderink and van Doorn 1980, 1982a; Pentland 1982b).

The Approach of Classical Differential Geometry

Differential geometry provides a ready made calculus of surfaces with promising applications to visual perception[2]. Differential geometry is a formal mathematic discipline, however, and some of its procedures might not be realizable in living visual systems. This section highlights a few of the most basic assumptions and methods which come into play when differential geometry is employed in perceptual models.

The simplest classical methods apply to surfaces which are sufficiently smooth to admit approximation in the limit by small planar segments. At any point on a surface, then, the surface can be said to have a local orientation, described by the components of a unit normal to the tangent planar segment at that point, as the planar segment is shrunk infinitesimally. Starting from a given point, moving continuously on a curved surface to nearby points produces smooth changes in the orientations of local surface normals. The rate of change or surface orientation in a given direction is the curvature of the surface at that point in that direction. In general shape information at a point is captured by knowing the direction and magnitudes of the two orthogonal "principle curvatures", the extreme values of curvature at that point.

Points on a surface can be classified according to the signs of their principle curvatures. Elliptic points have principle curvatures with equal signs and occur in "bumps" or "dimples". Hyperbolic or "saddle" points have principle curvatures of opposite sign, while parabolic points have at least one principle curvature of zero magnitude.

All of the terms used so far refer to the *local* differential geometry of surfaces, specifying quantities defined in a small neighborhood of a point. Results from *global* differential geometry are of greater relevance to describing solid shape, however. Global theorems constrain what combinations of local surface orientations are physically realizeable in Euclidean 3-space. For example, no actual surface enclosing a finite volume can be composed entirely of hyperbolic points. Global theorems also specify what conditions must obtain if others are known to exist. For example, regions of positive and negative curvature on a surface are always separated by parabolic lines, that is, curves composed entirely of parabolic points.

2 For readers without a background in differential geometry, an excellent nontechnical introduction appears in Hilbert and Cohn-Vossen's classic Geometry and the Imagination

Differential Geometry for Vision

The mathematics of differential geometry proceeds in whatever coordinate system the mathematician finds convenient, and all points on an object are equally accessible for analysis, so long as their positions can be rigorously specified. Visual perception occurs with reference to an eye's viewpoint, however, and the relation of a point of observation to a visible surface is the bedrock of analyses of visual information for solid shape perception. The perceiver thus imposes a perspective or coordinate system on the structure of light which is reflected by an object's surface. Optical information is considered next in this paper, but it is worthwhile here to briefly examine the role of coordinate systems in the geometric concepts just discussed.

The values of the components of surface normals, which define surface orientation at a point, are dependent on the orientation of the chosen coordinate system, but the values of curvature at that point remain the same regardless of coordinate system orientation. A perceiver's task in shape perception is often said to involve a transition from eye-centered to object-centered coordinates. More is required, however. Local surface orientation defined with respect to a line of sight changes as an observer moves about or as the perceived object rotates, but the curvature at a point on a rigid object remains the same, even over such motions. Thus, apprehension of curvature involves not just a coordinate transformation, but a transformation from a coordinate-bound to a more coordinate-free description.

Analyzing Information in Shaded Displays

Sources of Variation in Shading

The experiment described in this paper employed computer generated shaded images of ellipsoids. The process of computing the appropriate intensity at each image point to simulate a specified arrangement of surfaces under point-light illumination is by now a familiar one in the computer graphics literature, and an introduction to such image generation for perceptual experiments appears in Todd and Mingolla (1983). The core issue for perceptual psychology is that the mapping from viewpoint, illuminant position, surface layout and surface reflectance characteristics to resulting image intensities is locally many to one, so that no obvious route from image intensities to perceived object shape can be stated.

Great difficulties can be encountered in even the most simple image environments. Consider, for example, a small window in the interior of a shaded image of a sphere. Let that sphere have Lambertian reflectance and be illuminated by a point source. Next perturb the image in one of three ways. 1) change the position of the point source, or 2) let the reflectance of the sphere be non-Lambertian, or 3) distort the shape of the sphere into that of an ellipsoid. The intensity gradient within the small window will change in all three cases, and infinite families of parameters exist for which the changes in intensity are locally indistinguishable.

Formal Analysis of Information for Shape in Shading

At present there are no models rooted in perceptual data of how humans perceive shape from shading. What analyses exist are almost exclusively formal accounts of how one *could* determine shape, given a pattern of intensities and perhaps some additional knowledge or assumptions. The proposed procedures can be grouped into two categories.

The first approach is essentially a frontal attack on the local many-to-one mapping of display variables to intensities. The perceiver's task is taken to be the computation of the inverse of the image generating process. Since a strict inverse does not exist, auxiliary constraints or explicit knowledge of certain display variables is added until a unique solution is found. This is the approach of Horn's pioneering work in the last decade (Horn 1975, 1977). Todd and Mingolla (1983) reviews the assumptions of Horn's procedure and those of a related approach of Pentland (1981, 1982a). Horn makes no claims about psychological validity and invokes vast amounts of externally supplied knowledge concerning surface reflectance functions and illumination conditions in order to compute shape.

Pentland (1982b) has modified some details of his earlier work, but the nature of his analysis remains the same. Pentland instead assumes Lambertian reflectance and does not require any knowledge of illuminant direction. Pentland does claim certain psychological validity for his model and, in that respect, requiring little prior knowledge is desirable, since humans routinely perceive shape without being supplied such knowledge. Notably, however, Pentland's procedure does require the prior computation of a perceived direction of illumination by one module in order to provide data to the shape module. Thus, while the information for illuminant direction is taken from the image itself, the shape from shading algorithm cannot procede without an estimate of illuminant direction.

Pentland's approach in fact employs concepts from local differential geometry. The strategy is to convert local variations of image intensity into local estimates of surface orientation, by considering what surface normals are consistent with the image intensities, given the estimated illumination direction and assuming Lambertian surface reflectance. Pentland argues that any local gradient of image intensities could have been produce by some Lambertian sphere at some orientation. Additional smoothing assumptions are then used to calculate what surface is most consistent with the surface orientation estimates derived from the initial sphere assumption. For a given sampling density, perhaps one per image point, the resulting grid of surface orientations can be integrated into a surface.

The second category of formal analyses has been contributed in a strong series of papers by Koenderink and van Doorn (1970, 1980, 1982a). Instead of trying to locally invert the image formation process, they show that the solid structure of an object modulates the structure of light available at station points in *globally* constrained ways. By structure of light is meant the positions of singularities of luminance (local maxima and minima) and the topology of connection and closure of isophotes, contours of equal luminance. The implication of their analysis is that perceivers need not bother determining what variables generated a given intensity at all. Instead, the nestings and inflections of possible luminance distributions are shown to be both highly constrained and highly specific to those configurations of hills, valleys, and saddles which make up solid objects. While their analysis does not purport to produce parametrically accurate surface orientation or surface position estimates across a whole object, there is no perceptual evidence that humans are capable of such performance. Moreover, from a theoretical perspective, the absence of such purported accuracy is more than compensated by the robustness of the mapping from luminance topologies to object structure. That is, unlike in the case of local intensity to shape mappings, many types of perturbations leave the structure of isophotes unaffected, while actual changes in isophote topologies can only occur in ways that are lawfully related to changes in shape or object position.

Although not reviewed here, another visual theory is also noteworthy for its emphasis on the importance of global patterns of optical inputs in its analyses of perceived depth, form, and lightness (Cohen and Grossberg 1984; Grossberg 1983).

Sadly, almost no perceptual data on the validity of the models just discussed exists. No experimental techniques to date have approached the complexity needed to challenge the models of Koenderink and van

Doorn, but the displays in the experiment described in this paper are rich by traditional standards and they begin to probe many of the issues already raised.

Measuring Performance of Solid Shape Judgement

The quantitative assessment of performance in solid shape judgments poses conceptual as well as practical problems for an experimenter. Global verbal descriptions such as "pear-shaped" are simply too vague for most experimental paradigms. In order to investigate the perception of two differently shaped pears, the experimenter would like to gauge the subject's impression of the relative differences in the pears' shapes over their entire visible surfaces. If the base of one is "rounder" than that of the other, or if one pear appears to have greater longitudinal symmetry, how can the *extent* of these perceived differences be assessed? The impression of shape "at a point" of the surface must be probed, but global solid shape as such does not exist at any point. The shape related quantities that are defined at a point are local surface orientation and curvature. Thus, as the discussion of differential geometry indicated, it may be necessary to measure the subject's impression of local surface orientation as a means of parametrically probing their impression of solid shape.

As a practical matter, assessing the subject's impression of surface orientation can probably only be carried out for a few dozens or, conceivably, hundreds of points. For an image of the resolution available on a typical computer graphics display (e.g. 640 X 512 pixels) it would be extremely difficult, given limitations in subjects' accuracy, to mathematically "reconstruct" an interpolation of the perceived shape by integrating over pointwise orientation judgments. While this would be a fascinating project, even if carried out successfully, the result would be a set of coordinate values, with precisely the limits mentioned previously. It would still remain to give some compact description of the nature of the transformation from depicted to perceived object.

The preceding discussion indicates that an experiment in which subjects are asked to indicate surface orientation is in danger of remaining an experiment about surface orientation rather than about shape. Clearly a subject's misperception of depicted shape would result in errors in orientation judgments, but the magnitudes and directions of orientation errors must be related to recognizable aspects of the depicted object's shape. In other words, sizable orientation errors could result from the subject experiencing the object's shape entirely accurately, but mistaking the object's overall orientation to the line of sight. The account of the experiment reported below describes how this ambiguity can be resolved.

Experiments

Introduction

The basic strategy of the experiment was to present the subjects with displays of two related but different solid shapes in a variety of conditions which affect image shading. The chosen shapes were ellipsoids, in part because they permit parametric control over displayed shape; two ellipsoids can, for example, differ only in the length of one semiaxis, making one ellipsoid thicker than the other. Subjects were asked to judge local surface orientation at a number of points; this permitted a fitting procedure to determine what shape in what orientation the subject actually experienced. Subjects were also asked to judge illuminant direction, which was one of the independent shading variables in the experiment. The other shading variables were the presence or absence of specular highlights on the ellipsoid surface and the presence or absence of cast shadows on a background.

The experimental manipulations were designed to investigate the following issues:

1. Is the formal assumption of Lambertian surfaces psychologically plausible for shape detection? If human visual processes somehow embody such an assumption, performance should be worse for shiny ellipsoids than for dull ones, because the external knowledge usually said to compensate for shininess is not provided to the subject.

2. Is performance on direction of illumination estimation related to performance on surface orientation estimation? In particular, is there evidence that an estimate of the illuminant direction is necessary for analysis of shape?

3. Is information for solid shape fundamentally local or global? This issue is complex, but the intensity variations produced by the two solid shapes employed in the experiment are indistinguishable locally, particularly in the context of the manipulation of surface reflectance and illuminant direction. A purely local procedure would therefore perform with the same accuracy for the two shapes employed in the experiment.

Method

Specifying Orientations. The positions of the objects and light sources in the experimental displays and the data collected from subjects are described in terms of *slant* and *tilt*. Figure 1 describes how these terms can be used to describe both local surface orientation and global object orientation. Similarly, although not displayed in Fig. 1, the same reference system can be used to describe a prevailing direction of illumination,

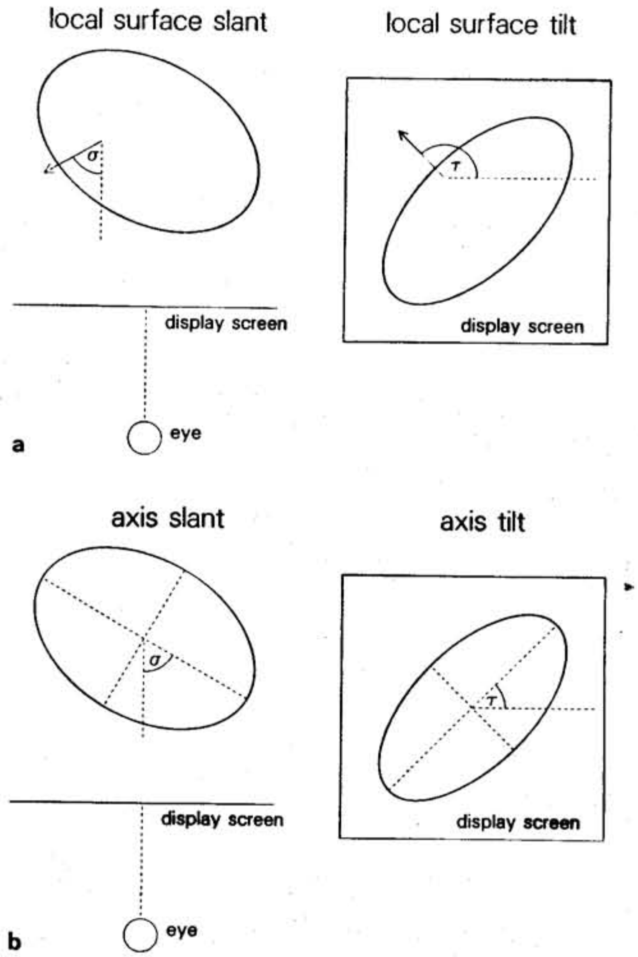

Fig. 1a and b. Slant and tilt are defined with respect to a coordinate system whose axes are parallel to the display screen and the line of sight. **a** shows how these terms apply to the local surface orientation of points on a solid object. (Only the outer rims of the solid object are depicted in this line drawing.) Local surface slant, σ, is the angle formed by the local surface normal (solid arrow) and the direction parallel to the line of sight to the display screen center. By convention, the direction along the line of sight is said to have zero slant. Tilt, τ, is the radial direction of slant relative to the picture plane. By convention, tilt is measured counterclockwise, with zero in the "three o'clock" position. **b** indicates how slant and tilt can also be used to indicate the global orientations of objects having axes of symmetry

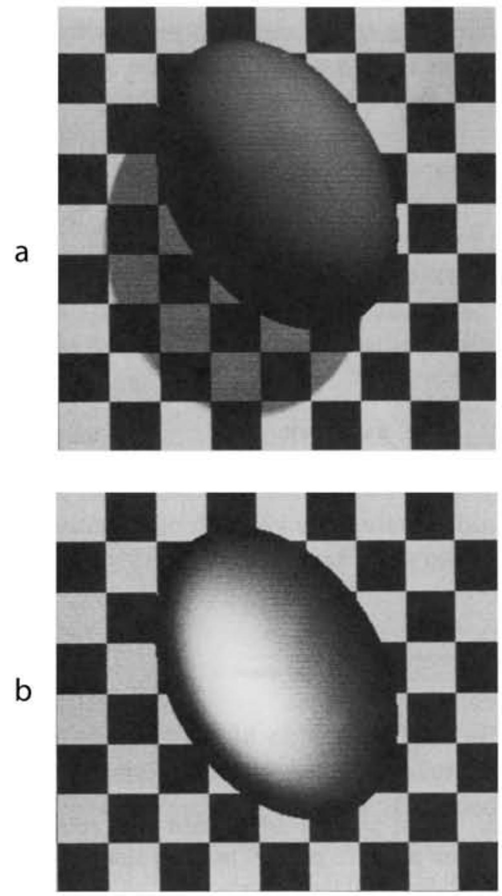

Fig. 2a and b. Two of the 16 ellipsoids in the experiment illustrate the variations of the four experimental factors. The ellipsoid in **a** is dull, has a cast shadow, is obliquely illuminated and has the low eccentricity shape. The one in **b** is shiny, does not have a cast shadow, is directly illuminated and has the high eccentricity shape. (It should be noted that because of distortions in the processes of photographic reproduction, these pictures appear considerably less realistic and three dimensional than they do on a video screen)

by imagining a vector positioned at the center of a display screen and pointing at a light source.

Subjects. Five male graduate students participated in four 90- to 120-min sessions and were paid $40.00. Four were naive concerning the design of the experiment, while one had casual conversations with the experimenter prior to participation.

Apparatus. Displays were presented on a Conrac 17-inch video monitor controlled by a NOVA minicomputer. The displays were viewed binocularly at a distance of approximately 100 cm. Head and body movements were not restricted.

Stimulus Displays. Each display consisted of an ellipsoid object in front of a checkerboard background, as exemplified in Fig. 2. A complete crossing for four two-level experimental factors produced 16 related but different displays. The four factors were: *surface reflectance, cast shadows, direction of illumination*, and *shape*. The chosen levels of these factors can be described with reference to the image generating procedure which was employed.

Todd and Mingolla (1983) describe the assumptions and procedures related to generating shaded displays on computer controlled video screens, and Mingolla and Todd (1984) describe the geometry of the modeled ellipsoids, observer, and display screen in detail. The heart of the display algorithm for the shaded images employed in the experiment is the image shading equation (Blinn 1977):

$$I_P = I_A s + I_L s(\mathbf{L} \cdot \mathbf{N}) + I_L g(\mathbf{H} \cdot \mathbf{N})n. \quad (1)$$

In (1), I_P refers to the intensity of a picture element or pixel on the video screen. That intensity is the result of three terms in the equation, the first describing the effects of ambient illumination on a surface, the second describing the effects of diffuse reflectance of a point source's illumination at a surface region, and the third describing specular reflectance at the same region. I_A is the intensity of ambient illumination and I_L is the point source intensity. The albedo of the surface is specified by s, ranging from 0 for perfectly black to 1 for perfectly white. The factor g specifies the proportion of incident light reflected specularly, and n describes the scattering of the specular reflectance; a perfect mirror would be modeled by an infinitely large n. **L** is a unit vector pointing from the modeled surface region to the point source illuminant, **N** is the outward pointing unit surface normal, and **H** is the unit vector which bisects the angle formed by **L** and the line of sight.

Two levels of *reflectance* were employed; ellipsoids could be shiny or dully. For shiny displays s, g, and n were 0.4, 0.5, and 9 respectively, and for dull displays s was 0.4 and g was 0. Consequently shiny ellipsoids had higher maximum intensities than dull ones.

For a given display a *cast shadow* was either present or absent. When shown, the cast shadows were computed according to the method outlined in Mingolla (1983). A checkerboard background pattern was modeled to be parallel to the display screen and perpendicular to the line of sight. Its squares were dull, with g equal to 0 and s values of 0.2 and 0.5 for dark and light squares, respectively. Illumination for shadows was taken to be from an infinitely distant cluster of five point sources in a cross pattern, with the cross points displaced 1° from their nearest neighbor. The cross illumination pattern produced somewhat fuzzy penumbras in the cast shadows.

Although homogeneous diffuse illumination of about one tenth the intensity of the prevailing point source illumination was included in all displays, a strong directional illumination was always present. The computed *direction of illumination* from the point source illuminating the ellipsoids was either "direct", with a slant of 10°, or "oblique", with a slant of 40°. Computed illuminant tilt for all displays was 60°.

Each display contained an ellipsoid with one of two *shapes*; the first had semiaxis lengths with ratios 100:84:50, and the second had ratios of 100:70:60. For reference in this article, objects of the first shape are referred to as "high eccentricity" and those of the second shape as "low eccentricity". These relative terms refer to the ratios of the lengths of the longest and shortest axes of the respective ellipsoids. Although the ellipsoids were of two solid shapes, the outer contours of the ellipsoids in the picture plane were always identical. This was achieved by finding an appropriate combination of semiaxis lengths and orientations as described in Mingolla (1983). In order to match the outer visible contours of ellipsoids with the dimensions just described, the semiaxis orientations of the high eccentricity ellipsoids were first rotated 45° about the horizontal, or x axis, and 120° about the line of sight, or z axis. Corresponding rotations for low eccentricity ellipsoids were 20° and 120°. The purpose of the outer contour control was to guarantee that differential shape information was contained solely in the interior intensity gradients.

Procedure. A within-subjects design was used, and each subject saw each display twice. Eight displays were presented in a randomized order in the first session and the other eight in the second session. The set of sixteen was similarly repeated in the third and fourth sessions.

At the start of the experiment each subject was given 10–20 min of instruction about slant and tilt. Each was shown a physical sphere with toothpicks stuck in it to illustrate surface normals and was quizzed about the slants and tilts of the toothpicks until the experimenter was satisfied that the concepts and frame of reference to be adopted relative to the display screen were clear to the subject. A physical model of an ellipsoid was also shown to each subject. Subjects were warned that an ellipsoid was not necessarily a surface of revolution, and that the orientation of displayed ellipsoids would not necessarily be such that their axes would be parallel to the display screen or to the line of sight. Subjects were run individually and keyed their responses into a terminal connected to the minicomputer that controlled the displays. The experiment was self-paced, and a new display generated a few seconds after the subject reported the last surface orientation for a given display. Subjects had a chart available at all times which labeled tilt directions in the picture plane in increments of 10°, from 0° counter-clockwise to 350°, with 0° in the standard "three o'clock" position. They were told that all displayed slants would be less than 90°, but were otherwise left to report slants and tilts of whatever ranges and in whatever increments they pleased. For the surface orientation judgments a

Table 1. Slants and tilts of probed points in degrees

0	30	60	90	120	150	180	210	240	270	300	330
60	30	30	10	20	60	20	10	50	50	40	40

subject's attention was directed to a three pixel high black cross, which flashed three times and then remained steady on the screen until the slant and tilt responses for that point were collected.

As each display was presented, subjects were prompted on the terminal to report the slant and tilt of the prevailing illuminant direction. They were then asked to report the slant and tilt of the surface normals for 13 points on each display. Sets of probed points were related for all displays by virtue of having the same slants and tilts, which were chosen as follows. Tilts corresponding to the directions of each of the twelve numbers on an analog watch face were chosen for each display. Slants were randomly assigned once for all displays from the set 10° through 60° in units of 10°, with the constraint that each slant appear twice, resulting in the combinations of Table 1.

In addition, the point having the slant 0° and tilt undefined was also displayed; its surface normal pointed along the line of sight. It may be noted that for high eccentricity ellipsoids the locations of the probed points on the screen varied somewhat from the locations of the probed points for low eccentricity ellipsoids, in keeping with the constraint that the depicted *orientations* of sampled points be constant for all displays. Details of the procedure for matching probed points in this way can be found in Mingolla and Todd (1984).

Before beginning the 16 experimental displays each subject practiced reporting direction of illumination and surface orientation judgments for two ellipsoids similar to those used in the experiment. In the practice trials feedback was given in two forms for surface orientation judgments. First the computer terminal displayed the deviation in degrees between the reported and computed orientations at a point. The display screen then showed which point on the ellipsoid surface actually had the computed slant and tilt reported by the subject in response to the original probed point. It was hoped that in this way subjects could calibrate their reports to their visual impressions. During the experiment itself, however, no feedback was given to subjects.

All the subjects found the task of reporting orientations in terms of slant and tilt to be somewhat difficult. Each agreed that the surface orientations themselves were easy enough to "see", but difficult to report in angular units. Subjects often took over a minute to make a single response, especially during the early trials, but all felt confident that they succeeded on the whole in representing their impressions through the slant and tilt responses.

Results

The results of the experiment can be viewed in two main ways. First, display variables are organized into a simple factorial design, and surface orientation and illuminant direction judgments can be analyzed by conventional ANOVA and correlation techniques. These techniques do not yield satisfying interpretations of perceived shape, however, and in the second part of this section the best fitting ellipsoids for subject's responses are described.

Statistical Analyses. The discrepancy in degrees between a subject's reported orientation and the displayed orientation was computed for all probed points. While subjects reported two numbers, slant and tilt, for each reported orientation, those numbers refer to one surface normal. The angle formed between that surface normal and the surface normal used by the display algorithm was the error at that point for all reported analyses. The mean of the 13 error scores for a given presentation of a given display was the score of overall performance on that display. The patterns of errors across displays were similar for all five subjects, and their data for each task were analyzed in a single repeated measured ANOVA.

Two shapes were displayed in two directions of illumination, with or without cast shadows, and with either shiny or dull surfaces. While ANOVA's for surface orientation and direction of illumination judgments yielded many statistically significant effects due to the large number of observations, only a few accounted uniquely for as much as 5% of total variance.

Surface Orientation Judgments. – Reports for high eccentricity ellipsoids were much more inaccurate than those for low eccentricity ellipsoids. Error means in degrees were 21.0 for high eccentricity ellipsoids and 13.4 for low eccentricity ellipsoids, yielding $F(1,4)=23.7$, $p<0.01$, accounting for 47.4% of the total variance.

– Judgements in oblique illumination were somewhat more difficult than those in direct, with error means of 18.7 and 15.6, respectively, $F(1,4)=22.6$, $p<0.01$, accounting for 7.5% of the total variance.

– The comparison of shiny displays with dull displays and cast shadow present displays with cast shadow absent displays produced nothing interpretable as even a trend. In fact the surface gloss and cast shadow factors *combined* accounted for less than 0.2% of the total variance.

The ANOVA just reported probed the tendency for the surface orientation judgments for an entire display to be affected by the presence of highlights. While no such effect was found, it would be logically possible for performance on those points within the highlight region of a display to be worse than performance for those points in dull regions. Accordingly, the intensity at each probed point on shiny surfaces was broken down into its Lambertian and specular components in order to determine whether surface orientation performance degraded as a function of increased shininess at a point. A simple linear regression of error on highlight intensity showed a tiny but statistically significant *improvement* of performance as highlight intensity increased. Highlight intensity was measured on the computer's 0–255 scale and error, as usual, in degrees. The regression yielded $B=-0.020$, $t(1038)=-2.88$, and $p<0.005$, accounting for only 0.8% of the total variance. Clearly performance was not adversely affected in highlight regions[3].

Direction of Illumination Judgments. – Cast shadows impaired performance in displays with direct illumination and improved performance in oblique illumination. Error means in degrees across subjects in the direct illumination condition were 13.0 for cast shadow absent and 18.4 for cast shadow present. Corresponding means for oblique illumination were 22.6 and 12.8, producing $F(1,4)=119.2$, $p<0.001$, and accounting for 11.1% of the total variance. Further analysis of the reported slant of the illuminant revealed that this interaction occured simply because cast shadows make any illumination appear more oblique (that is, to have higher slant). Increasing perceived slant would thereby improve performance in those displays where it was otherwise being underestimated, but it would adversely affect performance on those displays that actually had low displayed slant.

– High eccentricity displays were somewhat more difficult than low eccentricity ellipsoid displays. Error means for high eccentricity ellipsoids and low eccentricity ellipsoids were 20.0 and 12.8, respectively, with $F(1,4)=13.0$, $p<0.05$, and 8.9% variance accounted for.

– Highlights were detrimental. Means for shiny and dull displays were 19.5 and 13.3, respectively, with $F(1,4)=9.9$, $p<0.05$, and 6.5% variance accounted for.

Relation of Performance on the Two Tasks. For every presentation of a display the mean of 13 surface orientation error scores was compared with the error score for the illuminant direction judgment. The overall correlation of accuracy between the two tasks was 0.27, with $p<0.001$. When this correlation was run separately for low eccentricity ellipsoid and high eccentricity ellipsoid displays, however, a striking divergence was apparent. For low eccentricity ellipsoids the correlation was 0.31, $p<0.01$, while for high eccentricity ellipsoids the correlation was -0.07, not significantly different from zero.

The juxtaposition of the patterns of results for the two tasks analysed also has noteworthy features. First, the effect of cast shadows on illuminant direction judgments had no analog in surface orientation judgments. Second, highlights impaired direction of illumination judgement accuracy without affecting surface orientation judgment accuracy. Finally, the shape manipulation (high eccentricity ellipsoids vs. low eccentricity ellipsoids) produced a far greater effect on surface orientation judgments than on illuminant direction judgments.

Recovering Perceived Solid Shape. In order to understand how surface orientation judgments can be related to perceived shape, the display generation and data collection procedures must be examined in detail. The shapes and positions of ellipsoids were specified for the computer shading algorithm in the following way. The overall orientation of any ellipsoid in space can be specified through the slants and tilts of its semiaxes, as shown in Fig. 1. The shape of an ellipsoid is captured by the lengths of its semiaxes, and Fig. 3 shows orthographic views of the high eccentricity ellipsoid and low eccentricity ellipsoid shapes actually used in the experiment. Figures 4 and 5 depict the orientations of the axes of the ellipsoids relative to the display screen and the line of sight.

Both shape and orientation of ellipsoids are governed by appropriate choices of coefficients in the following equation for general quadric surfaces (Mingolla and Todd 1984):

$$a_{11}x^2+a_{22}y^2+a_{33}z^2+2a_{12}xy+2a_{13}xz+2a_{23}yz \\ +2a_{14}x+2a_{24}y+2a_{34}z+a_{44}=0\,. \qquad (2)$$

Moreover (2) is used to determine the surface normal at every image point for the computer shading algorithm, which employs (1).

3 Shininess and local mean curvature were somewhat correlated in the displays used in the present experiment, with flatter regions tending to be somewhat shinier. Describing the combined effects of these variables on surface judgments is beyond the scope of this article, but Mingolla (1983) analyzed this interaction in detail. Once again the effects were tiny, and consistent with the claim that shiny regions are no more problematic than dull ones

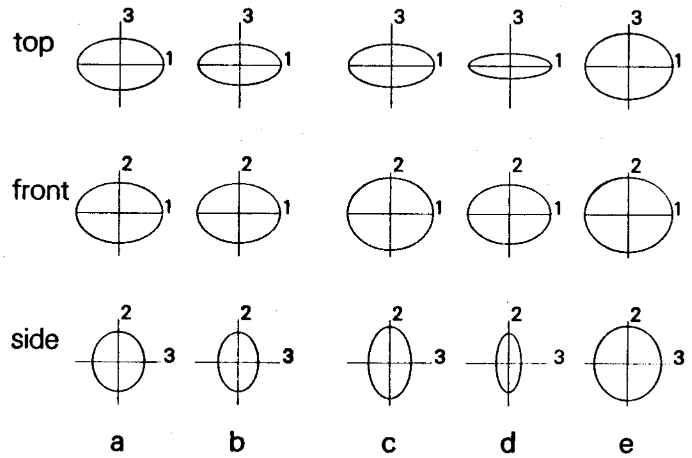

Fig. 3. Orthographic views depicting the shapes of the low and high eccentricity ellipsoids that were displayed to subjects are shown in columns a and c, respectively. Column b shows corresponding views for the recovered shape from all subjects' pooled data for low eccentricity displays. Columns d and E show the recovered shape from data given by Subjects 2 and 3, respectively, for high eccentricity displays. The most noteworthy comparison of columns a and b is the contraction of axis 3, the one most nearly aligned with the line of sight. Thus subjects saw less depth than was depicted. Columns d and e indicate the extent to which different subjects experienced different shapes when viewing the high eccentricity displays

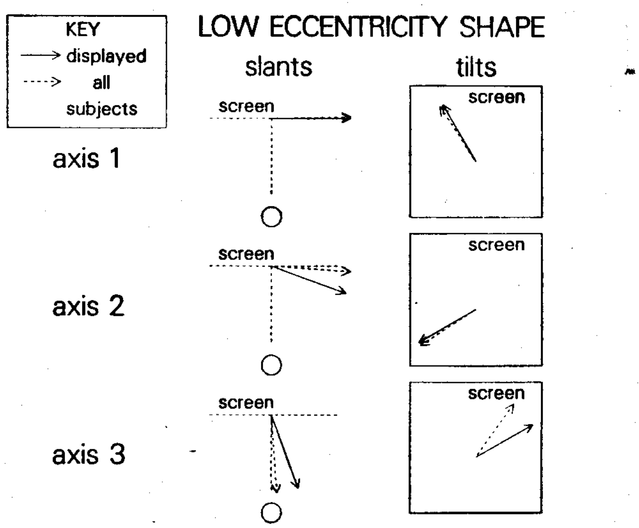

Fig. 4. The orientations of axes for the displayed and recovered shapes for low eccentricity displays show that subjects were quite accurate in determining the tilts of the ellipsoid axes. The comparison of displayed and recovered slants for axes 2 and 3, however, shows that subjects experienced the orientation of the axes as more closely aligned with the display screen and the line of sight than was the case in the display

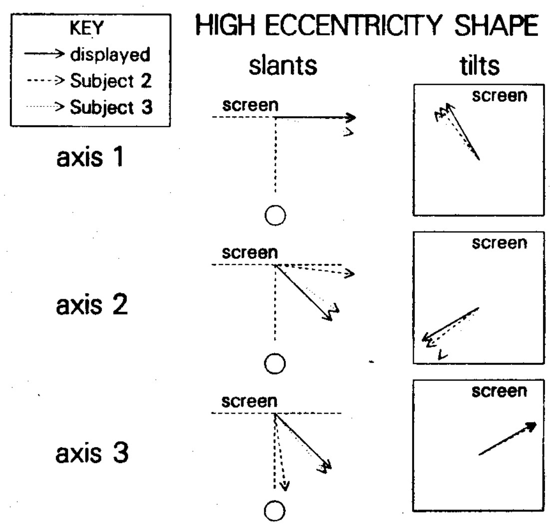

Fig. 5. As in Fig. 4, the orientations of axes for displayed and recovered high eccentricity shapes show accuracy in judgment of tilt. Subject 2 reported slants for axes 2 and 3 that were quite different from those of the displayed ellipsoid and nearly aligned with the display screen and the line of sight. The slants recovered from the data of subject 3 were more in accordance with the displayed slants

The subjects' task in estimating local surface orientation is precisely to indicate a local surface normal through its slant and tilt. For the 13 points probed of a given shape then, two sets of surface normals exist, those computed from (2) with appropriate coefficients, and those reported by subjects. The discrepancy in degrees between the pair of normals for any point is the error at that point. The fitting procedure seeks find new values for the coefficients of (2) such that the resulting discrepancy or "error" between the resulting normals and the subjects' reported normals is minimized. In other words, for display generation one starts with desired coefficients and computes normals; for perceived shape recovery one starts with normals and computes those coefficients most consistent with the normals. The resulting coefficients are most interesting. While space constraints prevent reporting all the results graphically, a few cases illustrate the power of the method.

For low eccentricity displays, the recovered shapes were so similar across subjects that data were pooled and a single composite perceived shape computed. As shown in Fig. 3, the reported shape was similar to that displayed, with the chief difference being a foreshortening of axis 3, the one most nearly aligned with the line of sight. Figure 4 displays the computed and reported axis orientations. The most noteworthy result was that the reported orientations of axes 2 and 3 were more nearly aligned with the display screen and line of sight, respectively, than was the case for the displayed orientations.

The data for high eccentricity displays were not nearly as consistent as those for low; in fact, two subject reported qualitatively different shapes than the other three. Accordingly, the two extreme cases are chosen for illustration. (Data for recovered shape is numerically summarized in Table 2.)

Subject 2 reported an extremely flattened shape, as shown in Fig. 3. While this fit the spirit of the experiment's shape manipulation, this subject's reported axis slants deviated markedly from those displayed. Here again the tendency for axis slants to align with the display screen and the line of sight was pronounced (Fig. 5). For high eccentricity ellipsoids, however, such an alignment was more at odds with the displayed orientation.

Subject 3, on the other hand, reacted differently to the high eccentricity displays. His reported shape was essentially a prolate spheroid, one axis being longer than two other nearly equal ones. This subject's reported semiaxis orientations, on the other hand, more nearly reflected the overall oblique orientation of the ellipsoid (Fig. 5).

The power of the fitting procedure can be appreciated by comparing the subjects' raw error scores with residual errors, as shown in Table 3. The residual error describes the discrepancy between the subjects' surface orientation judgments and the appropriate surface orientations for the best fitting ellipsoids. As Table 3 shows, the residual errors are quite small, both in proportion to the raw errors and as absolute angles[4].

The authors have often been asked whether the recovered ellipsoids are indistinguishable from the displayed ellipsoids when rendered according to the same computer graphics shading procedure. Certain paradoxes are latent in even asking this question, since one criterion of "distinguishability" would involve rerunning the entire experiment on shaded renditions of the recovered ellipsoids. Still, when a shaded rendition of a recovered shape is generated, its "appearance" is generally different from that of the

4 The shape recovery procedure was also performed on individual subjects' data, broken down by illuminant direction as well as by shape. The pattern of shapes and axis orientations was similar to those reported in Table 3. While, for a given subject, the displayed and recovered shapes and orientations usually differed somewhat, no pattern of such differences was evident

Table 2[a]. Axis lengths and orientations for displayed and recovered ellipsoids

		High eccentricity shapes			Low eccentricity shapes		
Data	Axis	Length	Slant	Tilt	Length	Slant	Tilt
Displayed	1	100	90	120	100	90	120
	2	84	45	210	70	70	210
	3	50	45	30	60	20	30
Subject 1	1	96	90	128	95	85	307
	2	69	77	218	67	86	216
	3	48	13	39	41	6	92
Subject 2	1	96	89	127	97	89	123
	2	69	83	217	66	84	214
	3	28	8	28	34	6	22
Subject 3	1	101	78	133	98	84	129
	2	77	53	232	76	82	28
	3	81	40	29	57	10	256
Subject 4	1	111	76	119	102	89	115
	2	77	84	27	71	80	205
	3	89	15	275	53	10	22
Subject 5	1	92	73	308	96	80	304
	2	67	71	212	71	90	214
	3	56	25	77	54	10	126

a Semiaxis lengths (distance from center to surface) are reported in screen pixel units, with 16 pixels equaling 1 cm. Slants and tilts are reported in degrees, as described in Fig. 1. The axes for each data set are ordered so that the first two are more nearly aligned with the picture plane than the third, which points most nearly along the line of sight. By symmetry, tilt values differing by 180° describe the same orientation relative to the picture plane. By convention the value given describes the direction in which the axis end that is closer to the observer points

Table 3. Mean raw and residual errors by shape and subject in degrees

Data set	High eccentricity shapes		Low eccentricity shapes	
	Raw	Residual	Raw	Residual
Subject 1	16.6	4.3	11.3	5.7
Subject 2	19.7	3.9	14.6	4.2
Subject 3	23.9	6.1	15.5	6.9
Subject 4	22.7	3.7	9.1	3.9
Subject 5	22.0	4.9	16.1	4.1

corresponding originally computed display – at times strikingly so. The differences are more evident in oblique illumination than direct, and generally much more so for shiny surfaces than for dull, as anyone with access to shaded graphics equipment can verify. This difference of "appearance" may be largely dependent on simultaneous display of the "original" and "recovered" ellipsoids, however, and pairwise comparison was not part of the present experiment's methodology. Moreover, a difference of "appearance" can mean, for example, simply a difference of the steepness of a luminance gradient near an ellipsoid's edge or a difference in the location of a highlight. Nearly identical perceived shapes can easily be recovered from displays with different "appearances", as in the case of shiny and dull displays in the present experiment. It seems that methodological considerations favor comparing displayed and recovered shapes using some medium other than the experimental shading variables themselves.

Discussion

The manipulation of shape and shading variables in an experimental procedure involves many factors, and the present experiment was able to employ only two levels of several key variables. While generalizations about many different shapes or shading conditions are not possible from the data, specific predictions of formal models can be measured against the subjects' performance. Each of questions raised in the experiment's introduction can now be examined.

Is the Formal Assumption of Lambertian Surfaces Psychologically Plausible for Shape Detection? The present experiment supports and extends earlier indications that humans are not led astray by the presence of specular reflection when judging shape or surface orientation (Todd and Mingolla 1983). Both the ANOVA for performance broken down by display variables and the regression of accuracy of surface orientation judgments on amount of highlight within a picture confirmed this. The absence of any performance deficit for shiny displays in an experiment clearly possessing the power to detect such deficits, coupled with the absence of other sources of knowledge for the subjects to compensate for the shininess, argues that humans simply do not employ a "default" Lambertian reflectance assumption. In its strongest interpretation, this finding would suggest that, when perceiving shape, the visual system is *not* attempting to solve for the parameters of surface reflectance and position which generate specific luminances. This conclusion does not entail a denial that the visual system can detect reflectance characteristics. Noticing that a surface is shiny, glossy, or dull is a natural part of seeing. The present experiment suggests that determining or knowing in advance an object's surface reflectance properties is not a prerequisite for determining shape.

Is an Estimate of the Illuminant Direction Necessary for Analyzing Shape? Patterns of errors on the two tasks, illuminant direction and surface orientation judgments, indicated a dissociation of the two processes, since cast shadows, highlights, and object shape had different effects for each task. Moreover the correlation of performance on illuminant direction estimation with surface orientation was weak overall and zero for high eccentricity ellipsoids. The design of the experiment is such that, if there had been a high correlation of performance on the two tasks, serious obstacles would exist to imputing causal priority to illuminant direction estimation. The theoretically important point, however, is that models which take illuminant parameters as essential inputs to surface orientation estimation implicitly predict strong relations in performance on the two tasks. Such relations were scarcely if at all present in the reported experiment. Once again with illuminant direction as with surface reflectance, the visual system does not seem to be projecting hypotheses about a key element of the physics underlying the generation of local intensity gradients.

The indications of the present experiment regarding the human visual system's fundamental approach to determining solid shape are in some ways counterintuitive. Since these results contradict certain of the most basic assumptions of existing computational analyses, a brief review and restatement is in order. Many possible values of surface orientation, illuminant direction and surface reflectance can result in identical values of luminance at any pixel in an image. The local mapping from pixel intensity to surface orientation, illuminant direction, and reflectance is therefore inherently one-to-many and unsolvable, unless knowledge external to the image itself is brought to bear or unless the mapping from image to shape, illuminant, and reflectance properties is further constrained. A typical move in computational analyses is therefore to introduce assumptions or knowledge about the image formation process until the number of unknowns can be reduced to the number of equations available from the image. This approach rests on the firm conviction that the image formation process *must* be inverted in order to determine shape. Often a shape from shading algorithm incorporates assumptions about illuminant direction or surface reflectance. This ordering is arbitrary, however; thus Pentland (1981) makes weak assumptions about the overall distribution of surface orientations in an image in order to infer illuminant direction.

The present experiment suggest that the human visual system does not attempt to invert the image formation process *at all*, whether it is determining surface orientation, illuminant direction, or surface reflectance. It seems that the visual system extracts impressions of these visual properties directly from the input, with no mediating steps that attempt to represent what condition in the world *could* have produced the observed luminance pattern according to some internally embodied theory of image generation. If this is so, the one-to-many mapping problem still needs accounting. For this reason analyses of global and contextual constraints on visual information and on visual processes are paramount.

Is Information for Solid Shape Fundamentally Local or Global? The present experiment does not frontally address this difficult question, but it strongly suggests that shape information is fundamentally global. To review briefly, local approaches, such as those of Horn (1975, 1977) and Pentland (1982b), take the unit of shading information to be the gradient of image intensities at a point. For a global approach, such as that of Koenderink and van Doorn (1980, 1982a), the unit of analysis is a pattern of isophotes whose size is determined by closure and connectivities of the contours themselves.

The high vs. low eccentricity shape manipulation addresses the theoretical issue indirectly. When the two shapes were chosen at the outset of the experiment, the intent was to insure that there would be some way to know if subjects could reliably distinguish two related shapes at all. No attempt was made to make one shape more difficult than the other, and the great discrepancy in performance on high vs. low eccentricity shapes came as somewhat of a shock. Further, the total pattern of performance summarized in Table 2 indicates that subjects were not able to distinguish the two shapes very well. Moreover, as shown in Figs. 3–5 and indicated in Table 2, subjects showed a strong tendency to see shapes as flatter than they were computed to be along the line of sight, resulting in a sort of "regression into the picture plane." This flattening was also accompanied by an even stronger tendency to underestimate the overall orientational asymmetries of the ellipsoids' axes.

These findings can be examined from the perspective of the two types of formal models. The local approach predicts none of the effects just stated. That is, the local intensity values and gradients for the two shapes have essentially similar distributions, so there is no basis for predicting performance differences. Likewise, flattening and orientational symmetry effects have no reasonable correlates in local computations; the effects are almost by definition global and configurational. With respect to the Koenderink and van Doorn (1980, 1982a) approach, it should first be noted that the two shapes employed in the experiment generate isophote patterns with the same topology (Fig. 6). In fact, qualitatively different isophote profiles of the kind analyzed by Koenderink and van Doorn were not manipulated in the present experiment, and they do not offer explicit predictions about the effects of slightly varying the placements of isophotes without changing the topology, as was done in the present experiment. Still it is fair to conclude from the lack of such attention that they would not expect such manipulations to reliably help distinguish shapes. A

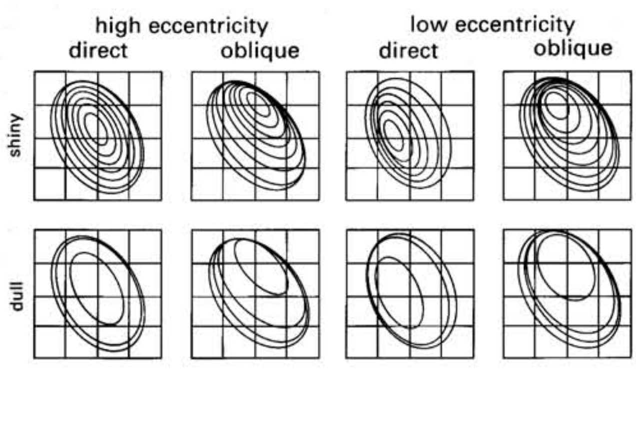

Fig. 6. Isophote patterns for the surfaces displayed in the present experiment are grouped by three experimental factors. (Cast shadows are ignored.) Digitization effects cause discontinuities in certain isophotes; ideally all would be closed contours. The thickening of the isophotes in certain regions is another digitization artifact, but it does serve to indicate *shallowness* in the local intensity gradient

similar negative inference is available regarding the tendency for the perceived ellipsoids to be oriented more symmetrically with respect to the display screen that the computed ones, for no information supporting such asymmetry is claimed to exist in the isophotes. Moreover, as the displayed shape whose axes had the more oblique orientations, high eccentricity ellipsoids would be expected to be more difficult to accurately perceive.

The preceding discussion has overstated the "local" nature of the Horn and Pentland approaches for the sake of clarity. Both require global consistency across local parameter estimates and can thus support propagation of image information across an inferred surface. Koenderink and van Doorn, however, have identified certain natural and necessary constraints of ecological optics. The constraints of the Horn and Pentland approaches arise from treatment of boundary conditions at image extrema or from smoothing procedures and are not intrinsic to the character of the initial local analyses. On the other hand, Koenderink and van Doorn have not described an actual implementation of their formal analysis. Running competing models on the same images and comparing the reconstructed shapes with those obtained from the performance of human observers could produce definitive answers concerning the models' psychological validity. Until recently marshalling the resources for such an effort would have been prohibitive; the present experiment is a step toward full implementation of this method.

Conclusion

Notwithstanding the disappointing history of research in shape perception, the present work shows that the perception of solid shape from shading is experimentally tractable. Moreover, with present technology experiments which challenge and inform the most sophisticated formal analyses can be conducted. Through manipulation of shading and other optical variables in the manner of the present experiment, psychologists can probe the perception of objects and surfaces that are considerably more complex than those manipulated up until now in the laboratory. Fitting procedures such as the one employed in the present experiment enable the psychologist to visualize the shapes that subject experience in ways that no numerical measure of experimental "error" can afford.

Although by nature a preliminary study, the present work already strongly indicates that several common assumptions of certain formal models are not psychologically valid. The most notable of these assumptions are that the visual system initially assumes that all surfaces have Lambertian reflectance and that illuminant direction must be known before shape detection can proceed. These assumptions are often accompanied by a third assumption that surface orientation is detected locally, and global shape determined by smoothing over local surface orientation estimates. The present experiment indicates that an alternative approach offered by Koenderink and van Doorn may be more psychologically accurate, as it avoids all three assumptions. Its alternative assumptions and predictions were not explicitly challenged by the present experiment, however. More experiments of the type described in this paper will help further assess the perceptual validity of shape from shading models.

Acknowledgements. The authors express gratitude to Domina Eberle Spencer and Michael Cohen, without whom the reported display manipulation and shape recovery procedures, respectively, would not have been performed.

References

Blinn JF (1977) Models of light reflection for computer synthesized pictures. Comput Graphics 11 : 192–198

Braunstein ML, Andersen GJ, Riefer DM (1982) The use of occlusion to resolve ambiguity in parallel projections. Percept Psychophys 31 : 261–267

Cohen MA, Grossberg S (1984) Neural dynamics of brightness perception: features, boundaries, diffusion, and resonance. Percept Psychophys 36:428–456

Cutting JE, Millard RT (1984) Three gradients and the perception of flat and curved surfaces. J Exp Psychol [Gen] 113 : 198–216

Grossberg S (1983) The quantized geometry of visual space: the coherent computation of depth, form, and lightness. Behav Brain Sci 6 : 625–692

Hilbert D, Cohn-Vossen S (1952) Geometry and the imagination. Chelsea, New York

Horn BKP (1975) Obtaining shape from shading information. In: Winston PH (ed) The psychology of computer vision. McGraw-Hill, New York

Horn BKP (1977) Understanding image intensities. Artif Intell 8 : 201–231

Koenderink JJ, van Doorn AJ (1979) The structure of two-dimensional scalar fields with applications to vision. Biol Cybern 33 : 151–158

Koenderink JJ, van Doorn AJ (1980) Photometric invariants related to solid shape. Opt Acta 27 : 981–996

Koenderink JJ, van Doorn AJ (1982a) Perception of solid shape and spatial lay-out through photometric invariants. In: Trappl R (ed) Cybernetics and systems research. North-Holland: Amsterdam

Koenderink JJ, van Doorn AJ (1982b) The shape of smooth objects and the way contours end. Perception 11 : 129–137

Lappin JS, Doner JF, Kottas B (1980) Minimal conditions for the visual detection of structure and motion in three dimensions. Science 209 : 717–719

Mingolla E (1983) Perception of shape and illuminant direction from shading. Unpublished doctoral dissertation, University of Connecticut

Mingolla E, Todd JT (1984) Computational techniques for the graphic simulation of quadric surfaces. J Exp Psychol [Hum Percept] 10:740–745

Pentland AP (1981) Finding the direction of illumination. J Opt Soc Am 72:448–455

Pentland AP (1982a) The visual inference of shape: computation from local features. Unpublished doctoral dissertation, Massachusetts Institute of Technology

Pentland AP (1982b) Local shading analysis. Technical Note 272, SRI International

Rock I (1973) Orientation and form. Academic Press, New York

Todd JT, Mingolla E (1983) Perception of surface curvature and direction of illumination from patterns of shading. J Exp Psychol [Hum Percept] 9:583–595

Zusne L (1970) Visual perception of form. Academic Press, New York

Received: February 20, 1985

Dr. Ennio Mingolla
Research Associate
Center for Adaptive Systems
Boston University
111 Cummington Street, Second Floor
Boston, MA 02215
USA