

Estimating local shape from shading in the presence of global shading

RODERIK G. F. ERENS, ASTRID M. L. KAPPERS, and JAN J. KOENDERINK
University of Utrecht, Utrecht, The Netherlands

In theory, global shading may help with the estimation of local surface structure from shading (e.g., in specifying the illuminant direction). Empirically, we do not know whether human observers combine the information given by the local and global shading to estimate local shape.

Observers had to indicate the orientation of a local elongated perturbation with or without global shading information provided by a background surface. Our psychophysical results show the following:

1. Observers do not estimate the orientation of the local perturbation more accurately with global shading information than they do in the absence of such information.

2. Responses depend dramatically on the inclination between the illuminant direction and the viewing direction. For an inclination of 20°, observers indicate more or less the orientation of the local ridge; however, for an inclination of 40°, they indicate either the direction of the illuminant or an orientation close to the shadow edge of the perturbation. Most subjects show some combination of these behaviors. This behavior is not altered by global shading information.

We conclude that in our paradigm, global shading information does not aid the estimation of local shape.

In most theoretical work on "shape from shading," the three-dimensional (3-D) structure of a surface is calculated from the local luminance distribution. (Horn, 1975, 1977; Ikeuchi & Horn, 1981; Oliensis, 1991; Pentland, 1982, 1984, 1989). However, the shading of an object is not determined just by the surface orientation but also by the illuminant direction and the nature of the surface material. It is therefore impossible to estimate the local shape uniquely from a local shading pattern. Constraints necessary to find local surface structure from patterns of shading can be given by the global shading of the surrounding background surface. For example, the global shading can be used to estimate the illuminant direction (Oliensis, 1991; Pentland, 1982).

Some of these shape-from-shading theories are proposed for human vision (Pentland, 1982, 1984, 1989). However, we do not know empirically whether human observers combine both local and global shading to estimate local shape. In our experiments, we report on the extent to which observers can indicate the orientation of a local ridge on a background surface (examples of the set of stimuli are given in the first column of Figure 1). The global shading information is controlled by the background surface; a spherical background surface reveals global shading information, whereas a planar one does not (equal gray value). In this paper, we investigate

whether the responses of observers are altered by the nature of the background surface.

Furthermore, we investigate how the responses of the observers depend on the illuminant direction. What is the influence of the inclination between the illuminant direction and the viewing direction on the ability of observers to indicate the orientation of the local ridge? Do observers respond differently when the illumination is aligned with the orientation of the ridge than if it is perpendicular to it?

The results of our experiment will be discussed in the view of several shape-from-shading models proposed for human vision.

Global Shading Information

Several computational shape-from-shading models use the global shading as an aid to the estimation of local shape (Ikeuchi & Horn, 1981; Oliensis, 1991; Pentland, 1982, 1984, 1989).

A model proposed by Pentland (1982) uses the shading at the occluding boundary of the shape to estimate the illuminant direction. Near the occluding boundaries of a solid shape, the surface normals are known—they are perpendicular to the viewing direction. By integrating the shading information along the occluding boundary, one obtains a rather good estimate of the illuminant direction (Pentland, 1982). Another shape-from-shading model that uses the shading pattern at the occluding boundary to estimate local shape was proposed by Ikeuchi and Horn (1981). The luminance distribution near the occluding boundary is used to start the iteration process of the surface reconstruction.

Other global shading information is given by the shadow edge (or shadow contour) on the surface. The shadow edge

This research was supported by the SPIN project, "3D-Computer Vision," of the Dutch Ministry of Economic Affairs. The authors are affiliated with the Utrecht Biophysics Research Institute. Correspondence should be addressed to R. G. F. Erens, University of Utrecht, Princetonplein 5, 3584 CC Utrecht, The Netherlands (e-mail: roderik@fysae.fys.ruu.nl).




Figure 1. The first column in Figure 1 gives an impression of the three types of surfaces that we used in our experiment. The second and third columns show the shaded stimuli for an inclination of 20° and 40°, respectively. The azimuth of the illuminant direction is 45° in each image. The elongation of the bump is 1.4, and it is oriented in all images along the y-axis. In the experiment, the orientation of the elongation and the azimuth of illuminant direction are randomly chosen from a uniform distribution (see text). The stimuli are displayed frontoparallel on an 8-bit color monitor and viewed monocularly.

is the boundary of the illuminated area. Here the illuminant direction is tangent to the surface. Oliensis (1991) proposed a shape-from-shading algorithm that uses this information.

Thus, the global shading of the background surface provides information that in theory helps in the estimation of the local surface structure. Generally, the direction of the illumination is obtained from the global shading pattern (Oliensis, 1991; Pentland, 1982). Pentland (1984, 1989) has applied his theory of local shape from shading to human vision. Psychophysical experiments have shown that human observers can estimate the illuminant direction from a shading pattern (Pentland, 1982, 1984; Todd & Mingolla, 1983). This does not imply that the human visual system uses the information about the illuminant direction to find the local shape of an object or that observers use the global shading information together with the local shading pattern to estimate local shape.

METHOD

Stimulus

In our experiments, the 3-D shapes consisted of a local perturbation on either a spherical or planar background surface. An example of the shapes is given in the first column of Figure 1. The spherical background surface reveals shading information, whereas

the planar background surface (uniform gray tone) does not. The perturbation is either an elongated Gaussian bump or a dent on the background surface.

The Gaussian perturbation is defined by three parameters that can be varied independently of the background surface: the width (σ), the height (H), and the elongation ($E = R_{long}/R_{short}$) of the Gaussian. R_{long} and R_{short} specify the values of the long and short axis of the elliptical Gaussian, respectively.

$$z(x,y) = H e^{-\frac{1}{2\sigma^2}(R_{long}x^2 + R_{short}y^2)} \tag{1}$$

The z-axis is toward the viewer. Equation 1 describes a Gaussian perturbation on a planar surface. The Gaussian perturbation on the spherical background surface is given by the following equation:

$$z(x,y) = H e^{-\frac{1}{2\sigma^2}(R_{long}x^2 + R_{short}y^2)} + \sqrt{R_{sphere}^2 - x^2 - y^2} \tag{2}$$

R_{sphere} is the radius of the sphere—namely, 4.15° (7.2 cm from a viewing distance of 100 cm). The Gaussian perturbation is added along the z direction to the sphere. Because of this definition, the Gaussian perturbation disturbs the boundary of the sphere for large values of sigma. This will never occur when the Gaussian term is added along the normal directions of the sphere instead of along the z direction. However, for the σ of 2.2° in our setup, the disturbance of the boundary of the spherical background surface is less than 4%.

The height of the Gaussian perturbation is 4.3 cm. It is chosen so that the surface normals are maximally 45° inclined with the z direction at the inflection points of the Gaussian ridge. The bump and dent are defined on the spherical background by a positive and a negative sign of the height, respectively. In the planar situation, we have only a bump on the surface; a dent on a flat surface generates exactly the same luminance pattern by a reversal of π degrees of the azimuth of the illuminant.

The view of the surface is frontoparallel, hence the perturbation does not occlude the background surface. The diameter of the stimulus is 8.3° (14.5 cm from a viewing distance of 100 cm). The outline of the stimulus is obtained from the occluding boundary of the sphere in the case of the spherical background surface. The outline in the planar situation is given by a circular mask of the same diameter as that of the occluding boundary of the sphere.

The stimuli are rendered on an Apollo DN590 computer and displayed on an 8-bit color screen (1,024 × 1,280 pixels, 256 gray tone). The illumination contains a diffuse term with the light source at infinity and an ambient term from homogeneous background lighting. As a result of the ambient term, the luminance distribution is zero nowhere on the surface, which would be very unrealistic in daily life scenes. The “luminance value” of a pixel is calculated by the following equation:

$$I_p(x,y) = A_s L_s \max(0, N \cdot L) + A_a L_a \tag{3}$$

A_s and A_a are the attenuation factors of the infinite and the ambient light sources, respectively. L_s and L_a are the intensity values of the infinite and ambient light sources, respectively. In our setup, the product of the attenuation and the light source intensity will be constant: $A_s L_s = 0.75$, and $A_a L_a = 0.2$. The luminance value of pixel $I_p(x,y)$ will be between 0.2 and 0.95. The mean “screen luminance” is 26 cd/m² ($I_p = 0.5$). N is the surface normal and L is the unit vector pointing towards the illuminant. $L = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$, where θ is the inclination of the illuminant direction and ϕ is the azimuth (Figure 2). Thus, the inclination is the angle between the illuminant direction and the positive z-axis, and the azimuth is the angle between the projection of the illuminant direction in the frontoparallel plane and the x-axis. In our setup, the inclination of the illuminant direction is either 20° or 40°. Examples of the shaded stimuli for both values of the inclination are given in Figure 1.




Figure 2. The definition of the several angles. For reasons of clarity, the orientation of the perturbation is along the y-axis. In the experiment, the orientation is randomly chosen (see text). The inclination of the illuminant direction is the angle between the z-axis and the illuminant direction. The azimuth of the illuminant direction is the angle between the projection of the illuminant direction on the frontoparallel plane and the positive x-axis.

The Gaussian perturbation disturbs the spherical shape of the background surface. So, the shading pattern on the spherical background deviates somewhat from the shading pattern at the boundary of an undisturbed sphere under the same illuminant conditions. The deviation in shading patterns is less than 5% at the boundary of the spherical background. Thus, the shading pattern of the spherical background provides almost identical information, as in the case of an undisturbed sphere.

Experimental Procedure

The experiments were performed by 6 subjects with normal or corrected-to-normal vision. They viewed the stimulus monocularly with the dominant eye and natural pupil. Head movements were restricted by a chinrest. The experiment was performed in a totally dark room.

The subjects' task was to indicate the orientation of the elongated Gaussian perturbation. They did so by a sort of a compass needle (Figure 3), which was displayed on the computer screen together with the stimulus. The needle did not cross the stimulus itself, and the subjects rotated it by moving the computer mouse. The subjects were asked to align the compass needle with the longest axis of the perturbation. They confirmed their decision by pressing the mouse button. This task was very clear for all subjects, and no training sessions were needed. During the measurements, no feedback regarding the validity of the answers was given.

Experiment 1

In Experiment 1, we investigated the influence of the illuminant direction on the ability of observers to indicate the orientation of the local perturbation. This was done for the three stimulus types—a bump or a dent on a spherical background or a bump on a planar

background. The angle between the azimuth of the illuminant direction and the long axis of the perturbation can vary maximally between 0° (aligned with the axis) and 90° (perpendicular to the axis; see Figure 3). We refer to this angle as the "relative angle" of the illuminant direction. The inclination of the illuminant direction was either 20° or 40° . The values of the inclination of the illuminant direction were chosen so that the perturbation did not cast a shadow on the surface.

We measured the angular deviation between the subject's response and the actual orientation of the elongated perturbation. If a subject adjusted the orientation of the elongation toward the azimuth of the illuminant direction, a positive sign was assigned to the deviation; if the response deviated toward an orientation perpendicular to the azimuth of illuminant direction, a negative sign was assigned to the deviation.

In Experiment 1, we used Gaussian perturbations with two aspect ratios of the elongation ($E = 1.2$ and $E = 1.4$). For each trial, the azimuth of the illuminant direction was selected randomly from a uniform distribution, and the perturbation was randomly oriented in the center of the background surface. During a series of measurements, both values of elongation and both inclinations were presented in random order. During a series, the same type of background surface was presented. Each subject performed two series of 200 measurements.

Experiment 2

In Experiment 2, we examined how the accuracy with which subjects indicate the orientation of the perturbation depends on its elongation. We were especially interested in the influence of the spherical or planar background surface on the responses of the subjects.

We used the same stimulus set as in Experiment 1, but with a larger range of elongation of the perturbation ($E = 1.3$ to $E = 2.0$).




Figure 3. Schematic drawing of the setup of the stimulus. The "relative angle" is the angle between the orientation of the elongated perturbation and the azimuth of the illuminant direction. The deviation is the angle between the subject's response and the actual orientation of the axis of elongation. If the deviation is toward the azimuth of the illuminant direction, a positive sign is assigned to it; otherwise, a negative sign is assigned. Figure 3 shows an example in which a negative sign is assigned to the deviation. We measured the deviations of the subject's responses as a function of the "relative angle" and the inclination of the illuminant direction.

The inclination of the illuminant direction was fixed at 20°, and the azimuth was drawn randomly from a uniform distribution.

The task of the subjects was the same as in Experiment 1. Three subjects who had performed the previous experiment participated again. Three series of 200 measurements were performed for each of the three different stimuli. During a series of measurements, stimuli with randomly chosen values from the set of elongations were presented. For each trial, the orientation of the perturbation was random in the center of the background surface. The background surface remained constant during a series. Again, the angular deviation between the subject's response and the actual orientation of the elongated perturbation was measured.

RESULTS

Experiment 1

The orientation responses for 2 of the 6 subjects can be seen in Figures 4-7. As a result of our definition, the deviations in the responses of the subjects will be maximally between -90° and 90°. The responses are plotted as a function of the relative angle between the illuminant direction and the orientation of the elongation. The result will vary around 0° when a subject indicates the orientation of the perturbation correctly. The standard devia-

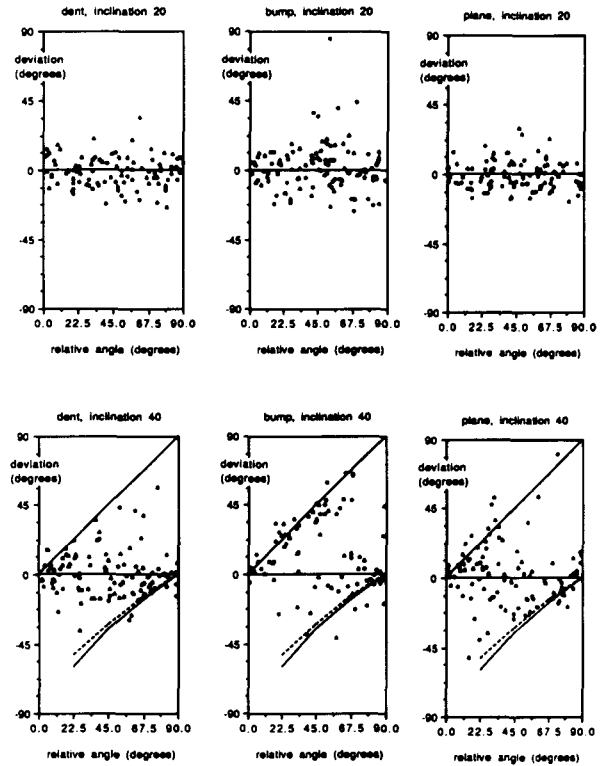


Figure 5. The orientation responses of Subject R.E. for an elongation of 1.4 of the perturbation. The top row shows that for an inclination of 20°, the orientation of the perturbation is correctly indicated. The spread in the responses is somewhat decreased in comparison with the 1.2 elongation. In the bottom row of the figures, the data from Subject R.E. indicate the axis of elongation and the illuminant direction for an inclination of 40°.

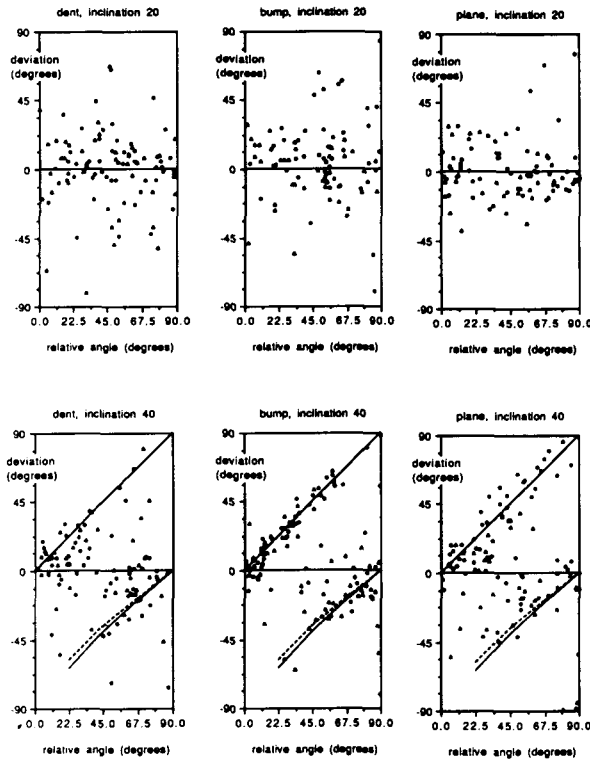


Figure 4. The orientation responses of Subject R.E. for an elongation of 1.2. The top row shows that for the inclination of 20°, the orientation of the perturbation is correctly indicated, although the spread in the responses is quite large. The bottom row shows that for an inclination of 40°, Subject R.E. no longer indicates the orientation of the perturbation but instead indicates the illuminant direction for the lower values of the relative angle and almost the shadow edge for the higher values of the relative angle.

tion of the responses is then a measure of the performance of the subject.

The top and bottom rows of Figures 4-7 show the orientation responses for the inclinations of 20° and 40°, respectively. The three columns present the three stimuli—a dent on a sphere, a bump on a sphere, and a bump on a plane. Figures 4 and 6 show the results of both subjects for an elongation of 1.2, whereas Figures 5 and 7 present the results for an elongation of 1.4.

In the top rows of Figures 4-7, with the inclination of 20°, the responses of both subjects scatter around 0°. For the elongation of 1.4 (top row in Figures 5 and 7), the response variability is somewhat less than it is for the elongation of 1.2 (top row of Figures 4 and 6). Of course, one expects that for a circular perturbation (elongation of 1.0), the responses will tend to scatter uniformly over the whole scale of -90° to 90°, because no information specifying the elongation is available. When we compare the results for the planar and spherical background surface in the case of the inclination of 20°, we find that the type of background surface hardly alters the subject's responses.

For the inclination of 20°, the subjects indicate more or less the orientation of the elongated perturbation. However, for an inclination of 40°, the responses of the sub-

jects are quite different. For this inclination, we find a "systematic deviation" from the true direction (compare the top and bottom row in Figures 4-7). Apparently, the subjects indicated some other feature in the stimulus instead of the axis of elongation. We first discuss this systematic deviation for each of the two subjects and later compare them with the response behavior of the other subjects.

One type of systematic deviation tends to the azimuth of the illuminant direction instead of to the axis of elongation. This is best shown by the results of Subject R.E. In Figures 4 and 5, the data deviate to the positive side, which means that the subject adjusted the orientation of the compass needle toward the illuminant direction. In an extreme case, when the direction of the illuminant is indicated exactly, the results will be on a diagonal from 0° to 90°. This situation is represented by the solid line in Figures 4-7. Figures 4 and 5 show that the *indication of the illuminant direction* as the orientation of the elongation occurs especially for the small relative angles between the illuminant direction and the actual orientation of the perturbation.

From the bottom row of Figures 4 and 5, one can see that the systematic deviation indicating the azimuth of the

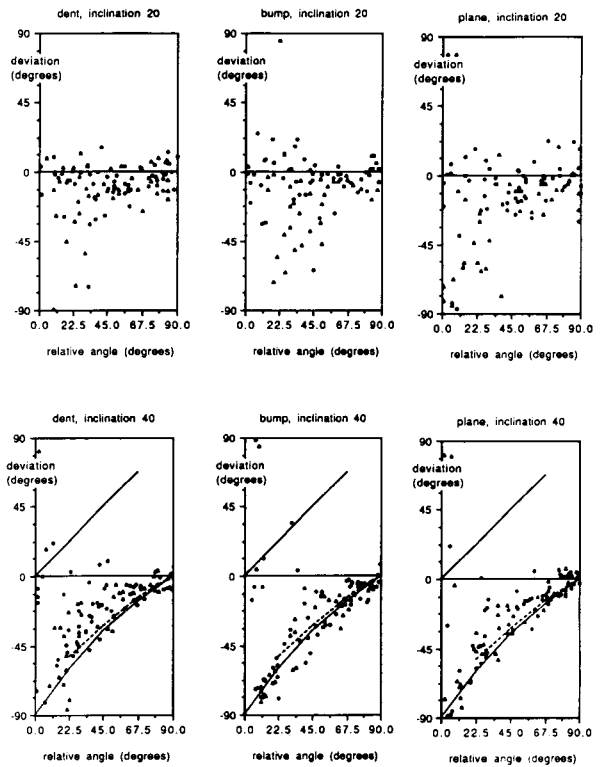


Figure 7. The data from Subject S.P. for an elongation of 1.4. The top row shows that the orientation of the perturbation is more or less indicated for the inclination of 20°. The spread is somewhat decreased in comparison with the 1.2 elongation. The bottom row shows that the data from Subject S.P. also indicate the shadow edge for the larger ratio of the elongation.

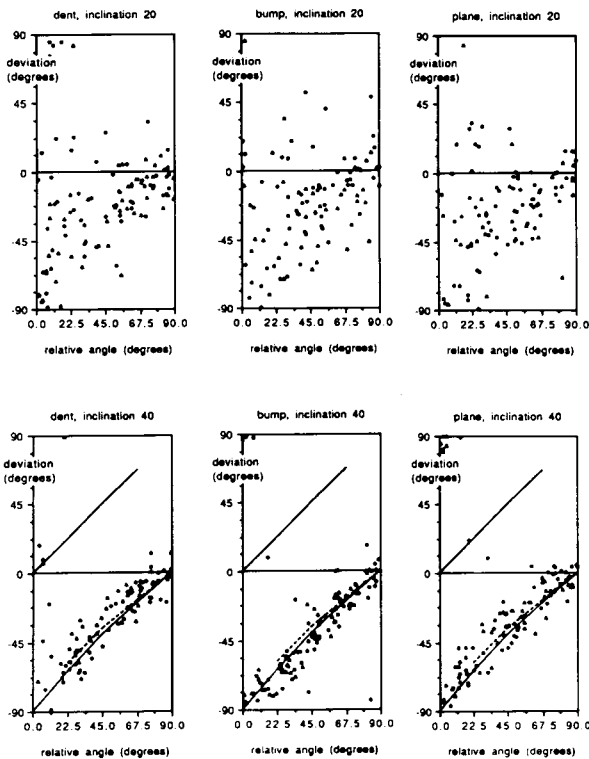


Figure 6. The data from Subject S.P. for an elongation of 1.2 of the perturbation. The top row shows that for an inclination of 20°, the orientation of the perturbation is more or less correctly indicated. There is a slight tendency to adjust the orientation of the perturbation toward the shadow edge. The bottom row shows that for the inclination of 40°, Subject S.P. indicates an orientation close to the shadow edge.

illuminant direction is stronger for the elongation of 1.2 than it is for 1.4. This systematic deviation is obtained for all three stimulus types, but is somewhat stronger for a bump on the spherical background surface than it is for a bump on the planar background. This is also found for Subject M.H. (Table 1), who tended to indicate the illuminant direction slightly stronger in the case of the spherical background surface than she did in the planar one.

The other type of systematic deviation tends to an orientation *almost perpendicular to the illuminant direction*, in which case the data would deviate to the negative side, which can be seen in Figures 6 and 7 for Subject S.P. If a subject were to make an accurate adjustment of the orientation perpendicular to the illuminant direction, the data would be on a diagonal from -90° to 0°. However, we find that the deviation is less extreme.

We suggest two obvious interpretations for the latter type of observed behavior.

A first explanation is that the *shadow edge* of the shape, rather than the orientation of the elongation, is indicated. The shadow edge is the locus of points where the illuminant direction is tangent to the elliptically elongated perturbation. The shadow edge on a spherical perturbation is always perpendicular to the illuminant direction. However, on an elliptical perturbation, it will be more aligned

Table 1
Typical Response Behavior of Subjects in Experiment 1

Subjects	Dent 1.2	Bump 1.2	Plane 1.2	Dent 1.4	Bump 1.4	Plane 1.4
M.H.	E L -	E L -	E L -	E - -	E - -	E - -
R.E.	E L S	- L S	- L S	E - -	E - -	
A.N.	- - -	E - S	E - S	- - -	E - -	E - -
J.W.	- L S	- L S	E L S	E L S	E - S	E L S
A.K.	- - -	- - S	- - S	- - -	E - S	- - S
S.P.	- - S	- - S	- - S	E - S	- - S	- - S

Note—The rows represent the typical responses for each subject. Each column represents the elongation (long/short axis) for the three stimuli. We observed three types of responses: Indication of the elongation (E), indication of the light direction (L), and indication of an orientation almost aligned with the shadow edge (S). The table contains only the results for the large inclination of 40°. All subjects indicate the orientation of the perturbation for the 20° inclination. Subjects A.N. and A.K. did not measure the dent on a sphere.

with the axis of elongation (as can be seen in the shaded pictures in Figure 1). In the extreme situation of a cylinder-like perturbation, the shadow edge and cylinder axis will be aligned exactly. In Experiment 1, this means that the shadow edge and the axis of elongation are more aligned for the elongation of 1.4 than for the ratio of 1.2. The dotted lines in Figures 4-7 represent what the subjects' response would be according to this explanation.

A second explanation is that the subjects indicated the *direction perpendicular to the largest luminance gradient* in the center of the stimulus instead of on the axis of elongation. This explanation is shown by the dashed lines in Figures 4-7.

It is clear that either explanation may account for our results. Figures 6 and 7 show that this type of systematic deviation (the *indication of the shadow edge*) occurs especially for the larger angles of the relative illuminant direction.

The data for Subject R.E. (bottom row of Figures 4 and 5) show that he indicated the illuminant direction for small relative angles and the shadow edge for the large ones. Subject S.P. indicated only the shadow edge.

In Table 1, the systematic deviations of all of the subjects are summarized. Only the results for the inclination of 40° are included, because for the inclination of 20°, all of the subjects indicated the orientation of the perturbation more or less accurately. The systematic deviations are divided into three types: indication of the axis of elongation (E), indication of the illuminant direction (L), and indication of the shadow edge (S). We assigned the responses of the subjects to one of these three categories using the following criterion. From the results shown in Figure 2, we calculated the distances between each plotted deviation angle and the lines of the suggested response types. The deviation was categorized in the response type for which the smallest distance is calculated. We assigned a category type to the responses of a subject when at least 20% of the deviations were allotted to it.

Table 1 shows that for almost all of the subjects, we found systematic deviations of the responses for the incli-

nation of 40°. This systematic deviation did not occur for the inclination of 20°. Another conclusion to be drawn from Table 1 and Figures 5-7 is that the systematic deviation of the responses of the subjects occurred for both background surfaces. The "false" indication of the azimuth of the illuminant direction was slightly stronger on the spherical background. We found no dependency on the nature of the background surface for the "false" indication of the shadow edge.

Experiment 2

For an inclination of 20°, all of the subjects tended to indicate the orientation of the axis of elongation and apparently did not show the systematic deviation that was observed for the inclination of 40°. However, their responses scattered over most of the scale. Figure 8 shows the variability of the orientation responses for 3 subjects. The standard deviations in the responses are plotted as a function of the aspect ratios of the elongations. The error bars in Figure 8 indicate the standard deviation in three repeated series of measurements. The three lines represent the three stimulus types—a bump or dent on a sphere or a bump on a plane.

Figure 8 shows that for increasing elongation, the variability of the responses decreased. For the smaller elongations, the standard deviation differed somewhat among the subjects and was around 15°. For the more elongated perturbations, the responses of the 3 subjects had similar standard deviations of about 7°. It is intuitively evident that for increasing elongations, the standard deviations in responses will decrease.

More interesting is that the standard deviations of the responses are the same for the three stimulus types. Thus, the nature of the background surface hardly alters the responses of the subjects for the range of elongations in Experiment 2.

DISCUSSION

The results of our experiments show the following:

1. The global shading information provided by the spherical background surface hardly influenced the subjects' estimates of the orientation of a local perturbation. In both Experiments 1 and 2, we found that the nature of the background surface did not alter the orientation responses of the subjects.

2. The orientation responses of the subjects depended on the inclination of the illuminant direction. For an inclination of 40°, the subjects usually indicated either the azimuth of the illuminant direction or an orientation close to the shadow edge, rather than the orientation of the elongation. This systematic deviation varied among subjects.

This second finding shows that the inclination of the illuminant direction had large effects on judgment of shape. Systematic deviations were especially evident for the smaller ratios of the perturbation. Although the luminance distribution is then almost the same as that on a




Figure 8. The standard deviation in the responses of Subjects R.E., J.W., and M.H. are given as a function of the elongation of the perturbation. The inclination of the illuminant direction is 20°; in this case, the orientation of the elongation is more or less indicated, and therefore the responses will vary around 0°. The error bars show the standard deviation in the three repeated series. The three stimulus types are represented by separate lines.

circular symmetric perturbation, we found that some of the subjects could still indicate the orientation of the elongation. This shows that the necessary information was available in the shading pattern.

It is not clear why some subjects' responses systematically deviated from the orientation of the perturbation when the inclination of the illuminant direction was 40°. The difference in shading pattern between the 20° and 40° inclination is that in the latter situation, the shadow edge is more pronounced and appears closer to the center of the perturbation (compare the images of the second and third columns in Figure 1). For large elongations (cylinder-like elongations) the shadow edge will be aligned more or less with the axis of elongation. Although this is not true for the smaller aspect ratios, some subjects still indicated the shadow edge as the axis of elongation.

On the one hand, we have found that the global shading information is apparently not used to help with the estimation of local shape from shading; on the other hand, we have found that some subjects indicate the azimuth of the illuminant direction as the axis of elongation. Theoretically, we expect that in the case of the planar background surface, the illuminant direction will be harder to find than it is on the spherical background. In the latter case, the global shading specifies the illuminant direction. This is exactly what we can see from the results of Subjects R.E. and M.H. These observers were even more wrong by indicating the azimuth of the illuminant direction instead of the orientation of the perturbation in the presence of global shading than in the absence of it. Thus, their erroneous indications of the illuminant direction were even stronger with global shading information.

In most theoretical work on shape from shading, a knowledge of the illuminant direction is essential to reconstruct the surface structure (Oliensis, 1991; Pentland, 1982, 1984, 1989). An incorrect estimate of the illuminant direction will lead to erroneous shape reconstruction. In our experiments, we found that human observers apparently did not use the global shading information to estimate the orientation of the elongation and that if they did so, their responses got even worse, as is shown by the results of Subjects R.E. and M.H. We found that all of the subjects indicated the axis of elongation quite inaccurately in the case of 20° inclination and often even completely incorrectly for a 40° inclination of the illuminant direction. Therefore, observers seem to make an incorrect interpretation of the local surface shape, and their interpretation depends on the inclination of the illuminant direction. We will compare our experimental results with several shape-from-shading models proposed for human vision.

A first model, proposed by Pentland (1984), calculates the surface orientation locally from the first and second order derivatives of the local luminance distribution. The surface reconstruction relies heavily on an accurate estimate of the direction of the illuminant. However, even with explicit knowledge about the direction of the illuminant, the surface cannot be solved uniquely (Blake &

Bülthoff, 1991; Erens, Kappers, & Koenderink, 1993). Hyperbolic and elliptic surfaces cannot be distinguished on the basis of local shading alone. An extra boundary condition in Pentland's theory is, therefore, that surfaces are locally elliptic.

Lehky and Sejnowski (1988, 1990) have proposed a neural network model that estimates the principal directions and curvatures of a surface patch from the luminance distribution. The network model is based on the same shape-from-shading theory as is Pentland's (1984) model.

Both models depend strongly on accurate information about the illuminant direction. Several psychophysical experiments have shown that human observers are able to indicate the illuminant direction quite precisely (Pentland, 1982; Todd & Mingolla, 1983). This does not imply that the illuminant direction is used by observers to estimate the surface structure from shading. Our experiments show that human observers do not use the global shading information, which specifies the illuminant direction, to estimate local shape. The human visual system does not seem to use an accurate measure of the direction of the illuminant.

Recently, Pentland (1989) proposed a shape-from-shading model, based on the implementation of a linear approximation of the reflectance function, in biological terms of receptive fields. This model accounts for his psychophysical results. The linear approximation is more correct for larger inclinations of the illuminant. For our experiments, this model implies that observers would judge the orientation of the perturbation better in the situation of 40° .

This is in contradiction with the result we obtained from our experiments. We found that for a small inclination of 20° , the observers were able to indicate the orientation of the perturbation, although the variance in their results was quite large. But for the larger inclination of 40° , the subjects no longer indicated the orientation of the perturbation. They now indicated either the illuminant direction or a direction close to the shadow edge. Thus, all subjects responded differently for the larger inclination, and their behavior depended strongly on the azimuth of the illuminant direction relative to the orientation of the perturbation. In conclusion, Pentland's (1989) model does not explain our results satisfactorily.

One theory that directly couples the shading pattern to the surface geometry without the use of the illuminant direction has been proposed by Koenderink and van Doorn (1980, 1982). The luminance distribution on a solid shape will contain several types of singularities (local extrema, such as minima, maxima, and saddle points), two of which pertain to the stimulus we have used. One type of singularity appears in the luminance distribution in which the surface normal is aligned with the illuminant direction (this is always an intensity maximum). The other type of singularity (this can be either extrema) appears on curves between elliptic and hyperbolic regions (so-called parabolic lines) and therefore conveys information about the local surface geometry.

The surface geometry of the stimulus with the spherical background contained two closed parabolic lines, whereas the stimulus with the planar background contained one closed parabolic line. Therefore, the luminance distribution of the stimuli contained singularities of the latter type. These were different for the two background surfaces. We found that the response behavior depended strongly on the inclination of the illuminant direction and was quite specific for each subject. It is not clear how these results can be explained with respect to Koenderink's and van Doorn's (1980, 1982) theory.

Several psychophysical experiments indicate that although the 3-D impression of 2-D images can be very compelling, human observers are rather poor at estimating the local surface structure in images (Koenderink, van Doorn, & Kappers, 1992; Mingolla & Todd, 1986; Todd & Reichel, 1989). Mingolla and Todd (1986) have found that observers estimate the local surface orientation of a shaded ellipsoid quite inaccurately. Recently, Koenderink et al. have shown that especially the subject's settings of the slant direction are very unprecise, whereas the tilt direction is very reliably reproduced. In our experiments, we found that the settings of the orientation of a local perturbation are surprisingly inaccurate, too. The dependency of the observers' responses on the inclination of the illuminant direction was never obtained in any of the other investigators' psychophysical experiments. However, our experimental paradigm was quite different from the other ones, so the results are hard to compare.

In summary, we conclude that the human observers were rather poor at indicating the orientation of a local perturbation on the basis of shading. The inclination of the illuminant direction has a great influence on an observer's estimate of the local shape. Global shading, which in theory can be used for determining the exact illuminant direction, does not seem to be used by the human visual system for finding the surface structure. We believe that our results cannot be explained satisfactorily by any of the models proposed for the analysis of shape from shading.

REFERENCES

- BLAKE, A., & BÜLTHOFF, H. H. (1991). Shape from specularities: Computation and psychophysics. *Philosophical Transactions of the Royal Society of London B*, *331*, 237-252.
- ERENS, R. G. F., KAPPERS, A. M. L., & KOENDERINK, J. J. (1993). Perception of local shape from shading. *Perception & Psychophysics*, *54*, 145-156.
- HORN, B. K. P. (1975). Obtaining shape from shading information. In B. K. P. Horn & M. J. Brooks (Eds.), *Shape from shading* (pp. 123-171). Cambridge, MA: MIT Press.
- HORN, B. K. P. (1977). Understanding image intensities. *Artificial Intelligence*, *8*, 201-231.
- IKEUCHI, K., & HORN, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, *17*, 141-184.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1980). Photometric invariants related to solid shape. *Optica Acta*, *27*, 981-996.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1982). Perception of solid shape and spatial lay-out through photometric invariants. In R. Trappell

- (Ed.), *Cybernetics and systems research* (pp. 943-948). Amsterdam: North-Holland.
- KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, **52**, 487-496.
- LEHKY, S. R., & SEJNOWSKI, T. J. (1988). Network model of shape-from-shading: Neural function arises from both receptive and projective fields. *Nature*, **333**, 452-454.
- LEHKY, S. R., & SEJNOWSKI, T. J. (1990). Network model of visual cortex for determining surface curvature from images of shaded surfaces. *Proceedings of the Royal Society of London B*, **240**, 251-278.
- MINGOLLA, E., & TODD, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics*, **53**, 137-151.
- OLIENSIS, J. (1991). Shape from shading: A partially well-constrained problem. *CVGIP: Image Understanding*, **54**, 163-183.
- PENTLAND, A. P. (1982). Finding the illuminant direction. *Journal of the Optical Society of America (A)*, **72**, 448-455.
- PENTLAND, A. P. (1984). Local shading analysis. *IEEE: Transactions on Pattern Analysis Machine Intelligence*, **6**, 170-187.
- PENTLAND, A. P. (1989). Shape information from shading: A theory about human perception. *Spatial Vision*, **4**, 165-182.
- TODD, J. T., & MINGOLLA, E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception & Performance*, **9**, 583-595.
- TODD, J. T., & REICHEL, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, **96**, 643-657.

(Manuscript received July 14, 1992;
revision accepted for publication February 8, 1993.)