Journal of Experimental Psychology:
Human Perception and Performance
1983, Vol. 9, No. 4, 583-595

Perception of Surface Curvature and Direction of Illumination From Patterns of Shading

James T. Todd
Brandeis University

Ennio Mingolla
University of Connecticut

Three experiments examine the perceptual salience of shading information for the visual specification of three-dimensional form. The observers in these experiments were required to estimate the surface curvature and direction of illumination depicted in computer-synthesized images of cylindrical surfaces, both with and without texture. The results indicate that the shininess of a surface enhances the perception of curvature, but has no effect on perceived direction of illumination; and that shading is generally less effective than texture for depicting surfaces in three dimensions. These and other findings are used to evaluate the psychological validity of several mathematical analyses of shading information that have recently been proposed in the literature.

> Shadows appear to me to be of supreme importance in perspective, because, without them opaque and solid bodies will be ill-defined; that which is contained within their outlines and their boundaries themselves will be ill-understood unless they are shown against a background of a different tone from themselves. And therefore in my first proposition concerning shadow I state that every opaque body is surrounded and its whole surface enveloped in shadow and light. (Leonardo da Vinci, cited in Richter, 1970)

The study of how image shading provides information about three-dimensional form has long been neglected by perceptual psychologists. Although there has been a growing amount of research on three-dimensional form perception during the past few decades, it has focused almost exclusively on other sources of information, such as motion, texture, and binocular vision. This apparent lack of interest in the effects of shading is somewhat surprising considering that manipulations of light and shade have been used to create the impression of "depth" in representational art for over 500 years. Leonardo da Vinci, for example, considered the use of shading to be of "supreme importance" for the realistic depiction of objects in space, and he wrote extensively on this topic in his now famous *Notebooks.* Indeed, the theoretical analysis presented by Leonardo has been little improved over the past five centuries.

The Experimental Control of Shading

One important factor that has inhibited research on the perception of shading is the lack of an appropriate technology for the systematic manipulation of shaded images. In order to study how patterns of light and shade provide information about three-dimensional form, a researcher must be able to control those patterns with mathematical precision. Until recently, the only practical techniques for producing a shaded image of an object were to paint a picture or take a photograph, but neither of these techniques could allow a researcher to control the lightness of each point on an image independently. During the past few years, however, these methodological difficulties have been greatly alleviated by the increasing availability of computer video systems. Unlike photographs or paintings, the images created on a computer video display are generated directly from mathematical models. The basic procedure is quite simple. An image is broken down into a rectangular grid of arbitrarily small regions called picture elements or *pixels,* and the lightness of each region is stored within a portion of the computer's memory called a *video frame buffer.*[1] These

Requests for reprints should be sent to James T. Todd, Department of Psychology, Brandeis University, Waltham, Massachusetts 02254.

---

[1] When generating an image on a color display it is necessary to specify separate intensity values for the red, green, and blue elements at each picture location, thus requiring a threefold increase in the size of the frame buffer. Some systems are designed to reduce this memory

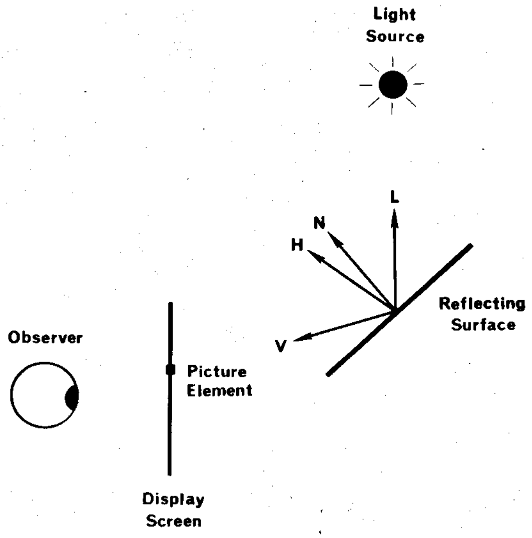

Figure 1. The geometric variables needed to simulate an image of a surface in three-dimensional space. (**L**, **H**, **V**, and **N** are all unit vectors from a point on the surface: **L** lies in the direction of the light source; **V** lies in the direction of a picture element and the point of observation; **H** bisects the angle formed by **L** and **V**; and **N** is perpendicular to the surface.)

lightness values are scanned in rapid succession by a digital-to-analog (D/A) converter and the resulting video signal is then output to the display monitor.

The main difficulty in generating a shaded image on a computer graphics display is determining the appropriate lightness values for depicting a desired scene. A number of techniques for addressing this problem have been developed in recent years, and it is now possible to simulate images of objects that are as perceptually compelling as images produced by photographic techniques (see Blinn, 1977; Blinn & Newell, 1976; Cook & Torrance, 1981; Kay & Greenberg, 1979; Phong, 1975; Whitted, 1980). The basic ingredients needed to generate a realistic shaded image are depicted in Figure 1. A surface in the environment is illuminated by light that originates from a luminous body such as an incandescent lamp or the sun, and some of the incident light energy is reflected toward the point of observation. It is assumed that a video-display screen is located somewhere between the observer and the observed surface. The intensity of each picture element on this display screen is determined mathematically by estimating the amount of light energy that would pass through its boundaries in the direction of the observation point. (For the picture element depicted in Figure 1, this direction is indicated by the vector **V**.)

The amount of light that reflects from a surface region in a given direction depends on a wide variety of physical variables, including the orientation and chemical composition of the surface, and the position and spectral composition of the light source. If a surface region is one that diffuses incident light equally in all directions, then the intensity of its corresponding picture element (I_P) can be simulated using Lambert's law: $I_P = I_L s(\mathbf{L}\text{-}\mathbf{N})$, where I_L is the intensity of the light source; s is the shade or albedo of the surface ranging from 0 (black) to 1 (white); **L** is a unit vector in the direction of the light source; **N** is a surface normal (i.e., a unit vector that is perpendicular to the surface); and "$\cdot$" denotes the vector inner product. This relation is a reasonable approximation of how light reflects from pure matte (Lambertian) surfaces, but it does not capture the highlights that are typically observed on shiny surfaces. Suppose, for example, that the surface depicted in Figure 1 were a perfect mirror. In that case reflected light would only reach the point of observation if $\mathbf{H}\cdot\mathbf{N} = 1$, where **H** is a unit vector that bisects the angle formed by **L** and **V**. For an imperfect mirror, some of the light rays would be deflected slightly from this direction of maximum highlight. The intensity of a picture element could then be simulated by $I_P = I_L g(\mathbf{H}\cdot\mathbf{N})^n$, where g is the proportion of incident light reflected in the direction of maximum highlight and n represents the sharpness of the highlight (e.g., for a perfect mirror $n = \infty$). The reflectance properties of many natural surfaces can be adequately described as an additive combination of a diffuse reflector (the Lambertian component) and an imperfect mirror (the

---

load by restricting the number of colors that can be presented at any one time. In these systems, the numbers stored in the frame-buffer memory represent addresses in a "video lookup table" where the actual intensity values are stored.

specular component.[2] Most simulations also add in a third component to represent the intensity of ambient light (I_a) that falls on a surface uniformly from all directions. Taking all of these factors into account, the intensity of a picture element is given by the following:

$$I_P = I_aS + I_LS(\mathbf{L}\cdot\mathbf{N}) + I_Lg(\mathbf{H}\cdot\mathbf{N})^n. \quad (1)$$

The model is easily extended, moreover, to accommodate multiple sources of illumination. For a configuration of k light sources, the intensity of a picture element is given by the following:

$$I_P = I_aS + \sum_{J=1}^{k} I_{L_J}(S[\mathbf{L}_J\cdot\mathbf{N}] + g[\mathbf{H}_J\cdot\mathbf{N}]^n). \quad (2)$$

It is important to keep in mind that Equations 1 and 2 can only describe how light is reflected from a single surface in isolation. The analysis of image shading becomes considerably more complex when the environment is cluttered with many different objects because the amount of light reflected from one surface can be dramatically influenced by the presence of another. One of the most obvious effects of surface interactions is the appearance of cast shadows. Shadows occur when the light rays headed toward a visible surface are occluded by an opaque object. If all of the light rays are occluded, the shadow is called an *umbra;* if only some of the light rays are occluded, the shadow is called a *penumbra* (see Figure 2). A related phenomenon occurs when viewing transparent surfaces. For example, consider the pattern of shading produced by a pool of clear water. Some light rays will be reflected from the surface of the pool. Others will be transmitted through the water (though their directions will generally change due to refraction) and will be reflected from the bottom. Both sets of reflected rays will combine to determine the intensity of a picture element (see Figure 3). Another way that patterns of shading can be affected by surface interactions is through the process of indirect illumination. Whenever a surface is illuminated, some of the incident light energy is reflected in many different directions. These reflected light rays can illuminate other objects in exactly the same way as direct illumination from a luminous body (see Figure 4). A good example of indirect illumination is the reflection of objects in a mirror. Light from a luminous body is reflected off a surface toward the mirror, which then reflects it a second time toward the point of

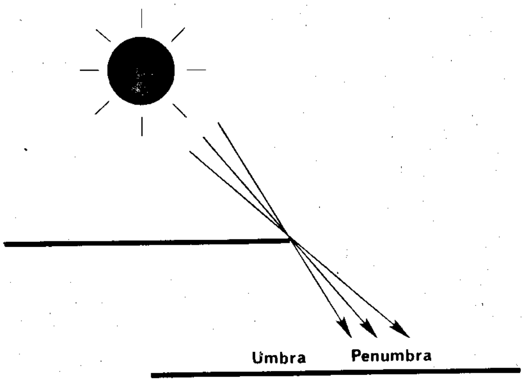

Figure 2. The formation of shadows. (A shadow is formed when the illumination of one surface is occluded by another. If the illumination from a given source is totally occluded, the shadow is called an umbra; if it is only partially occluded, the shadow is called a penumbra.)

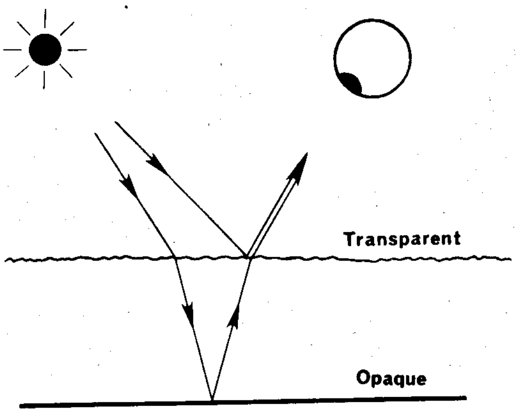

Figure 3. The effects of transparency. (When a transparent surface is illuminated, some of the incident light rays will be reflected; others will be transmitted through the surface, though their directions will generally change due to refraction.)

---

[2] The surfaces simulated with this simplified model of reflection tend to have a rather "plastic" appearance. There are other more elaborate models, however, that also take into account the roughness of a surface and its specific pattern of spectral absorption (e.g., see Cook & Torrance, 1981).

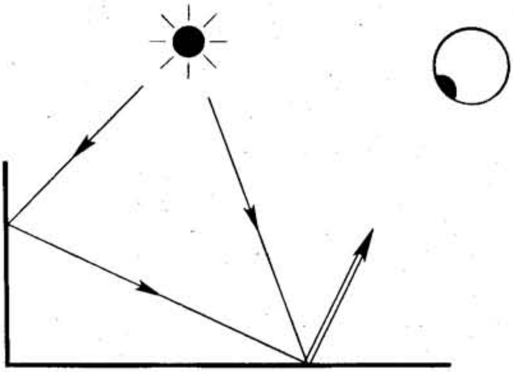

Figure 4. The effects of indirect illumination. (The light rays that illuminate a surface can come directly from a luminous body, or, they can also arrive indirectly by being reflected from some other nearby surface.)

observation (i.e. the mirror is indirectly illuminated). Another example of indirect illumination is the ambient light (I_a) described in Equations 1 and 2. Because light rays can be reflected many times in a cluttered environment, they tend to be scattered at random in all directions. The intensity of this ambient illumination is considerably reduced relative to the original source, however, because some amount of energy is lost each time a light ray is reflected. (A computer-generated image that illustrates many of these effects is provided in Figure 5.)

The Analysis of Shading

In generating a shaded image with a computer video system, one begins with a detailed description of a configuration of surfaces and light sources in three-dimensional space and then determines the pattern of intensities that would be produced on a visual-display screen when light is reflected from those surfaces toward the point of observation. The problem for a perceptual system, however, is exactly reversed; one begins with a pattern of

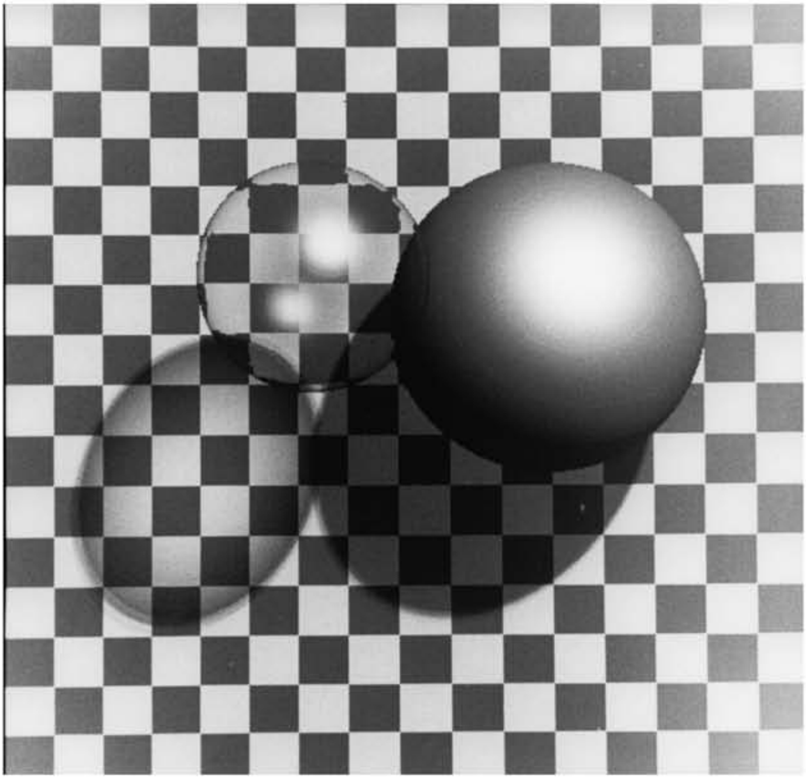

Figure 5. A computer-generated image depicting an opaque sphere with specular highlights, a transparent sphere with specular highlights for both its front and back surface, and the shadows cast by these objects on a checkerboard background.

light intensities on a visual-display screen and then determines the configuration of surfaces and light sources in three-dimensional space that would have produced that pattern in a natural environment. Were it not for the effectiveness of shading information for depicting depth in representational art, it would be tempting to conclude that this latter problem is insolvable because it involves a many-to-one mapping. That is to say, the specific intensity of any given picture element can be generated by an infinite number of possible combinations of surface position, surface orientation, surface reflectance, light-source position, and light-source intensity. If we also allow for surface interactions and multiple light sources, the problem becomes even worse. Thus, a fundamental issue for psychological theory is to adequately explain how the visual system is able to tease apart the separate effects of all these different variables.

Much of the existing research on how patterns of shading provide information about three-dimensional form has been performed by Berthold Horn and his colleagues at the MIT Artificial Intelligence Laboratory (see Horn, 1975, 1977, 1981; Ikeuchi & Horn, 1981; Pentland, 1982a, 1982b, Woodham, 1981, Note 1). Horn has proposed a number of techniques for analyzing a shaded image that are all based on a common set of simplifying assumptions. To apply these techniques, an observed surface must be far away from the observer and the sources of illumination so that its image will approximate an orthographic projection, and the intensity of each image point must be unaffected by surface interactions such as cast shadows or transparencies. These conditions ensure that a given orientation for a particular surface material can produce only one possible image intensity. Horn's analysis assumes that this set of correspondences between orientations and intensities is known for a given surface and a given configuration of light sources. However, because it is possible for many different orientations to be associated with the same image intensity, there is still not enough information to uniquely determine an object's three-dimensional structure. An adequate solution can only be obtained with the additional assumption that an observed surface is smooth, and with some sophisticated mathematical techniques for solving sets of differential equations.

Horn's pioneering research on image analysis is justly applauded in the field of artificial intelligence, but its underlying assumptions are difficult to justify for a theory of human vision. Not all surfaces in a natural environment are perfectly smooth; they need not be a long distance from the point of observation; they may be semitransparent or have shadows cast upon them; and, most importantly, the observer may have no knowledge whatsoever about the reflectance properties of any given region or the specific pattern of illumination. Because of these limitations, Pentland (1982a) has recently proposed an alternative analysis of image shading that is specifically designed to model the performance of human visual systems. This analysis has several desirable properties from the standpoint of perceptual theory. First, it requires no prior knowledge about the direction of illumination (see also Pentland, 1982b). Moreover, unlike Horn's analysis, it is based on a set of mathematical operations that could easily be implemented using physiological mechanisms that are known to exist within mammalian visual systems. When applied to actual shaded images, Pentland's analysis tends to underestimate the relief of surfaces, but his research suggests that the judgments of human observers often err in a similar manner. The analysis has other properties, however, whose perceptual validity has yet to be examined. For example, because it assumes that the environment is composed of relatively homogeneous Lambertian surfaces, and that the pattern of shading in an image is unaffected by surface interactions, the analysis is unable to cope with shaded images that contain specular highlights, indirect illumination, transparent surfaces, or cast shadows.

Although the analyses proposed by Horn and Pentland differ somewhat in their underlying assumptions, they are both based on a similar strategy in which the three-dimensional structure of an object is inferred using simplifying assumptions about the surface reflectance and pattern of illumination. There is, unfortunately, little evidence to suggest that the human visual system incorporates

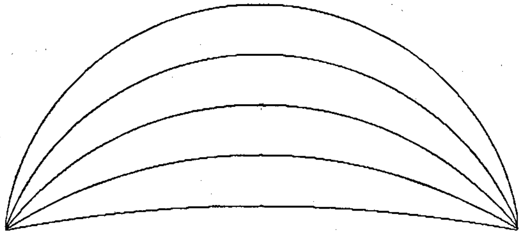

Figure 6. The five different curvatures simulated in Experiments 1–3.

a similar strategy. The few investigators who have studied the perception of three-dimensional form from shading information (e.g., Benson & Yonas, 1973; Bergström, 1977; Yonas, Kuskowski, & Sternfels, 1979; Bergström, Note 2) have not tried to determine if the accuracy of observers' judgments is significantly impaired by certain types of surface materials or illumination patterns. One set of findings that appears to contradict the analyses proposed to date is that human observers can, under certain circumstances, reliably interpret how images are affected by transparent surfaces (Metelli, 1974; Beck & Prazdny, Note 3), indirect illumination (Gilchrist, 1979, in press), or cast shadows (Hagen, 1976; MacLeod, 1932). The theoretical significance of these findings is difficult to interpret, however, because they are not specifically concerned with the perception of three-dimensional form. The present series of experiments was designed, therefore, to examine the perception of shaded images in which object structure, surface reflectance, and the pattern of illumination were all varied independently of one another. The goal of this research was to determine empirically if the mathematical limitations of existing analyses are consistent with the perceptual limitations of actual human observers.

Experiment 1

Experiment 1 examined the ability of observers to judge the curvature of cylindrical surfaces with varying degrees of specular highlights and varying directions of illumination.

Method

Subjects. Nineteen naive observers participated in the experiment to fulfill the requirements of an introductory psychology course.[3]

Apparatus. Displays were presented on a Conrac 17-in. video monitor controlled by a NOVA minicomputer. Contrast and brightness levels were set by the experimenter so that a test band of 10 intensities representing the monitor's possible range was clearly visible. The displays were viewed binocularly at a distance of approximately 100 cm. Head and body movements were not restricted.

Stimulus displays. Each video display consisted of a horizontal light-intensity gradient, approximately 28 cm wide and 17 cm high, which simulated a cylindrical surface illuminated by a narrow, vertically extended light source. Fifty displays comprised all combinations of 5 levels of curvature, 5 angles of illumination, and 2 surface reflectance functions (shiny and dull). The simulations were computed according to the shading model described by Equation 1 with the following particulars:

1. Curvature values were chosen with computational ease in mind. The width of the video screen was 512 pixels and the cylinders were modeled so that the distance from the plane of the screen to the point closest to the viewer would be 25, 75, 125, 175, or 225 pixels (see Figure 6). The virtual cylinders thus displayed had radii of 75.4, 27.2, 18.5, 15.7, and 14.7 cm. These values were chosen so that the occluding left and right edges of the cylinder were never visible. Because no top or bottom contours were displayed either, the resulting image formed a single, uninterrupted intensity gradient.
2. The angles of incident illumination were 0°, 10°, 20°, 30°, and 40° to the observer's right measured from the light source to the image center to the simulated eye position, which was always perpendicular to the image center.
3. Shiny surfaces were simulated by setting the parameters s, g, and n in Equation 1 to .5, .4, and 5, respectively, the corresponding parameter values for the dull surfaces were .5, .1 and 5. Consequently, the shiny surfaces had higher maximum intensities than the dull surfaces. In addition, all of the simulations included a low level of homogeneous diffuse illumination of approximately one tenth the maximum intensity level.
4. The simulated eye position was approximately 105 cm from the image center as was the simulated light-source position. Because the displays were designed to simulate illumination from a vertically extended light source, the intensity calculations only had to be performed once for each column of pixels on the display screen. If the displays had been modeled on a point light source, the number of required computations would have increased by a factor of 300 producing a delay of many minutes between each trial. (Figure 7 shows how image intensity varied as a function of horizontal position for a representative sample of the 50 displays.)
5. Finally, a 6-pixel border of homogeneous intensity at the middle of the range of possible intensities surrounded each image. This thin grey frame lent definition to the display gradient by providing a consistent standard

---

[3] Two additional subjects in Experiment 1 and one in Experiment 2 left the room complaining that they saw nothing like what the experimenter had described in the displays. The responses of these subjects appeared to be generated at random and were therefore excluded from the statistical analysis.

against which the gradient intensities could be better discerned. (A photograph of a typical display is given in the top portion of Figure 8.)

Procedure. Before an observer saw any of the computer-generated displays, the experimenter described verbally what they would depict. A homogeneously colored coffee can served as a model of a cylinder in an upright orientation. The observers were told that they were to judge the degrees of curvature for cross sections of different cylinders. The experimenter took a piece of pliant cardboard, bent it to approximate cylinder sections, and demonstrated how a nearly flat section would correspond to a large cylinder, whereas a highly curved section would correspond to a cylinder of smaller radius. The observers were told that they would see only a small frontal section of the surface, which the experimenter demonstrated by creating a small rectangular window with his fingers in front of the coffee can. They were then asked to imagine that the displays were life size, though this would require the cylinders to be quite large, ranging from about the size of the display monitor to larger than a trash can.

Next, the observers were shown five circular arc curves drawn on separate pieces of manila cardboard, labeled 1 through 5, whose curvatures corresponded to the ones that would be simulated in the video displays (see Figure 6). These five standard arcs were placed directly in front of the observers and remained there throughout the experiment. The observers were told that they would see a sequence of simulated cylindrical sections, and that they should rate the curvature of each one on a scale of 1 to 5 by pressing the appropriate key on the computer keyboard. They were told that the illumination direction would vary, generally coming from their right, and that the surfaces could appear either shiny or dull. Finally, the observers were warned that the displays were designed for experimental rigor rather than for pictorial realism. They were told not to worry if an image did not appear to match one of the five standards exactly or even if it did not particularly resemble a cylinder; they were merely to pick the number from 1 to 5 that best described their impression of curvature.

The observers saw a randomized sequence of the 50 displays three times in succession. The experimenter stayed in the room with the observer for the first 6 to 8 trials to answer any questions. The entire first pass of the 50 displays was treated as practice and not analyzed. The observers were told that the experiment was self-paced and encouraged to take their time without agonizing. Each display took approximately 1 sec to be completely generated on the monitor, beginning as soon as the response to the previous trial was recorded. Observers took from 25 to 50 min. to complete the experiment, including about 5 min. for instruction.

Results

Figure 9 shows the mean curvature ratings as a function of simulated curvature for both the shiny and dull surfaces. An analysis of linear regression was used to evaluate the accuracy of these ratings. For example, if the observers had been completely insensitive to shading information and had perceived all of the simulated cylinders as equally curved, the regression lines would have had a slope

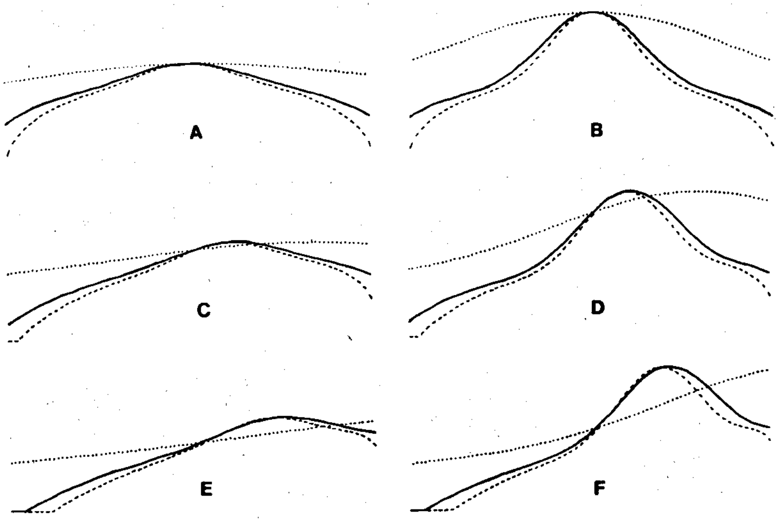

Figure 7. The effects of curvature, surface quality, and direction of illumination on the pattern of image intensities for a simulated cylindrical surface. (The height of each curve represents image intensity as a function of horizontal position for one of the displays used in Experiments 1 and 2. The low, middle, and high curvatures depicted in Figure 6 are represented by dotted lines, solid lines, and dashed lines, respectively. Dull surfaces are represented in the left column, a, c, e; shiny surfaces in the right column, b, d, f. The top row, a, b, middle row, c, d, and bottom row, e, f, represent illumination angles of 0°, 20°, and 40°, respectively, from the observer's line of sight.)

Figure 8. A representative example of (top) a shiny, untextured surface used in Experiments 1 and 2, and (bottom) a textured surface used in Experiment 3. (Both surfaces are maximally curved and illuminated from 20° to the right of the observer's line of sight. However, because of error introduced by the process of reproduction, these images appear considerably less three-dimensional than the actual video displays from which they were taken. They are included here only to provide specific examples of how the displays were constructed.)

of zero. If, on the other hand, they had responded with perfect accuracy, the regression lines would have had a slope of 1. The actual data fell somewhere in between these two extremes. The regression lines for the shiny and dull surfaces had slopes of .67 and .47, respectively.

An analysis of variance (ANOVA) confirmed all of the differences that are apparent in Figure 9. The effect of curvature accounted for 42.67% of the total sum of squares, $F(4, 72) = 323.559$, $p < .001$; the effect of specular highlights accounted for 10.24% of the sum of squares, $F(1, 18) =$

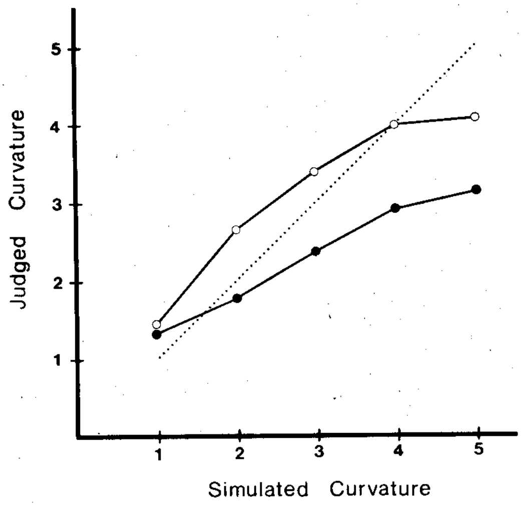

Figure 9. The mean curvature ratings of 19 naive observers as a function of simulated curvature for the untextured surfaces used in Experiment 1. (The units along each axis are arbitrary values on a 5-point subjective rating scale assigned to represent the five curvatures depicted in Figure 6. Ratings for shiny surfaces are indicated by open circles; ratings for dull surfaces are indicated by filled circles. The dotted line represents the ratings that would be produced by a perfectly accurate observer.)

60.245, $p < .001$; and the interaction of these two effects accounted for 2.36% of the sum of squares, $F(4, 72) = 23.285, p < .001$. There was also a small but highly significant effect of light direction, $F(4, 72) = 24.011, p < .001$, which accounted for 4.13% of the sum of squares. The mean curvature ratings for the 0°, 10°, 20°, 30°, and 40° directions of illumination were 2.434, 2.484, 2.629, 2.834, and 3.142, respectively.

Experiment 2

Experiment 2 used the same displays as in Experiment 1 to examine the ability of observers to judge the direction of illumination depicted in a shaded image.

Method

The apparatus and general procedure were roughly equivalent to those used in Experiment 1. Twenty naive observers participated in the experiment to fulfill the requirements of an introductory psychology course. The observers were told that they would see a sequence of video images and that they would be required to rate the direction of illumination depicted in each one on a scale of 1 to 5. In explaining the task, the experimenter placed a coffee can about 100 cm from the observer in a dimly lit room. A small light source was swung from directly over the observer's head through a horizontal arc of roughly 40° while the experimenter called out the numbers 1 through 5 at roughly equal arc intervals. The observers were cautioned that they would see only small portions of the simulated cylinders and that secondary information from cast shadows and the reflections of neighboring objects would be absent. They were also informed that the curvature of the cylinders would vary and that their surfaces could appear either shiny or dull.

Results

Figure 10 shows the mean direction of illumination ratings as a function of the simulated direction for both shiny and dull surfaces. Statistical analysis revealed that the specular highlights had no effect on the observers' judgments, $F(1, 19) = .397, p > .1$; the regression lines for the shiny and dull surfaces had slopes of .63 and .68, respectively. As is evident in Figure 10, there was a clear-cut effect of direction of illumination, $F(4, 76) = 208.042, p < .001$, which accounted for 47.13% of the total sum of squares. There were also significant interactions between surface curvature and direction of illumination, $F(16, 304) = 4.531, p < .001$, and between surface curvature and level of specular

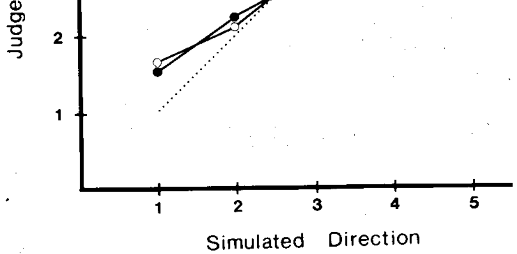

Figure 10. The mean direction of illumination ratings of 20 naive observers as a function of simulated direction. (The units along each axis are arbitrary values on a 5-point subjective rating scale defined in the instructions to be roughly equivalent to 0°, 10°, 20°, 30°, and 40° from the observer's line of sight. Ratings for shiny surfaces are indicated by open circles; ratings for dull surfaces are indicated by filled circles. The dotted line represents the ratings that would be produced by a perfectly accurate observer.)

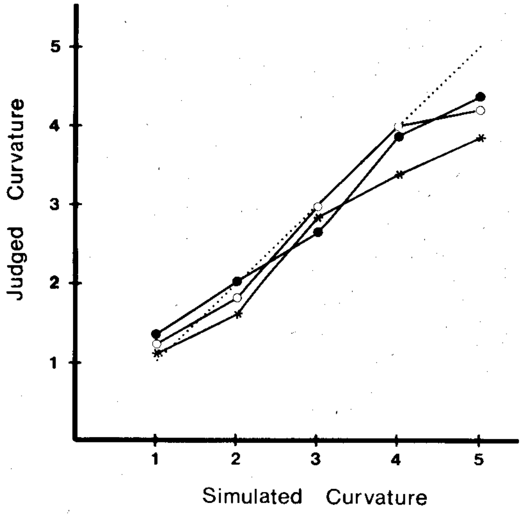

Figure 11. The mean curvature ratings of 16 naive observers as a function of simulated curvature for the textured surfaces used in Experiment 3. (The units along each axis are arbitrary values on 5-point subjective rating scale assigned to represent the five curvatures depicted in Figure 6. Ratings for shiny, dull, and unshaded surfaces are indicated by open circles, filled circles, and asterisks, respectively. The dotted line represents the ratings that would be produced by a perfectly accurate observer.)

highlight, $F(4, 76) = 9.782$, $p < .001$. These effects were quite small, however, together accounting for only 2.26% of the sum of squares.

Experiment 3

Experiment 3 examined the ability of observers to judge surface curvature from texture and shading information in combination, and from texture information in the absence of shading.

Method

Subjects. Sixteen naive observers participated in the experiment to fulfill the requirements of an introductory psychology course.

Stimulus displays. The stimuli consisted of 15 displays, each depicting a checkerboard cylindrical surface under polar projection (see the bottom portion of Figure 8). The possible curvatures of these surfaces were identical to those used in Experiments 1 and 2. Ten of these displays had shading gradients in addition to texture. Five had shiny surfaces with $s = .5$, $g = .4$, and $n = 5$ for the light squares, and $s = .3$, $g = .4$, and $n = 5$ for the dark squares. Five others had dull surfaces with $s = .5$, $g = .1$, and $n = 5$ for the light squares, and $s = .3$, $g = .1$, and $n = 5$ for the dark squares. The five remaining displays were generated with only two intensity values, one each for the light and dark texture elements. The simulated eye position was the same as in the previous experiments. The direction of illumination was assigned a value of 20° to the right of the line of sight in all displays having intensity gradients.

The addition of texture information in these displays created a slight problem with confounding variables. If the size of the texture elements were held constant across surfaces, then the number of elements would covary with curvature. Similarly, if the number of texture elements were held constant across surfaces, then their size would covary with curvature. To ensure that neither source of covariation would provide a short-cut strategy for the observers' curvature judgments, the number of texture elements per horizontal row was selected at random on every trial from possible values of 14, 16, 18, and 20. This procedure significantly reduced the covariation of curvature with texture size and density so that neither could be counted on as a reliable source of information.

Procedure. The procedure was almost identical to that of Experiment 1. The observers were told that they would see a sequence of checkerboard surfaces and that they should rate the curvature of each one on a scale of 1 to 5. They were informed of possible variations in the size and density of the texture elements and that the surfaces could appear either shiny or dull. Each observer saw three repetitions of the 15 displays. Because each display took from 5 to 7 sec to generate, the observers were told to wait until a display was fully completed before making a response. The entire experiment took from 15 to 30 min.

Results

Figure 11 shows the mean curvature ratings as a function of simulated curvature for the three types of surfaces (shiny, dull, and unshaded). An ANOVA revealed that the effect of curvature was highly significant, $F(4, 60) = 299.512$, $p < .001$, accounting for 71.46% of the total sum of squares. The differences between the three surface types were statistically significant, $F(2, 30) = 7.532$, $p < .005$, as was their interaction with curvature, $F(8, 120) = 3.190$, $p < .005$, but the relative size of these effects was extremely small accounting for less than 2% of the sum of squares.

It is important to note that the marginal influence of shading in this experiment is clearly the result of a ceiling effect. As is evident in Figure 11, the level of performance for the unshaded displays leaves little room for improvement. The experiment is still useful for understanding the perceptual analysis of shading information, however, because it provides a convenient standard for evaluating the results of Experiment 1. The regression lines for the shiny, dull, and unshaded surfaces in Experiment 3 had slopes of .81, .79, and .77, respectively—all considerably higher than the slopes of .67 and .47 that were obtained in Experiment 1 when shading information was presented in isolation. There was also a smaller within-cell variance in the present experiment. All of this suggests that the information provided by patterns of shading in these displays is less perceptually salient than other forms of information that are potentially available in a natural environment (e.g., texture, motion, and binocular vision).

Discussion

A wide variety of physical variables can affect the structure of light at a point of observation. The amount of light that reflects from a surface in any given direction depends on the orientation of the surface and its material composition, the positions and spectral compositions of all luminous bodies that can illuminate the surface directly, and the positions, orientations, and material compositions of all nonluminous bodies that can illuminate the surface indirectly through the process of reflection. The intensity of visual stimulation at each point on the retina is determined by the combined actions of all of these different variables, yet human observers are somehow able to tease apart their separate effects to obtain reliable information about the structure of the environment. Indeed, a single shaded image can often provide sufficient information to perceive the three-dimensional shape of an object, whether its surface is dull, shiny, or transparent, the direction of illumination, and whether the illumination is partially occluded by other objects.

Although a number of mathematical models have recently been developed for how patterns of shading could be analyzed in principle, there are almost no data to suggest the specific properties of shaded images that are perceptually informative to actual human observers. The present series of experiments was designed, therefore, to achieve two goals: first, to establish some baseline measures of how specific physical variables such as surface curvature, surface reflectance, and direction of illumination can influence the perceived three-dimensional structure of objects depicted in shaded images; and second, to evaluate empirically whether an observer's ability to use shading information is subject to the same limitations as the mathematical analyses of image shading that have been reported in the literature.

To appreciate the theoretical significance of these experiments, it is necessary to consider the empirical implications of some alternative mathematical models. For example, one possible method of analyzing shading information proposed by Horn (see also Ikeuchi & Horn, 1981; Woodham, Note 1; Woodham, 1981) is sufficiently powerful to determine the three-dimensional structure of an object with arbitrary precision, provided that the reflectance of its visible surface and the direction of illumination are known. An alternative method proposed by Pentland (1982a, 1982b) requires slightly less stringent assumptions, but can only determine the approximate structure of an object. The results of Experiment 1 suggest that Pentland's analysis is much more descriptive of human perception in this regard. The observers' judgments were especially inaccurate when the displays simulated the pattern of reflection for an idealized Lambertian surface. As shown in Figure 9, they consistently underestimated the simulated curvatures in those displays by over 50%. This low level of performance can only be attributed to inherent limitations in the specific strategies by which the human visual system makes use of shading information. The ability of observers to understand the rating task and to make reasonably accurate curvature judgments under appropriate conditions is clearly demonstrated by their responses to shiny surfaces in Experiment 1 and to textured surfaces in Experiment 3.

Another important aspect of environmental structure that can be depicted in a shaded image is the direction of illumination. The results of Experiment 2 suggest that the available information about how an object is illuminated is largely independent of information about surface curvature. Unlike curvature judgments, the perceived direction of illumination is unaffected by specular highlights. When presented with an image of a relatively dull (matte) surface, for example, an observer is likely to underestimate the curvature of the surface yet correctly judge the position of the light source. These findings are theoretically important primarily because of the particular surfaces on which the observers' judgments were based. Of all the mathematical analyses that have been proposed in the literature, the problem of perceived direction of illumination has only been addressed by Pentland (1982a, 1982b). Pentland has shown, however, that his analysis will not work for certain types of "degenerate" surfaces—specifically, planes and cylinders—that do not conform to its underlying assumptions. Thus, the analysis would have to be modified in order to explain how observers could accurately judge direction of illumination from the simulated images of cylindrical surfaces used in the present experiments.

Perhaps the most important theoretical issue addressed in these experiments concerns the effect of surface quality on the perceptual analysis of shading information. The mathematical models that have been suggested to date have all imposed severe constraints on the types of surfaces to which they can be applied. The analyses proposed by Horn (1975) and Pentland (1982a, 1982b), for example, are derived from an assumption that surfaces in a natural environment will generally approximate a hypothetical Lambertian surface, for which the amount of incident light reflected in any direction is proportional to the cosine of the angle of incidence.[4] The results of the present experiments suggest, however, that this type of analysis may be difficult to justify for a theory of human vision. If the perceived three-dimensional interpretation of a shaded image were based on an assumption that light reflected from a surface must always obey Lambert's law, then it would be reasonable to expect that increasing the specular component of a surface's reflectance function (see Equation 1) should decrease the accuracy of observers' judgments about surface curvature and direction of illumination. The data clearly do not support this prediction. The specular highlights had no effect whatsoever on the perceived direction of illumination, and their effect on perceived curvature was in the opposite direction. It would appear, therefore, that the analysis of shading information used by human observers is considerably more general than any of the theoretical models that have yet been proposed.

Many other aspects of shading information remain to be examined by future research. One important problem that the present research does not address concerns the manner in which different parts of a shaded image are globally organized. Consider the relative salience of three-dimensional structure for the two images presented in Figure 5 and the top portion of Figure 8. Although both of these images were generated with the same basic model of reflection, the one in Figure 5 is considerably more complex; it depicts the boundaries of objects, a textured background, cast shadows, and transparencies. Most of these properties would be treated as unwanted noise by existing mathematical analyses, but they have exactly the opposite effect on human perception. Indeed, the impression of three-dimensional form created by Figure 5 is much more striking than the relatively weak effect that is obtained with the top portion of Figure 8. A fundamental difficulty for analyzing the image of a complex scene is that the contours produced by opaque objects, transparent objects, texture, and shadows are completely indistinguishable within a single local region. To obtain reliable information from these higher order properties, it would be necessary to adopt a method of analysis that makes use of the global organization of an image in addition to its locally defined gradients. The human visual system has apparently evolved an adequate solution to this problem, but the specific principles on which it is based have yet to be discovered.

---

[4] More recent variations of Horn's analysis are not restricted to Lambertian surfaces. It is assumed instead that each possible surface orientation can only correspond to a single image intensity. The set of these correspondences, called a reflectance map, must be determined empirically for a given surface before any images of that surface can be analyzed (see also Woodham, Note 1).

Reference Notes

1. Woodham, R. J. *Reflectance map techniques for analyzing surface defects in metal castings* (Tech. Rep. AI-TR-457). Cambridge, Mass.: MIT, Artificial Intelligence Laboratory, June 1978.

2. Bergström, S. S. *Illumination, colour and 3-D form.* (Report No. 3). Umea, Sweden: Umea University, Department of Applied Psychology (DAPS), November 1980.
3. Beck, J., & Prazdny, K. *The perception of transparency with achromatic colors* (Tech. Rep. TR-1240). College Park: University of Maryland, Computer Science Center, January 1983.

References

Benson, C., & Yonas, A. Development of sensitivity to static pictorial depth information. *Perception & Psychophysics,* 1973, *13,* 361–366.

Bergström, S. S. Common and relative components of reflected light as information about the illuminations, colour and three-dimensional form of objects. *Scandinavian Journal of Psychology,* 1977, *18,* 180–186.

Blinn, J. F. Models of light reflection for computer synthesized pictures. *Computer Graphics,* 1977, *11,* 192–198.

Blinn, J. F., & Newell, M. E. Texture and reflection in computer generated images. *Communications of the ACM,* 1976, *19,* 542–547.

Cook, R. L., & Torrance, K. E. A reflectance model for computer graphics. *Computer Graphics,* 1981, *15,* 307–316.

Gilchrist, A. L. The perception of surface blacks and whites. *Scientific American,* 1979, *240,* 112–124.

Gilchrist, A. L., & Jacobsen, A. Perception of lightness and illumination in a world of one reflectance. *Perception,* in press.

Hagen, M. The development of sensitivity to cast and attached shadows in pictures as information for the direction of the source of illumination. *Perception & Psychophysics,* 1976, *20,* 25–28.

Horn, B. K. P. Obtaining shape from shading information. In P. H. Winston (Ed.), *The psychology of computer vision.* New York: McGraw-Hill, 1975.

Horn, B. K. P. Understanding image intensities. *Artificial Intelligence,* 1977, *8,* 201–231.

Horn, B. K. P. Hill-shading and the reflectance map. *Proceedings of the IEEE,* 1981, *19,* 14–47.

Ikeuchi, K., & Horn, B. K. P. Numerical shape from shading and occluding contours. *Artificial Intelligence,* 1981, *17,* 141–184.

Kay, D. S., & Greenberg, D. Transparency for computer synthesized images. *Computer Graphics,* 1979, *13,* 158–165.

MacLeod, R. B. An experimental investigation of brightness constancy. *Archives of Psychology,* 1932, no. 135.

Metelli, F. The perception of transparency. *Scientific American,* 1974, *230,* 90–98.

Pentland, A. P. Finding the direction of illumination. *Journal of the Optical Society of America,* 1982, *72,* 448–455. (a)

Pentland, A. P. *The visual inference of shape: Computation from local features.* Unpublished doctoral dissertation, Massachusetts Institute of Technology, 1982. (b)

Pentland, A. P. *The visual inference of shape: Computation from local features.* Unpublished doctoral dissertation, Massachusetts Institute of Technology, 1982.

Phong, B. T. Illumination for computer-generated pictures. *Communications of the ACM,* 1975, *18,* 311–317.

Richter, J. P. *The notebooks of Leonardo da Vinci.* New York: Dover, 1970.

Whitted, T. An improved illumination model for shaded display. *Communications of the ACM,* 1980, *23,* 343–349.

Woodham, R. J. Analyzing images of curved surfaces. *Artificial Intelligence,* 1981, *17,* 117–140.

Yonas, A., Kuskowski, M., & Sternfels, S. The role of frames of reference in the development of responsiveness to shading information. *Child Development,* 1979, *50,* 493–500.

Received October 7, 1982
Revision received January 26, 1983 ■