Photometric Method for Determining Shape from Shading*

by

ROBERT J. WOODHAM**

Technical Report 84-10
July 1984

Laboratory for Computational Vision
Department of Computer Science
University of British Columbia
Vancouver, B.C., CANADA

Abstract

A smooth opaque object produces an image in which brightness varies spatially even if the object is illuminated evenly and is covered by a surface material with uniform optical properties. Photometric methods relate image irradiance to object shape and surface material using physical models of the way surfaces reflect light. A reflectance map allows image irradiance to be written as a function of surface orientation, for a given surface material and light source distribution. Shape from shading algorithms use a reflectance map to analyze what is seen.

The development of photometric methods for determining shape from shading is discussed, beginning with examples from lunar astronomy. The results presented delineate shape information that can be determined from geometric measurements at object boundaries from shape information that can be determined from intensity measurements over sections of smooth surface. Recent work of Ikeuchi and Horn is presented which relaxes the requirement that the image irradiance equation be satisfied exactly. Instead, the image irradiance equation specifies one constraint that is combined with another constraint derived from general surface smoothness criteria. Shape from shading is expressed as a constrained minimization problem.

Another method uses multiple images in a technique called photometric stereo. In photometric stereo, the illumination is varied between successive images while the viewing direction remains constant. Multiple images obtained in this way provide enough information to determine surface orientation at each image point, without smoothness assumptions.

* This is a pre-publication version of a chapter to appear in *Image Understanding 1984*, S. Ullman & W.A. Richards (eds.), Ablex Publishing Co., Norwood, NJ.

** R.J. Woodham is a Fellow of the Canadian Institute for Advanced Research.

1. Introduction

A smooth opaque object will produce an image in which brightness varies spatially even if the object is illuminated evenly and is covered by a surface material with uniform optical properties. Computational vision often does not exploit shading information, using intensity measurements only to segment an image at locations where brightness changes abruptly. The purpose of this chapter is to show that shading provides essential information about object shape and that algorithms have been developed to determine shape from shading.

The key observation, derived from principles of physical optics, is that the apparent brightness of a surface element depends on the orientation of that element relative to the viewer and light sources. Different surface elements of a nonplanar object reflect different amounts of light toward an observer because of their differing attitude in space. In computer graphics, shaded images are generated by choosing a suitable reflectance model and illumination to assign a brightness value at each visible surface element. For a given material, viewer position and distribution of light sources, image brightness is determined as a function of surface orientation. Figure 1 is a synthesized image of a sphere. In this example, the object is assumed to be a perfectly diffuse reflector illuminated by a single distant light source oriented slightly above and to the right of the viewer. Sections 2.2 and 2.5 discuss this reflectance model in more detail. Realistic images have been produced for a wide variety of surface materials and illumination conditions. Cook and Torrance (1982) present examples of current work in the computer generation of shaded images.

In computational vision, the goal is to determine surface properties from image brightness. Shape from shading assigns a surface orientation to each visible surface element. The problem is difficult because surface orientation has two degrees of freedom and image irradiance, the quantity measured in an image, has only one. Image irradiance results from the interaction of several factors, some of which are properties of the objects in view and some of which are not.

The effects of shape and surface material must be separated from each other and from the effects of illumination, shadows, viewing direction and path phenomena. In general, there are trade-offs that cannot be resolved.

Figure 2 illustrates one trade-off. Figure 2(a) seems to show a container holding five eggs. But there are no eggs in the container. We are fooled at first because we are unaccustomed to interpreting scenes illuminated from below. Scrutiny is required to see the shapes as hollows. In Figure 2(b), obtained under more natural lighting conditions, two eggs have been added and the three remaining hollows are easily identified. In the absence of other cues, the perception of convexity and concavity is reversed when a corresponding change occurs in the direction of illumination.

Figure 3 illustrates another trade-off. Often, it is impossible to disambiguate smooth changes in surface material from smooth changes in surface shape. If one assumes that surface material is constant then the perception of shape will be altered. This is exploited in the application of cosmetic makeup. Figure 3(a) shows a model's face without makeup. Figure 3(b) demonstrates that makeup can alter the perception of shape. The face, with makeup, appears narrower and the cheek bones have been accentuated.

To deal with these trade-offs, it is necessary to understand how images are formed. One aspect is the geometry of image projection. Of equal importance is the radiometry of image formation. Relating image irradiance to object shape and surface material requires a model of the way surfaces reflect light. Fortunately, the needed analytic tools now exist. Section 2 develops an image irradiance equation to model how the physical world determines what we see. The reflectance map, originated by Horn (1977), allows image irradiance to be written as a function of surface orientation, for a given surface material and light source distribution.

In the photometric method for determining shape from shading one uses the image irradiance equation to analyze what is seen. A reflectance map is not directly invertible since, as noted, surface orientation has two degrees of freedom and image intensity provides only one measurement. Additional information is required to reconstruct the underlying surface. Sometimes, special reflectance properties of the surface material simplify the analysis. The moon, for example, has remarkable photometric properties arising from the special structure of the lunar surface. Figure 4 is an image obtained at full moon. That is, it is obtained when the light source is directly behind the viewer. At full moon, the distribution of brightness over the lunar disk is about the same everywhere. Brightness varies spatially as a function of surface material but there is no general shading owing to the spherical shape of the moon.

Figure 5 is an image obtained using a scanning electron microscope (SEM). Although different from a natural scene, shading in an SEM image provides a clear perception of object shape. In SEM images, there is a simple analytic relationship between image irradiance and surface orientation. Therefore, shape from shading methods can be applied to SEM images. The method of Ikeuchi and Horn (1981), described in Section 3.3, has been demonstrated on a sample SEM image.

Following the development of the image irradiance equation in Section 2, Section 3.1 traces the historical development of photometric methods for determining shape from shading, beginning with examples from lunar astronomy. Methods for determining object shape from a single view embody assumptions about surface curvature. For several classes of surface, surface orientation can be determined locally. This is demonstrated, in Section 3.2, for planar surfaces forming trihedral corners. These results help to delineate shape information that can be determined from geometric measurements at object boundaries and shape information that can be determined from intensity measurements over sections of smooth surface. Recent work relaxes the

requirement that the image irradiance equation be satisfied exactly. Instead, the image irradiance equation specifies one constraint that is combined with another constraint derived from general surface smoothness criteria. Finally, in Section 3.3, shape from shading is expressed as a constrained minimization problem.

Another photometric method uses multiple images. Section 4 discusses the technique called photometric stereo. Binocular stereo determines range by relating two images of an object viewed from different directions. If the correspondence between picture elements is known, then distance to the object can be calculated by triangulation. Unfortunately, it is difficult to determine this correspondence. The idea of photometric stereo is to vary the illumination between successive images, while holding the viewing direction constant. Multiple images obtained in this way provide enough information to determine surface orientation at each image point, without smoothness assumptions. Since the imaging geometry is not changed, the correspondence between image points is fixed. The technique is photometric because it uses the image irradiance values recorded at a single image location, in successive views, rather than the relative position of displaced features.

2. Developing the Tools

Image formation is modeled by an image irradiance equation that determines image brightness as a function of surface orientation. The development presented here was first given by Horn (1977) and incorporates extensions from Woodham (1981) and from Ikeuchi and Horn (1981). Section 2.1 describes the orthographic image projection. Section 2.2 defines a viewer-centered surface reflectance function. Section 2.3 presents three representations for surface orientation: gradient space, spherical coordinates and stereographic coordinates. Section 2.4 defines the Hessian matrix representation for surface curvature. Section 2.5 introduces the

reflectance map. Finally, Section 2.6 uses the reflectance map to define the image irradiance equation.

2.1 Image Projection

In general, optical systems perform a perspective projection. If the size of the objects in view is small compared to the viewing distance, then a perspective projection can be approximated by an orthographic projection as illustrated in Figure 6. To standardize the geometry, it is convenient to choose a left-handed Cartesian coordinate system with the viewing direction aligned with the negative z-axis. Without loss of generality, the image plane can be scaled so that surface point (x, y, z) projects to image point (u, v) with $u = x$ and $v = y$. With these assumptions, image coordinates (x, y) and surface coordinates (x, y) can be used interchangeably. Image projection simply discards the z coordinate of a visible surface point (x, y, z).

2.2 Surface Reflectance

The amount of light reflected by a surface element in a given direction depends on its optical properties, on its microstructure and on the distribution of scene irradiance. For most surfaces, the fraction of illumination reflected in a particular direction depends only on the surface orientation. The reflectance characteristics of such a surface can be represented as a function $\phi(i, e, g)$ of the three angles i, e and g defined in Figure 7. These angles are called the *incident, emergent* and *phase* angles. The reflectance function $\phi(i, e, g)$ determines the ratio of surface radiance to irradiance measured per unit surface area, per unit solid angle, in the direction of the viewer. That is, $\phi(i, e, g)$ determines how bright a surface element appears when illuminated in a particular way and viewed from a particular direction. Surface orientation, as given by a surface normal vector, determines the angles i and e. (With a distant point light

source and an orthographic projection, the phase angle g is constant at all surface points.) Although the angles i and e are defined in an object-centered coordinate system, the function $\phi(i,e,g)$ itself is defined with respect to the viewer. As such, it is a viewer-centered definition of surface reflectance.

The reflectance function $\phi(i,e,g)$ is assumed to be symmetric about the surface normal. There are some materials that are not isotropic in this way. Rotation of the surface about the surface normal can be included by making reflectance a function of four angles instead of three. The general case is discussed in Nicodemus et al. (1977) and leads to what has been called the *bidirectional reflectance distribution function* (BRDF).

Surface reflectance characteristics can be measured empirically, derived from analytic models of surface microstructure or developed from phenomenological models of surface behavior. To illustrate, consider the model of a perfect specular or mirrorlike surface illuminated by a single distant point source. In specular reflection, the incident angle equals the emergent angle ($i = e$) and the incident, emergent and normal vectors lie in the same plane ($i+e = g$). Expressing this formally, one obtains the function $\phi_1(i,e,g)$ given by

$$\phi_1(i,e,g) = \begin{cases} 1 & \text{if } i = e \text{ and } i+e = g \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

Equation 1 implies that there is only one surface orientation correct for reflecting the light source towards the viewer. This is familiar to anyone who has used a pocket mirror to direct rays from the sun towards a desired target.

A second model is given by

$$\phi_2(i,e,g) = \begin{cases} \rho \cos(i) & \text{if } i < \frac{\pi}{2} \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

This reflectance function describes a diffuse surface illuminated by a single distant point source.

The cosine of the incident angle i accounts for the variation in irradiance due to the inclination of the surface element from the light source. The reflectance factor ρ is a constant between 0 and 1 and accounts for the general brightness of the surface. When $\rho = 1$, all the incident light is reflected and the surface is called perfectly diffuse. Since $\phi_2(i,e,g)$ does not depend on the emergent angle e, a diffuse surface appears equally bright from all viewing directions. As we see in the third model below, this is not equivalent to a surface that reflects an equal amount of light in all directions.

A third model, similar to that of material in the maria of the moon, is given by

$$\phi_3(i,e,g) = \begin{cases} \frac{\rho \cos(i)}{\cos(e)} & \text{if } i < \frac{\pi}{2} \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

This reflectance function describes a surface, illuminated by a single distant point source, that reflects equal amounts of light in all directions. The terms ρ and $\cos(i)$ are as described for equation (2). Recall that $\phi(i,e,g)$ is defined with respect to the viewer. A surface appears foreshortened as a function of emergent angle e. That is, as the surface becomes more inclined from the viewer a larger surface area is captured per unit solid angle subtended at the viewer. If reflection is equal in all directions then the surface appears brighter as it becomes more inclined from the viewer. The cosine of the emergent angle e accounts for surface foreshortening as seen by the viewer.

A fourth model, similar to that of scanning electron microscope (SEM) images, is given by

$$\phi_4(i,e,g) = \frac{1}{\cos(e)} \tag{4}$$

This reflectance function describes a surface for which the only consideration is surface foreshortening as seen by the imaging system.

2.3 Surface Orientation

There are many ways to specify surface orientation. For a point on a curved surface, one can use the tangent plane or, equivalently, the surface normal vector. If the equation of a surface is given explicitly as $z = f(x,y)$ then a vector in the direction of the surface normal is given by

$$\left[\frac{\partial f(x,y)}{\partial x}, \frac{\partial f(x,y)}{\partial y}, \quad -1\right]$$

If parameters p and q are defined by

$$p = \frac{\partial f(x,y)}{\partial x} \quad \text{and} \quad q = \frac{\partial f(x,y)}{\partial y} \tag{5}$$

then the vector can be written as $[p,q,-1]$. The quantity (p,q) is called the *gradient* of $f(x,y)$ and *gradient space* is the two-dimensional space of all such points (p,q). Gradient space is one representation for surface orientation.

Another way to represent surface orientation is to use points on the unit sphere, called the *Gaussian sphere* after Hilbert and Cohn-Vossen (1952). Each point on the Gaussian sphere identifies the unit vector formed by joining the center of the sphere to that point. Spherical coordinates (ϕ,θ) can then be used to specify a direction. Figure 8 illustrates.

There is, of course, a mapping between points on the Gaussian sphere and gradient space. If (ϕ,θ) is a point on the Gaussian sphere then the gradient (p,q) is given by

$$p = \cos(\theta)\tan(\phi) \quad \text{and} \quad q = \sin(\theta)\tan(\phi) \tag{6}$$

This mapping has the simple geometric interpretation illustrated in Figure 9. Let the viewer be located far above the north pole of the Gaussian sphere looking southward along the positive z-axis. If a plane is constructed tangent to the sphere at the north pole then gradient space is the projection of the Gaussian sphere formed by extending each unit vector outwards until it intersects this plane ($z = -1$). Gradient space represents the northern hemisphere of the Gaus-

sian sphere. That is, gradient space represents all surface orientations that have a positive component in the direction of the viewer. These surface orientations correspond to visible surface elements.

A third way to represent surface orientation uses another projection of the Gaussian sphere onto the plane tangent at the north pole. If the center of projection is the south pole, not the center of the sphere, then the *stereographic projection* is obtained. The stereographic projection is illustrated in Figure 10. The coordinates of the stereographic plane are called f and g to avoid confusion with the p and q coordinates of gradient space. Stereographic coordinates (f, g) can be determined from the corresponding gradient (p, q) as follows

$$f = \frac{2p(\sqrt{1+p^2+q^2}-1)}{p^2+q^2} \text{ and } g = \frac{2q(\sqrt{1+p^2+q^2}-1)}{p^2+q^2} \tag{7}$$

A representation for surface orientation is chosen for the particular task at hand. Spherical coordinates, gradient space and stereographic coordinates are all used in the specification of surface reflectance. The bidirectional reflectance distribution function (BRDF), mentioned in Section 2.2, uses spherical coordinates to represent directions. Often, one is faced with multiple or extended light sources. The resulting dependence of image brightness on surface orientation can be derived by integrating the particular BRDF over the full range of light source directions. Spherical coordinates simplify the formulation and evaluation of the required integrals. Horn and Sjoberg (1979) show how to derive the reflectance function $\phi(i, e, g)$ from the BRDF for both single and extended light source configurations.

Gradient space has been used in work on the scene analysis of linedrawings (Huffman (1971), Mackworth (1973), Draper (1981)). Many geometric properties have a simple formulation in gradient space. In shape from shading, gradient space is used because it corresponds directly to the first partial derivatives of surface height z, allowing height to be computed by integrating p and q.

Gradient space does have one serious drawback. Points where a surface smoothly disappears from view ($e = \pi/2$) form what is called an *occluding contour* after Marr (1977). Points on an occluding contour map to infinity in gradient space, as can be seen in Figure 9. Consequently, constraints at occluding contours cannot be expressed using the gradient (p, q). Stereographic coordinates (f, g) are used instead. In the stereographic projection, points on an occluding contour lie on the circle $f^2 + g^2 = 2^2$.

2.4 Surface Curvature

Three parameters are required to specify the curvature at a point on a surface. Again, there are several possible choices. One familiar representation uses the two orthogonal principal curvatures and associated azimuthal direction about the surface normal. Another representation is based on the second partial derivatives, with respect to x and y, of the surface $z = f(x, y)$. Let $\mathbf{H}$ be the 2x2 matrix given by

$$\mathbf{H} = \begin{bmatrix} \frac{\partial^2 f(x,y)}{\partial x^2} & \frac{\partial^2 f(x,y)}{\partial x \partial y} \\ \frac{\partial^2 f(x,y)}{\partial y \partial x} & \frac{\partial^2 f(x,y)}{\partial y^2} \end{bmatrix} \tag{8}$$

$\mathbf{H}$ is called the Hessian matrix of the function $z = f(x, y)$. The Hessian matrix is used in nonlinear programming and is the representation of surface curvature used in Woodham (1981). $\mathbf{H}$ is indeed a three parameter function of x and y since, for surfaces of class C^2, the order of differentiation can be interchanged and $\mathbf{H}$ is symmetric.[1] If the Hessian matrix $\mathbf{H}$ and the corresponding gradient (p, q) are known at a point, then the corresponding principal curvatures and associated direction are also determined (Woodham 1978 pg. 68).

[1] A function is of class C^k if it has continuous partial derivatives of order k.

The Hessian matrix **H** is used because **H** directly determines the change in gradient $[dp, dq]$ corresponding to a small movement $[dx, dy]$ in the image. That is

$$\begin{bmatrix} dp \\ dq \end{bmatrix} = \mathbf{H} \begin{bmatrix} dx \\ dy \end{bmatrix} \tag{9}$$

Conversely, if two linearly independent $[dx, dy]$'s and the corresponding $[dp, dq]$'s are known at a given point in the image, then **H** is determined by

$$\mathbf{H} = \begin{bmatrix} dp_1 & dp_2 \\ dq_1 & dq_2 \end{bmatrix} \begin{bmatrix} dx_1 & dx_2 \\ dy_1 & dy_2 \end{bmatrix}^{-1} \tag{10}$$

2.5 Reflectance Maps

The surface normal relates surface geometry to image irradiance because it determines the angles i and e appearing in the surface reflectance function $\phi(i, e, g)$. An ideal imaging device produces image irradiance proportional to scene radiance. Thus, for a fixed light source distribution and viewer geometry, the ratio of scene radiance to scene irradiance depends only on the surface normal. Provided each surface element receives the same irradiance, scene radiance, and hence image brightness, depends only on surface orientation.

A *reflectance map* determines image irradiance as a function of surface orientation. Typically, a reflectance map is expressed as a function $R(p, q)$ of the gradient (p, q). One way to interpret a reflectance map $R(p, q)$ is as a transformation of a surface reflectance function $\phi(i, e, g)$ into a function of p and q. Expressions for $cos(i)$, $cos(e)$, and $cos(g)$ can be derived using normalized dot products of a surface normal vector $[p, q, -1]$, the vector $[p_s, q_s, -1]$ that points in the direction of the light source and the vector $[0,0,-1]$ that points in the direction of the viewer. The reflectance map corresponding to $\phi_2(i, e, g)$, defined in equation (2), is given by

$$R_2(p,q) = \frac{\rho(1 + pp_s + qq_s)}{\sqrt{1 + p^2 + q^2}\ \sqrt{1 + p_s^2 + q_s^2}} \tag{11}$$

Similarly, the reflectance map corresponding to $\phi_3(i,e,g)$, defined in equation (3), is given by

$$R_3(p,q) = \frac{\rho(1 + pp_s + qq_s)}{\sqrt{1 + p_s^2 + q_s^2}} \tag{12}$$

(Note that the reflectance map $R_3(p,q)$ is linear in p and q.) It is convenient to represent $R(p,q)$ as a series of contours of constant brightness in gradient space. Figures 11 and 12 illustrate the reflectance maps $R_2(p,q)$ and $R_3(p,q)$, defined above, with $\rho = 1.0$, $p_s = 0.7$ and $q_s = 0.3$.

Figure 1 was generated using the reflectance map $R_2(p,q)$ with $\rho = 1.0$, $p_s = 0.7$ and $q_s = 0.3$ as in Figure 11. For each image point (x,y), the gradient (p,q) at surface point $z = f(x,y)$ was determined analytically. A brightness value was then assigned using the function $R_2(p,q)$. (The resulting value is normalized to match the dynamic range of the image display device.)

A reflectance function $\phi(i,e,g)$ can also be expressed as a function $R(f,g)$ of stereographic coordinates (f,g). Again, using the example of $\phi_2(i,e,g)$ defined in equation (2), one obtains

$$R_2(f,g) = \frac{16\rho(ff_s + gg_s) + (4 - f^2 - g^2)(4 - f_s^2 - g_s^2)}{(4 + f^2 + g^2)(4 + f_s^2 + g_s^2)} \tag{13}$$

A reflectance map provides a uniform representation for specifying the reflectance properties of a surface material, for a particular light source distribution and viewing geometry. Reflectance maps can be derived analytically from a known $\phi(i,e,g)$ and a given distribution of source radiance. A reflectance map can also be measured empirically. A calibration object of known shape determines the reflectance map for a particular imaging situation. A reflectance map measured in this way can be used to analyze other objects made of the same material and

viewed under the same conditions of illumination. Determining the reflectance map empirically has the added benefit of automatically correcting for the transfer characteristics of the imaging device. A sphere serves as a useful calibration object. All visible surface element orientations are present in the image. In addition, the orientation at each surface element is easily determined from the object boundary.

2.6 The Image Irradiance Equation

Image formation can now be described by a single equation called the *image irradiance equation*. If gradient space is used to represent surface orientation then the image irradiance equation becomes

$$I(x,y) = R(p,q) \tag{14}$$

where $I(x,y)$ is the brightness at image point (x,y) and $R(p,q)$ is the reflectance map value at the corresponding surface gradient (p,q). Photometric methods for shape from shading reconstruct a surface $z = f(x,y)$ to satisfy equation (14). $I(x,y)$ is the measured image. The reflectance map is known as a function of surface orientation. The problem is to determine the particular gradient (p,q) at each point (x,y). Recall that p and q, defined in equation (5), are the first partial derivatives of $f(x,y)$ with respect to x and y. Thus, equation (14) is a first order partial differential equation. Photometric methods that use gradient space to represent surface orientation have used techniques developed to analyze such differential equations.

Alternatively, one can view image irradiance equations of the form $I(x,y) = R(p,q)$, $I(x,y) = R(f,g)$ or $I(x,y) = R(\phi,\theta)$ as single equations in the two unknowns chosen to specify surface orientation. Measured image brightness determines a relation the unknowns must satisfy. But, to determine surface orientation precisely, additional constraint must be provided. Properties of surface curvature, surface material and surface smoothness are one source

of additional constraint. Multiple images are the other.

3. Development of the Photometric Approach

3.1 Historical Background

The study of the photometric properties of a surface is a standard method of investigation in astronomy, especially when applied to the lunar surface. At full moon when the source is directly behind the viewer ($i = e$), the distribution of brightness over the lunar disk is about the same everywhere. Galileo considered this particular property of the moon three centuries ago in his "Dialogues on Two Systems of the World".

Since then, many investigators have carefully measured the surface reflectance function $\phi(i,e,g)$ of the moon. The goal was to predict physical properties of the lunar soil. Investigators experimented with hundreds of terrestrial materials and surface microstructures in an attempt to replicate the photometric properties of the moon. These experiments led to correct predictions of the composition, particle size and porosity of the lunar soil. The conclusion was that the lunar surface is an assembly of closely packed, randomly pointed, deep tunnels of all sizes, superimposed and juxtaposed. The photometric properties of the moon are determined primarily by the shadows cast by surface detail and not by the reflection of light from surface material. Indeed, the lunar surface is uncommon in that the same material and structure would collapse under its own weight if it were replicated in the gravitational field of the earth. Some earth vegetation, however, can have similar photometric properties. One natural material that provides a good approximation is a thick layer of the lichen *Cladonia rangiferina*. This lichen has a porous macrostructure and an opaque, somewhat rough microstructure. Minnaert (1961) and Hapke (1971) provide excellent reviews of the study of the photometric properties of the

moon.

Properties other than surface microstructure were of interest. The idea that photometric measurements can be used to derive quantitative information about surface shape seems to have originated with van Diggelen (1951). His work did not make use of the surface reflectance function $\phi(i,e,g)$ but exploited properties of surface points near the terminator (the boundary between the light and dark hemispheres). Only one component of surface orientation could be determined along the terminator. Nevertheless, the reconstructed one-dimensional elevation profiles for hills in the maria delineated height differences as small as 10 to 20 m. Interestingly, van Diggelen noted that a second image corresponding to a different position of the sun would enable the determination of both components of surface orientation. This idea could not be pursued since regions were not near the terminator in both images.

A photometric method for lunar topography was developed by Rindfleisch (1966). He formulated an image irradiance equation as a first order partial differential equation and presented a solution method that applied if the surface reflectance function was constant for constant $cos(i)/cos(e)$, as happens for the moon. With this restriction the reflectance map is similar to $R_3(p,q)$, defined in equation (12), so that the image irradiance equation (14) is linear in p and q. The method was applied to images returned by the three Ranger impacting spacecraft. Rindfleisch also noted the possibility of using a second image in his photometric method. The idea was not pursued since only single views were available from the Ranger spacecraft. In any event, multiple views obtained by a moving spacecraft would first have to be registered. This begs the question since precise registration requires the determination of stereoscopic parallax. If stereoscopic parallax were known then elevation could be determined by traditional photogrammetric techniques.

Horn (1975) generalized the photometric method for determining shape from shading by removing the requirement for linearity in p and q. The image irradiance equation was formulated as a nonlinear first order partial differential equation. A modified method of characteristic strip expansion was used to transform this equation into a set of five ordinary differential equations that were solved numerically. The method was successfully applied to several sample images including one of a human face.

Horn's method can be interpreted using the reflectance map, the image irradiance equation and the Hessian matrix **H**. By taking partial derivatives of the image irradiance equation (14), first with respect to x and then with respect to y, two equations are obtained that can be written as the single matrix equation

$$\begin{bmatrix} I_x \\ I_y \end{bmatrix} = \mathbf{H} \begin{bmatrix} R_p \\ R_q \end{bmatrix} \tag{15}$$

I_x and I_y denote partial derivatives of $I(x,y)$ with respect to x and y. Similarly, R_p and R_q denote partial derivatives of $R(p,q)$ with respect to p and q. Note that $[R_p, R_q]$ defines a normal to the contour of constant brightness in the reflectance map at the current (p,q) and that $[I_x, I_y]$ defines a normal to the contour of constant brightness in the image at the current (x,y). If a small movement $[dx, dy]$ is made in the image in the direction $[R_p, R_q]$ then, by equations (9) and (15), the corresponding movement $[dp, dq]$ in gradient space is in the direction $[I_x, I_y]$. More precisely,

$$\text{if } [dx, dy] = [R_p, R_q]ds \text{ then } [dp, dq] = [I_x, I_y]ds \tag{16}$$

where ds is a differential element of path length. Thus, the image irradiance equation (14) determines one $[dx, dy]$ for which the corresponding $[dp, dq]$ can be calculated. From equation (10), we know that two linearly independent $[dx, dy]$'s and the corresponding $[dp, dq]$'s are required to determine **H** completely. It is not possible to choose an arbitrary direction for $[dx, dy]$ and determine the corresponding $[dp, dq]$.

Horn's method confined the choice of $[dx, dy]$ to the direction determined by equation (16). Suppose image point (x_0, y_0) is known to have gradient (p_0, q_0). $[R_p, R_q]$ and $[I_x, I_y]$ are calculated at this point. A small step size ds is selected and equation (16) is applied to move to the point $(x_1, y_1) = (x_0 + dx, y_0 + dy)$ with gradient $(p_1, q_1) = (p_0 + dp, q_0 + dq)$. This procedure is repeated to trace a path through the image along which the gradient is determined. This path is called a *characteristic strip*.

Several problems remained. Points where the gradient is known are required as initial conditions for characteristic strip expansion. But, initial conditions are difficult to establish. In particular, constraints at occluding contours cannot be used since they cannot be represented by the gradient (p, q). Full surface reconstruction requires interpolation of surface points between neighboring characteristic strips. Brightness measurements are also influenced by noise and uncertainty. Errors accumulate as the computation iterates so that characteristic strips begin to deviate more and more from their correct path. Attempts to control error propagation and to interpolate intermediate surface points embodied assumptions about surface smoothness.

When assumptions about surface smoothness are suitably formulated, relaxation or cooperative computation methods can be employed. Local constraints propagate to determine a global solution. Constraints flow in all directions so that errors don't accumulate in particular directions. This leads to better numerical stability in the presence of noise.

Several results followed Horn in an attempt to capture these ideas. Woodham (1977) used local assumptions about convexity and concavity to determine monotonicity relations between gradients at selected image points. These relations were embodied in a network consistency algorithm to restrict the possible surface orientations at each image point. (See Mackworth (1977) for a discussion of network consistency algorithms.)

Brooks (1979) noted that if the solution surface is of class C^2 then the integral of (p,q) along any closed curve must be zero. That is, if height z is determined along a curve then one must end with no net gain or loss in height when the curve returns to its starting point. Brook's iterative algorithm eliminated surface orientations that were locally possible, according to the image irradiance equation, but violated the loop integral criterion.

Fundamental questions about conditions that $R(p,q)$ and $I(x,y)$ must satisfy to guarantee the existence of a unique solution $z = f(x,y)$ remained open. Recently, Bruss (1981) has provided some formal answers. Her results are technical in nature and are not reproduced in detail here. Suffice to say that there are few imaging situations in which every different surface produces a different image. Without initial conditions, the solution is not unique. Results have been established to integrate the photometric method with constraints determined from geometric measurements at surface discontinuities. Section 3.2 presents results for the domain of planar surfaces with trihedral corners. Ikeuchi and Horn (1981) reformulated the shape from shading problem to include initial conditions obtained at occluding contours. Their approach is discussed in Section 3.3.

3.2 Planar Surfaces with Trihedral Corners

Objects whose surfaces are planar form the familiar polyhedral domain. Horn (1977) demonstrated that when three visible planes meet at a point the orientation of each plane can be determined locally. The technique combines the quantitative approach to interpreting linedrawings of polyhedra, developed by Mackworth (1973), with the quantitative constraint provided by image irradiance.

Suppose three planes A, B and C meet to form a vertex as illustrated in Figure 13(a). Mackworth (1973) showed that the line in gradient space joining the gradients of two planes is

perpendicular to the line in the image formed by the intersection of the two planes. A possible triangle in gradient space formed by the corresponding gradients GA, GB and GC is shown in Figure 13(b). The edges formed by the planes A, B and C provide three constraints. The position and scale of the triangle shown in Figure 13(b) remain to be determined. Three other constraints are provided by measurements of image intensity for the three planes. If the intensities measured on the three planes A, B and C are α, β and γ, then GA, GB and GC must lie on gradient space contours $R(p,q) = \alpha$, $R(p,q) = \beta$ and $R(p,q) = \gamma$. The six constraints, taken together, are generally enough to determine the orientation in space of the three planes A, B and C.

The idea is illustrated graphically in Figure 14. Two triangles that satisfy the geometric constraints have been superimposed onto the reflectance map of Figure 11. The correct triangle is the one whose vertices also satisfy the brightness constraints. The position and scale of the triangle can be adjusted until GA, GB and GC lie on the required iso-brightness contours. More than one triangle can satisfy all six constraints. Often, however, the solution is unique.

In an ideal polyhedral domain, surface orientation is discontinuous at each edge. This is not true in practice, however, since even carefully prepared polyhedra tend to have edge imperfections that result in sharp but nevertheless continuous changes in surface orientation from one plane to the other. Figure 15 shows edge imperfections. The change in surface orientation from plane B to plane C of Figure 14, for example, can be continuous, taking on all orientations lying on the line joining GB and GC in gradient space. Often, intensity measurements across the edge joining B and C will have values outside the range spanned by β and γ. Indeed, if the larger triangle of Figure 14 is correct then there would be a highlight on the edge joining B and C because the line joining GB and GC passes through brightness values larger than both β and γ.

Figure 2 is an example of edge effects. The initial ambiguity arises since, for most reflectance functions, the image produced by a surface $z = f(x,y)$ illuminated from direction (p_s,q_s) is identical to that produced by the surface $z = -f(x,y)$ illuminated from direction $(-p_s,-q_s)$. Intensity measurements across edges help to resolve the ambiguity. Horn (1977) suggested that an analysis of intensity profiles across edges can determine whether the edge is convex, concave or occluding. In Figure 2(a), a highlight can be seen around the lower edges of the egg hollows. Such a highlight is inconsistent with the interpretation that the edge is occluding, as required for an egg to be present. Instead, it suggests that the edge is convex, which is the correct interpretation.

3.3 Occluding Contour

An object's silhouette provides additional information that can be used to establish initial conditions. Parts of a silhouette may correspond to sharp edges on the surface, as in the polyhedral domain. Other parts of a silhouette correspond to places where the surface curves around smoothly. The latter is referred to as an occluding contour, as we have seen. At an occluding contour, surface orientation can be computed directly from the silhouette (Marr (1977), Barrow and Tenenbaum (1981), Ikeuchi and Horn (1981)). In an orthographic projection, a normal to the silhouette in the image plane is also a normal to the surface at the corresponding point on the occluding contour.

Ikeuchi and Horn (1981) developed an iterative algorithm to determine surface orientation using the image irradiance equation and a smoothness criterion as constraints. Surface orientation at occluding contours is the main source of initial conditions although information from singular points, specular points and self-shadow boundaries is also included. The stereographic projection is used to represent surface orientation.

Consider first the continuous case. The goal is to find functions $f(x,y)$ and $g(x,y)$ that make error in the image irradiance equation small, while keeping the solution surface as smooth as possible. Error in the image irradiance equation is given by

$$\int\int (I(x,y) - R(f,g))^2 \, dxdy \tag{17}$$

Departure from smoothness is given by

$$\int\int ((f_x^2 + f_y^2) + (g_x^2 + g_y^2)) \, dxdy \tag{18}$$

where f_x, f_y, g_x and g_y are the first partial derivative of f and g with respect to x and y. Ikeuchi and Horn minimized the error term e defined by

$$e = \int\int [(f_x^2 + f_y^2) + (g_x^2 + g_y^2) + \lambda(I(x,y) - R(f,g))^2] \, dxdy \tag{19}$$

Error in the image irradiance equation is weighted by λ compared to departure from smoothness. If the reflectance map is known accurately and if the brightness measurements are precise then λ can be made large. On the other hand, if λ is small then a smooth surface is determined despite noise and uncertainties about reflectance and illumination.

In the discrete case, the local error in the image irradiance is given by

$$r_{i,j} = (I_{i,j} - R(f_{i,j},g_{i,j}))^2 \tag{20}$$

where $I_{i,j}$ is the measured brightness at image point (i,j) and $(f_{i,j},g_{i,j})$ is the corresponding surface orientation in stereographic coordinates. The local departure from smoothness is given by

$$s_{i,j} = \frac{(f_{i+1,j} - f_{i,j})^2 + (f_{i,j+1} - f_{i,j})^2 + (g_{i+1,j} - g_{i,j})^2 + (g_{i,j+1} - g_{i,j})^2}{4} \tag{21}$$

The problem is to minimize the sum of the local error terms over all points (i,j). That is, we wish to minimize the error term e given by

$$e = \sum_i \sum_j (s_{i,j} + \lambda \, r_{i,j}) \tag{22}$$

where λ is as in the continuous case.

The solution method requires that equation (22) be differentiated with respect to $f_{i,j}$ and $g_{i,j}$. For a minimum, these partial derivatives are all set to zero resulting in a large, sparse set of equations. These equations are solved using an iterative method and the set of values for $f_{i,j}$ and $g_{i,j}$ determine the solution surface.

Ikeuchi and Horn applied their algorithm to several test surfaces. The algorithm works well when all information is precise. It continues to work reasonably well even when the reflectance map is only a crude approximation.

4. Using Multiple Images

Another approach uses multiple images to provide additional constraint. These images are taken from the same viewing direction, but under different conditions of illumination. This technique is called *photometric stereo* and it allows one to determine surface orientation locally without smoothness assumptions.

Suppose two images $I_a(x,y)$ and $I_b(x,y)$ are obtained by varying the direction of illumination. Since there is no change in the imaging geometry, each picture element (x,y) in the two images corresponds to the same object point and hence to the same gradient (p,q). This means that one does not have the problem of first identifying corresponding points in multiple views, as happens in binocular stereo. The effect of varying the direction of illumination is to change the reflectance map $R(p,q)$ that characterizes the imaging situation.

Let the reflectance maps for the two imaging situations be $R_a(p,q)$ and $R_b(p,q)$ respectively. Suppose the point (x_0,y_0) has measured intensities

$$I_a(x_0,y_0) = \alpha \text{ and } I_b(x_0,y_0) = \beta \tag{23}$$

One obtains two equations for the gradient (p,q)

$$R_a(p,q) = \alpha \quad \text{and} \quad R_b(p,q) = \beta \tag{24}$$

The intersection of the corresponding iso-brightness contours in gradient space determines the surface orientation at (x_0,y_0). There may be more than one solution. Additional information, such as a third image, may be needed to determine the answer. This idea is illustrated graphically in Figure 16.

Sometimes, three or more images also allow the determination of the reflectance factor ρ at each (x,y). For reflectance given by equation (2), three views are enough to determine both the surface orientation and the reflectance factor ρ at each image point, provided the three directions of illumination are not coplanar (Woodham (1980a)).

Photometric stereo is easy to implement. The stereo computation, after an initial calibration step, is purely local and may be implemented by table lookup, allowing realtime performance. Photometric stereo is a practical scheme for environments, such as industrial inspection, where the illumination can be controlled.

The multiple images required for photometric stereo can be obtained by moving a single light source, by using multiple sources individually calibrated or by rotating the object and imaging device together to simulate the movement of a single light source. An equivalent to photometric stereo can also be achieved in a single view by using multiple sources that can be separated by color.

Several practical implementations of photometric stereo have been attempted. Silver (1980) achieved high accuracy under conditions of precise calibration. Ikeuchi (1981) used extended light sources to improve accuracy by altering the shape of the reflectance map. This proved useful for specular surfaces. Coleman and Jain (1982) obtained qualitatively correct results using only an approximation to the correct reflectance map. Initial experimentation, however, has been confined to single convex objects to avoid the effects of mutual illumination.

4.1 Comparison with Binocular Stereo

Photometric stereo is complementary to methods based on the identification of corresponding points in two images obtained from different viewpoints:

(1) Binocular stereo allows the accurate determination of distance to the surface. Photometric stereo is best when surface orientation is to be found.

(2) Binocular stereo works well on rough surfaces with discontinuities in surface orientation. Photometric stereo is best when surfaces are smooth with few discontinuities.

(3) Binocular stereo works well on textured surfaces with discontinuities in surface reflectance. Photometric stereo is best when surfaces have uniform optical properties.

Photometric stereo does have some distinct advantages:

(1) Since the images are obtained from the same point of view, there is no difficulty identifying corresponding points in two images. This is the major computational task of binocular stereo.

(2) In certain circumstances, the surface reflectance factor ρ can also be found because the effect of surface orientation on image brightness can be removed. Binocular stereo does not provide this capability.

(3) In many applications, a description of object shape based on surface orientation is preferable to a description based on range or altitude above a reference plane.

5. Discussion

The fundamental problem in computational vision is to reconstruct a three-dimensional representation of the scene from its two-dimensional projection onto the image plane. Substantial progress has been made both to define the levels of representation required and to identify specific modules working at each level. Using the terminology of Marr (1982), three levels of

representation emerge. They are:

1. the primal sketch
2. the 2 1/2-D sketch
3. 3-D representation

The primal sketch represents locations in an image where brightness changes abruptly. The primal sketch is a description of image properties. Surface properties are assigned to the locations described in the primal sketch to form the raw 2 1/2-D sketch. Surface properties include: reflectance changes, illumination changes, discontinuities in depth and discontinuities in surface orientation. Information in a raw 2 1/2-D sketch is likely to be sparse. The raw 2 1/2-D sketch is interpolated to define surface properties everywhere in the image. The result is called the full 2 1/2-D sketch. A representation of properties of the visible surfaces in a scene has also been called intrinsic images by Barrow and Tenenbaum (1978). When the particular property is surface orientation, it has also been called a needle diagram by Horn (1982). Both primal sketch and 2 1/2-D sketch representations are defined in a viewer-centered coordinate system. That is, they define spatially varying functions over the plane forced by the imaging geometry.

Photometric methods attempt to determine surface orientation directly from image intensity at each location in the image plane. To the extent that the laws of physical optics are adequately represented, the image irradiance equation must, of necessity, be correct. At this level, it represents a theory of the problem, but not a theory of solutions to that problem. The requirements of photometric methods for determining shape from shading may seem overly restrictive in comparison to the robustness of the human visual system. Nevertheless, the theory confirms that there are trade-offs between shape, surface reflectance and illumination that cannot be resolved. One must accept the notion that the perception of shape, although apparently unambiguous, may sometimes be incorrect. Of course, the assumptions and constraints imposed by the methods discussed in this chapter say nothing about the assumptions and constraints

that may be found to underly the human visual system.

There are two paths to follow. One path involves restricting the scene domain. Generic knowledge of surface properties and illumination can lead to precise scene reconstructions in a variety of circumstances. This is the path taken by the photometric approach, the results of which have been summarized in this chapter. The other path involves discovering general perceptual principles that apply regardless of scene domain. Constraints imposed by the perceiver may imply a unique reconstruction even though the physics of the problem does not. It seems that this must apply in human vision. Current work in shape from shading and shape from stereo may be retreating from the notion that a precise numeric reconstruction of shape is required, arguing instead that qualitative symbolic descriptions are sufficient.

The photometric approach also provides a theory to account for certain qualitative effects that occur in real images. Section 3.2 argued that an analysis of intensity profiles across edges in the polyhedral domain can often lead to their direct interpretation as convex, concave or occluding. An image irradiance equation also helps to constrain what isn't seen. For example, Grimson (1981) uses the image irradiance equation in his work on binocular stereo to constrain surface interpolation in areas that did not produce assertions at the level of the raw 2 1/2-D sketch.

Photometric methods are being applied to problems in remote sensing. Computer-based analysis of remotely sensed data typically assumes that image irradiance can be related directly to ground cover. This has been successful for wheat and other crop inventories in the flat plains of North America. But, it has been largely unsuccessful for forest inventory in British Columbia and other areas of rugged terrain. Models of surface reflectance predict conditions of illumination and terrain that cause the effects of topography to dominate those of ground cover. The ability to predict image features for known terrain and illumination has been used to automati-

cally register Landsat images to digital terrain models (DTM's) (Horn and Bachman (1978), Little (1982)). Even when a remotely sensed image has been registered to a surface model, it can be difficult to separate changes in ground cover from topographic and illumination effects (Woodham (1980b), Sjoberg (1982) and Teillet et al. (1982)). Work in this area is continuing.

In conclusion, it seems that any algorithm for surface reconstruction that determines a unique solution must either embody assumptions about the scene domain or about the perceiver. The tools summarized in this chapter can help to make those assumptions explicit. They help to make computational vision a theoretical science as well as an experimental one.

Acknowledgment

Figure 2 is adapted from the book *Light and Vision* (copyright 1966, LIFE Science Library, Time Inc., New York, NY). Figure 3 is adapted from an advertisement (copyright 1979, Merle Norman Cosmetics, Inc., Los Angeles, CA). Figure 5 appeared as the cover of *Science,* May 7, 1976. These figures are reproduced here with the permission of the publishers. Figures 1 and 4 were produced at the UBC Laboratory for Computational Vision. The remaining figures were prepared by Nedenia M. Holm.

References

Barrow, H.G. & J.M. Tenenbaum (1978), Recovering intrinsic scene characteristics from images, in *Computer Vision Systems,* A.R. Hanson & E.M. Riseman (eds.), pp 3-26, Academic Press.

Barrow, H.G. & J.M. Tenenbaum (1981), Interpreting line drawings as three-dimensional surfaces, *Artificial Intelligence* (17)75-116.

Brooks, M.J. (1979) Surface normals from closed paths, *Proc. IJCAI-79,* pp 98-101, Tokyo, Japan.

Bruss, A.R. (1981), The image irradiance equation: its solution and application, AI-TR-623. MIT AI Laboratory, Cambridge, Mass.

Coleman, Jr., E.N. & R. Jain (1982), Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry, *CGIP* (18)309-328.

Cook, R.L. & K.E. Torrance (1982), A reflectance model for computer graphics, *ACM Trans. on Graphics* (1)7-24.

Draper, S.W. (1981), The use of gradient space and dual space in line-drawing interpretation, *Artificial Intelligence* (17)461-508.

Grimson, W.E.L. (1981), *From Images to Surfaces,* MIT Press.

Hapke, B. (1971), Optical properties of the lunar surface, in *Physics and Astronomy of the Moon (2nd edition),* Z. Kopal (ed.), pp 155-211, Academic Press.

Horn, B.K.P. (1975), Obtaining shape from shading information, in *The Psychology of Computer Vision,* P.H. Winston (ed.), pp 115-155, McGraw-Hill.

Hilbert, D. & S. Cohn-Vossen (1952), *Geometry and the Imagination,* Chelsea.

Horn, B.K.P. (1977), Understanding Image Intensities, *Artificial Intelligence* (8)201-231.

Horn, B.K.P. (1982), Sequins and quills - representations for surface topography, in *Representation of 3-Dimensional Objects,* R. Bajcsy (ed.), Springer-Verlag.

Horn, B.K.P. & B.L. Bachman (1978), Using synthetic images to register real images with surface models, *CACM* (21)914-924.

Horn, B.K.P. & R.W. Sjoberg (1979), Calculating the reflectance map, *Applied Optics* (18)1770-1779

Huffman, D.A. (1971), Impossible objects as nonsense sentences, in *Machine Intelligence 6,* B. Meltzer and D. Michie (eds.), pp 295-323, Edinburgh Univ. Press.

Ikeuchi, K. (1981), Determining surface orientation of specular surfaces by using the photometric stereo method, *IEEE-PAMI* (3)661-669.

Ikeuchi, K. & B.K.P. Horn (1981), Numerical shape from shading and occluding boundaries, *Artificial Intelligence* (17)141-184.

Little, J.J. (1982), Automatic registration of Landsat MSS images to digital elevation models, *Proc. Workshop on Computer Vision: Representation and Control,* pp 178-184, Rindge, New Hampshire.

Mackworth, A.K. (1973), Interpreting pictures of polyhedral scenes, *Artificial Intelligence* (4)121-137.

Mackworth, A.K. (1977), Consistency in networks of relations, *Artificial Intelligence* (8)99-118.

Marr, D. (1982), *Vision,* W.H. Freeman.

Marr, D. (1977), Analysis of occluding contour, *Proc. R. Soc. Lond. B.* (197)441-475.

Minnaert, M. (1961), Photometry of the moon, in *The Solar System Vol. III: Planets and Satellites,* G.P. Kuiper & B.M. Middlehurst (eds.), pp 213-248, Univ. Chicago Press.

Nicodemus, F.E., J.C. Richmond, J.J. Hsia, I.W. Ginsburg & T. Limperis (1977), Geometrical considerations and nomenclature for reflectance, NBS Monograph 160, National Bureau of Standards, Dept. of Commerce, Washington, D.C.

Rindfleisch, T. (1966), Photometric method for lunar topography, *Photogrammetric Engineering* (32)262-276.

Silver, W.M. (1980), Determining shape and reflectance using multiple images, SM Thesis, Dept. Electrical Engineering and Computer Science, MIT, Cambridge, Mass.

Sjoberg, R.W. (1982) Atmospheric effects in satellite imaging of mountainous terrain, AI-TR-688, MIT AI Laboratory, Cambridge, Mass.

Teillet, P.M., B. Guindon & D.G. Goodenough (1982), On the slope-aspect correction of multispectral scanner data, *Canadian Journal of Remote Sensing,* (8)84-106.

Van Diggelen, J. (1951), A photometric investigation of the slopes and the heights of the ranges of hills in the maria of the moon, *Bulletin of the Astronomical Institute of the Netherlands* Vol. XI, No. 423.

Woodham, R.J. (1977), A cooperative algorithm for determining surface orientation from a single view, *Proc. IJCAI-77,* pp 635-641, Cambridge, Mass.

Woodham, R.J. (1978), Reflectance map techniques for analyzing surface defects in metal castings, AI-TR-457, MIT AI Laboratory, Cambridge, Mass.

Woodham, R.J. (1980a), Photometric method for determining surface orientation from multiple images, *Optical Engineering* (19)139-144.

Woodham, R.J. (1980b), Using digital terrain data to model image formation in remote sensing, *Proc. SPIE* Vol. 238, pp 361-369.

Woodham, R.J. (1981), Analysing images of curved surfaces, *Artificial Intelligence* (17)117-140.

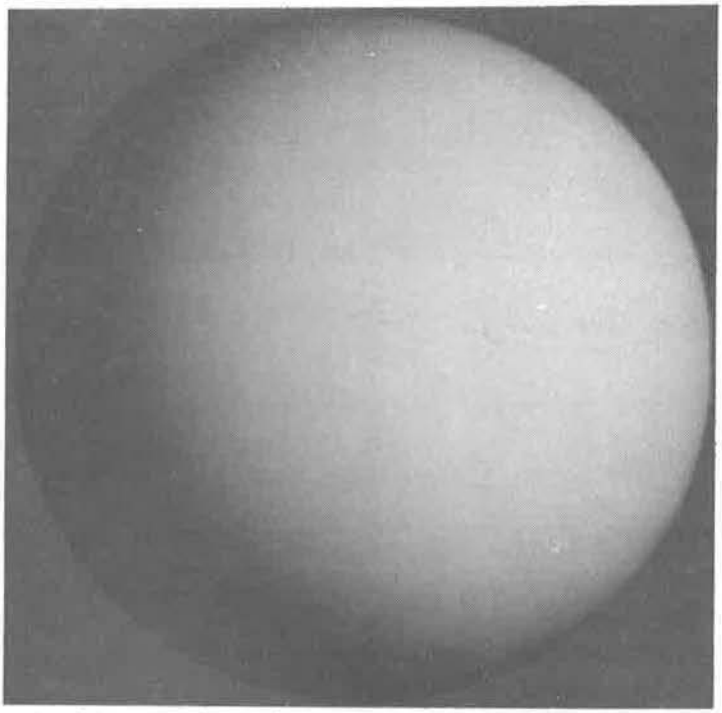

Figure 1. A synthesized image of a sphere. For a given object material, viewer position and distribution of light sources, image brightness is a function of surface orientation. In Computer Graphics, images are synthesized by using a suitable reflectance model and illumination to assign a brightness value at each visible surface element.

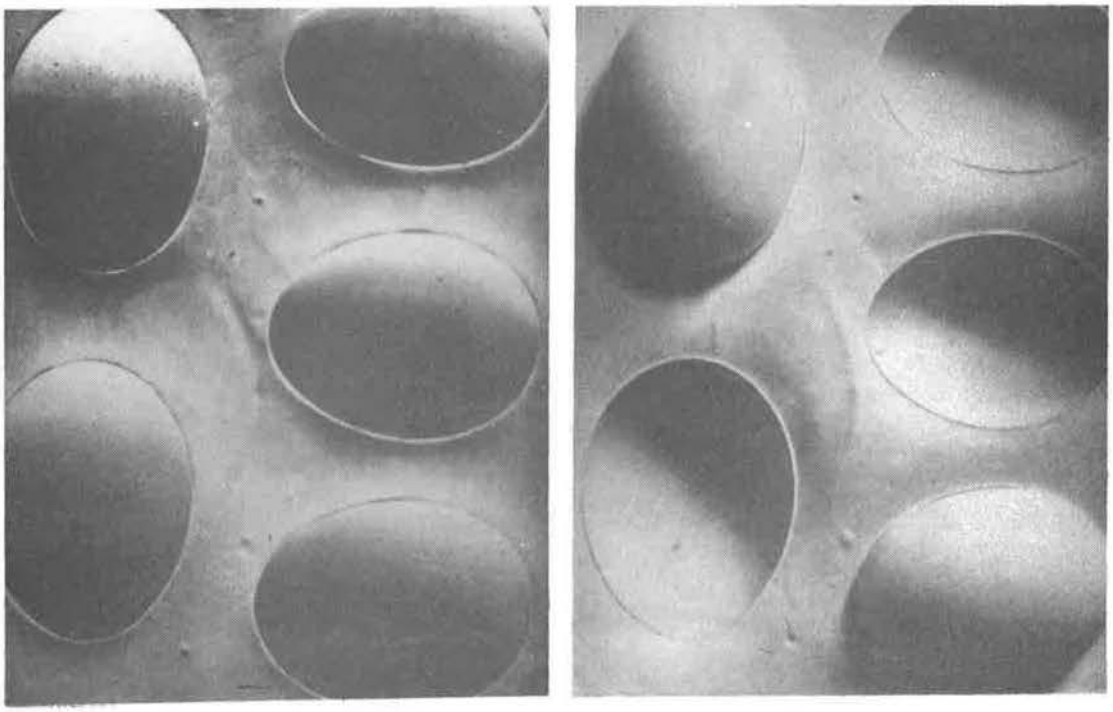

Figure 2. Lighting conditions alter the perception of shape. The image in (a) seems to show a container holding five eggs. But there are no eggs in the container. In (b), the image is obtained under more natural lighting conditions. Two eggs have been added and the three remaining hollows are easily identified. To see image (a) "correctly", turn it upside down. (Adapted from *Light and Vision,* LIFE Science Library, Time Inc., New York, NY, 1966)

Figure 3. Altering surface properties can alter the perception of object shape. The image in (a) shows a model's face without makeup. The image in (b) demonstrates the use of makeup to alter the perception of shape. (Adapted from an advertisement for Merle Norman Cosmetics, Inc., Los Angeles, CA, 1979)

Figure 4. The full moon. The moon has remarkable photometric properties arising from the special structure of its surface. Quantitative shape from shading methods were first developed to deal with the lunar surface.

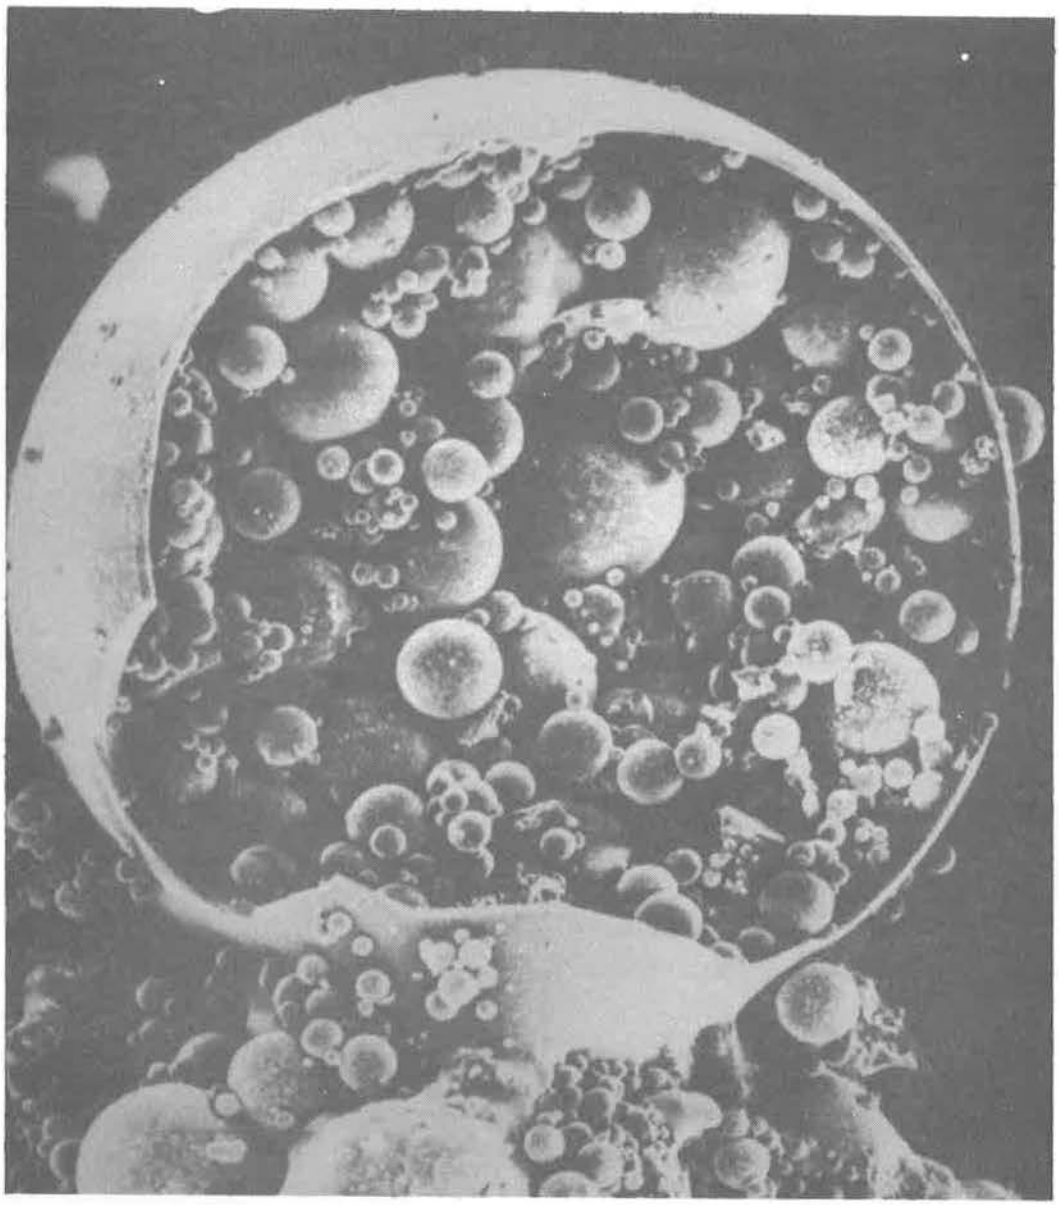

Figure 5. "Spheres within spheres" as seen by a scanning electron microscope (SEM). Shading in a SEM image, although not natural, still provides strong cues to shape. (Adapted from Fisher, G. L. et al., "Fly Ash Collected from Electrostatic Precipitators", *Science,* Cover, Vol. 192, May 7, 1976)

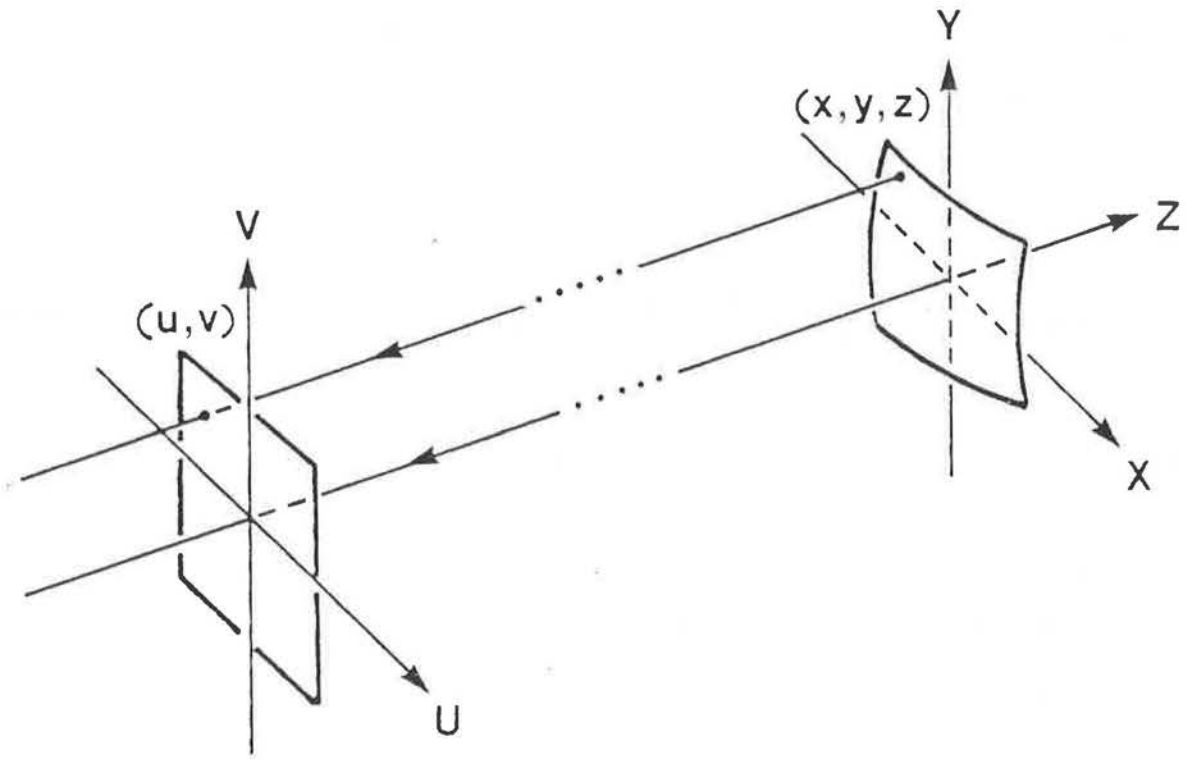

ORTHOGRAPHIC PROJECTION

u = x v = y

Figure 6. An orthographic projection. In an orthographic projection all rays from object surface to image are parallel. With appropriate scaling of the image uv-plane, image coordinates (x, y) and surface coordinates (x, y) can be used interchangeably.

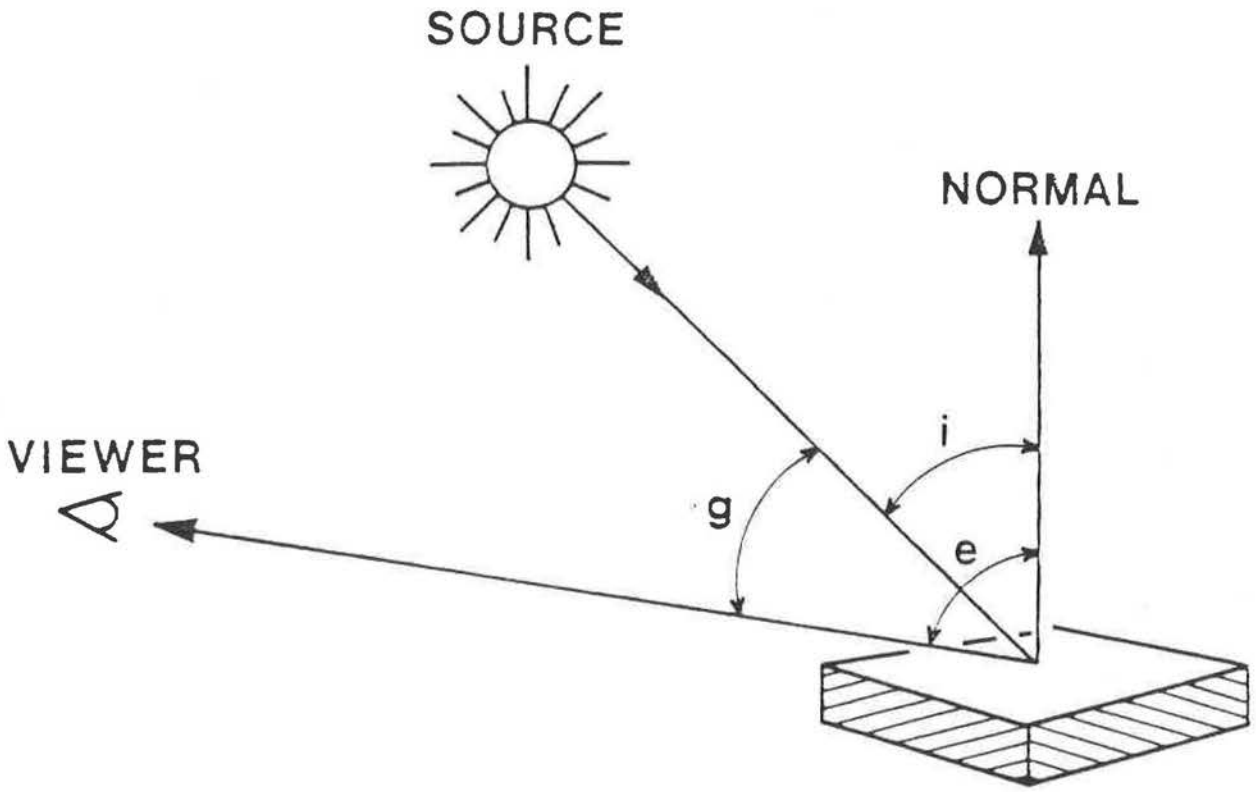

Figure 7. Defining the angles i, e and g. The incident angle i is the angle between the incident ray and the surface normal. The emergent angle e is the angle between the emergent ray and the surface normal. The phase angle g is the angle between the incident and emergent rays.

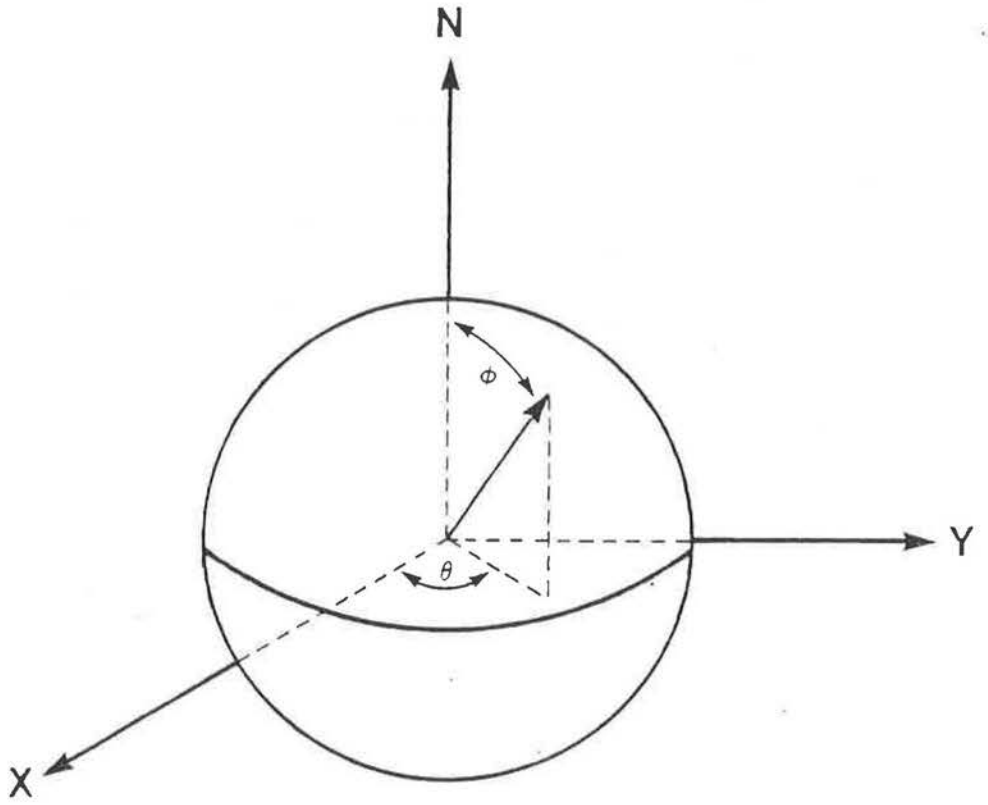

Figure 8. Spherical coordinates. The azimuth angle θ is measured counter-clockwise from the x-axis in the xy-plane and the polar angle ϕ is measured from the z-axis.

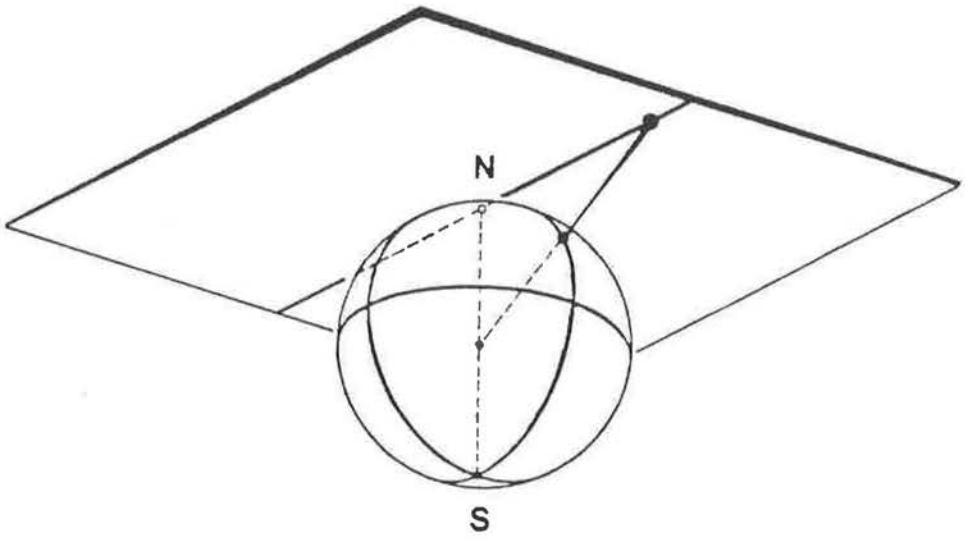

Figure 9. The gradient (p, q) is obtained by extending the point on the Gaussian sphere outward from the center (i.e., from the point (0,0,0)) until it intersect the plane $z = -1$. Only the northern hemisphere of the Gaussian sphere is represented in gradient space.

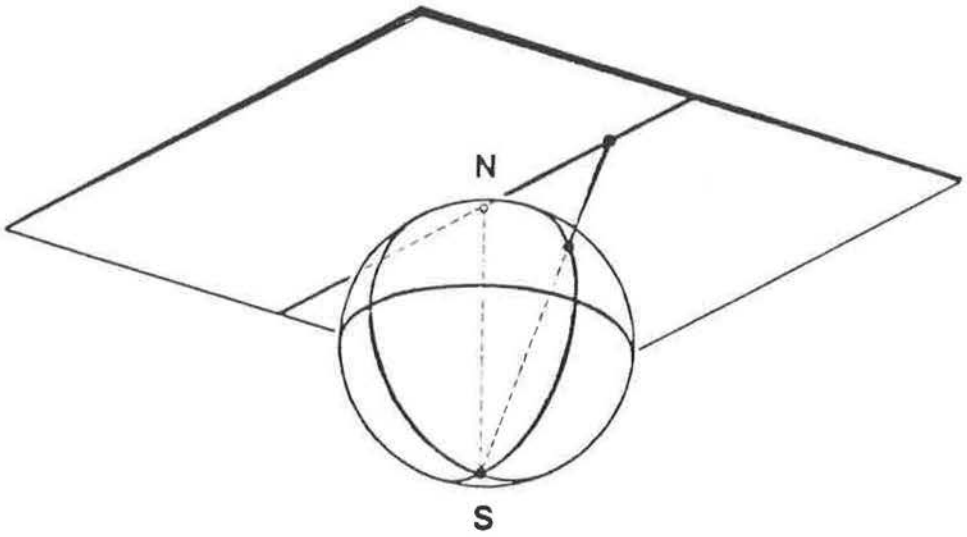

Figure 10. Stereographic coordinates (f, g) are obtained by extending each point on the Gaussian sphere outward from the south pole (i.e., from the point (0,0,1)) until it intersect the plane $z = -1$. All points on the Gaussian sphere, except the south pole itself, are represented in the stereographic projection.

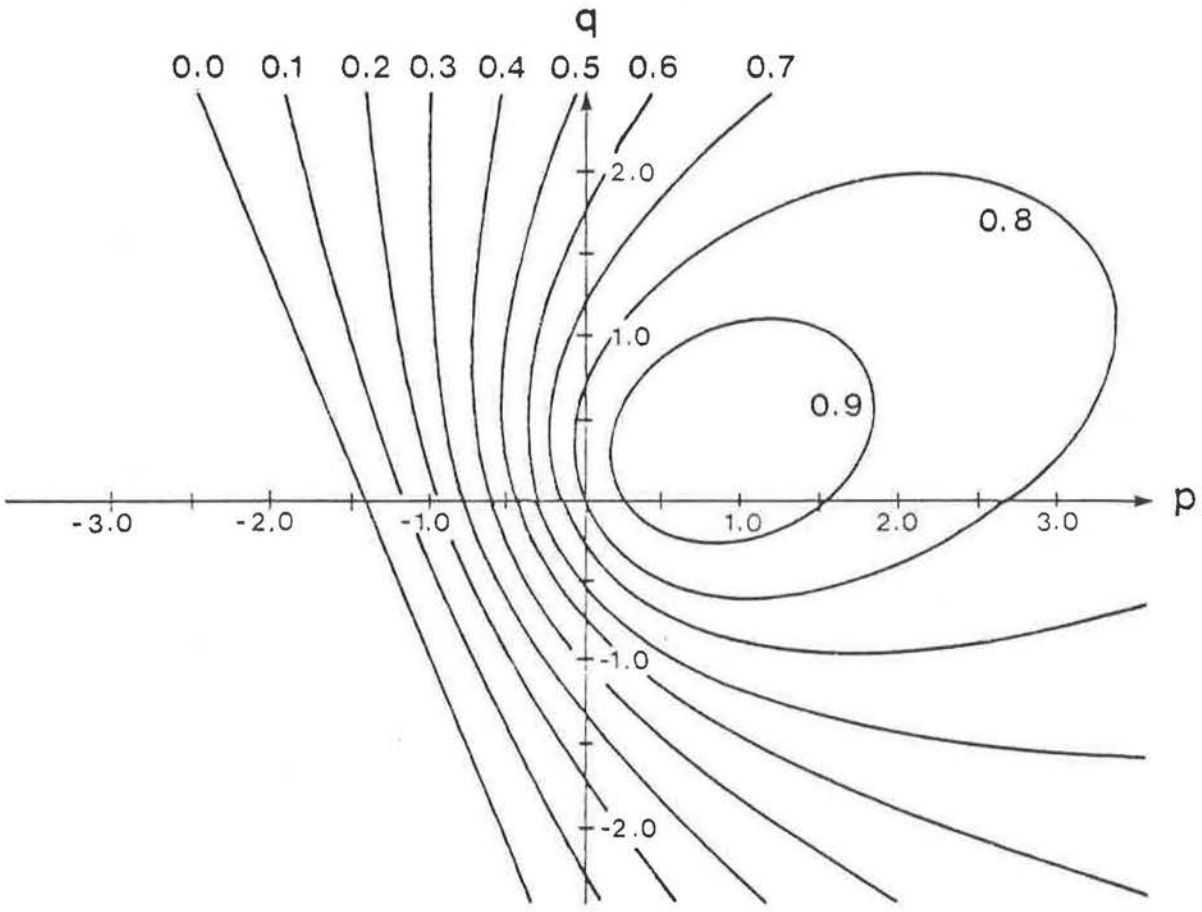

Figure 11. The reflectance map $R_2(p,q)$ for a perfectly diffuse surface illuminated from a distant point source with gradient $p_s = 0.7$ and $q_s = 0.3$ ($\rho = 1.0$). The reflectance map is plotted as a series of contours spaced 0.1 units apart.

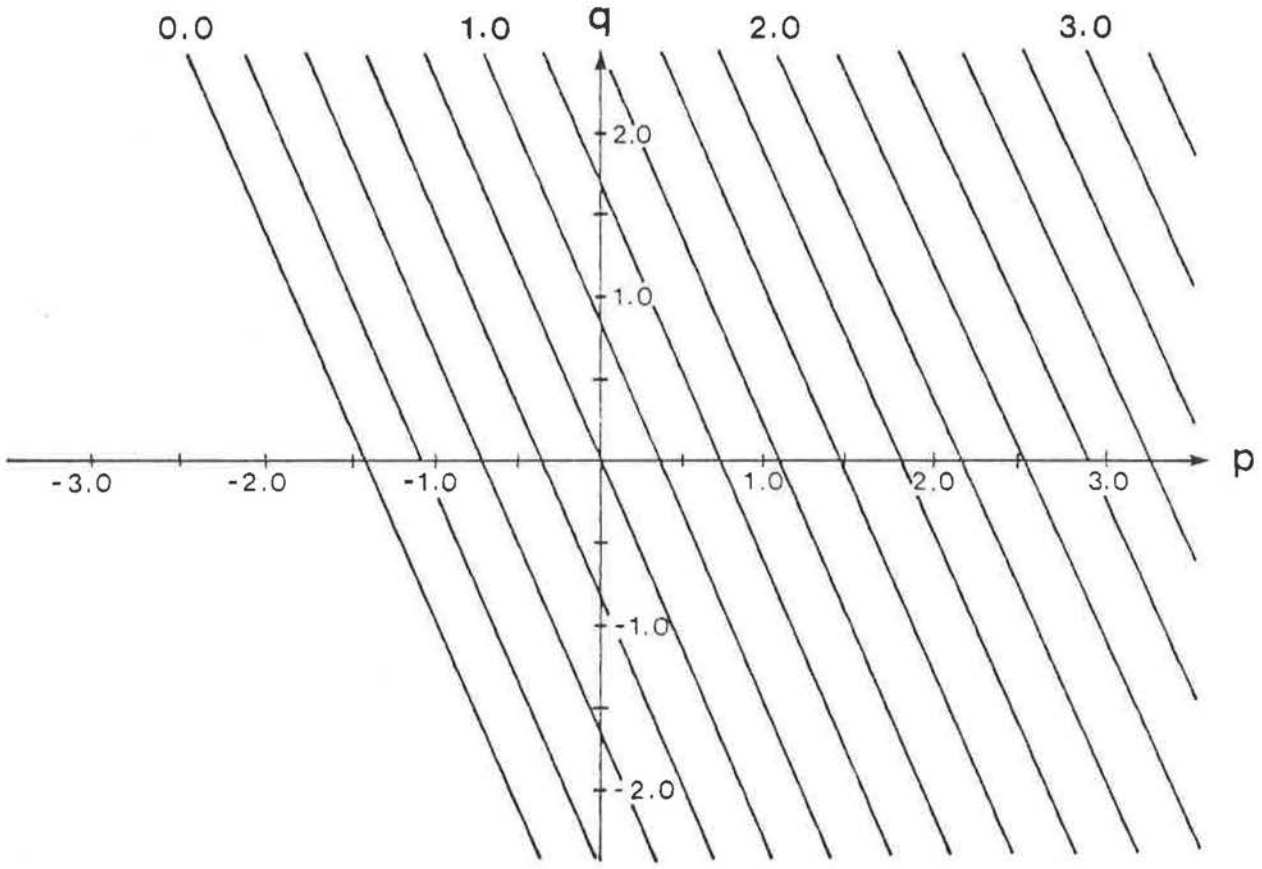

Figure 12. The reflectance map $R_3(p,q)$, similar to material in the maria of the moon, with illumination from a distant point source with gradient $p_s = 0.7$ and $q_s = 0.3$ ($\rho = 1.0$). The reflectance map is plotted as a series of contours spaced 0.2 units apart.

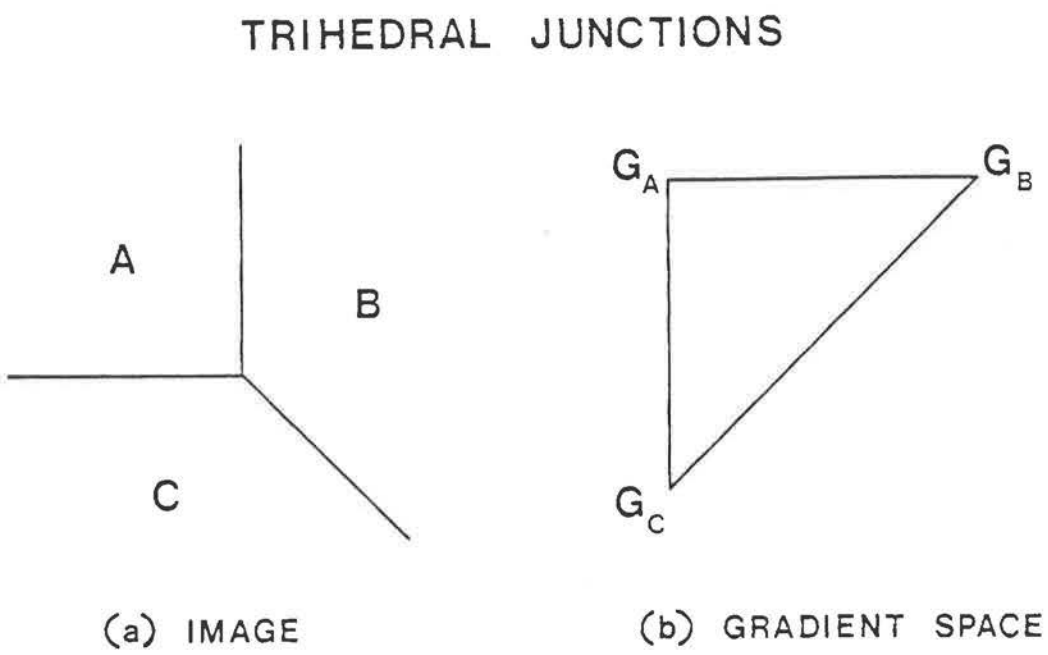

Figure 13. A trihedral vertex formed by planes A, B and C is shown in (a). The triangle joining the corresponding gradients GA, GB and GC is shown in (b). The line in gradient space joining the gradients of two planes is perpendicular to the line in the image formed by the intersection of the two planes.

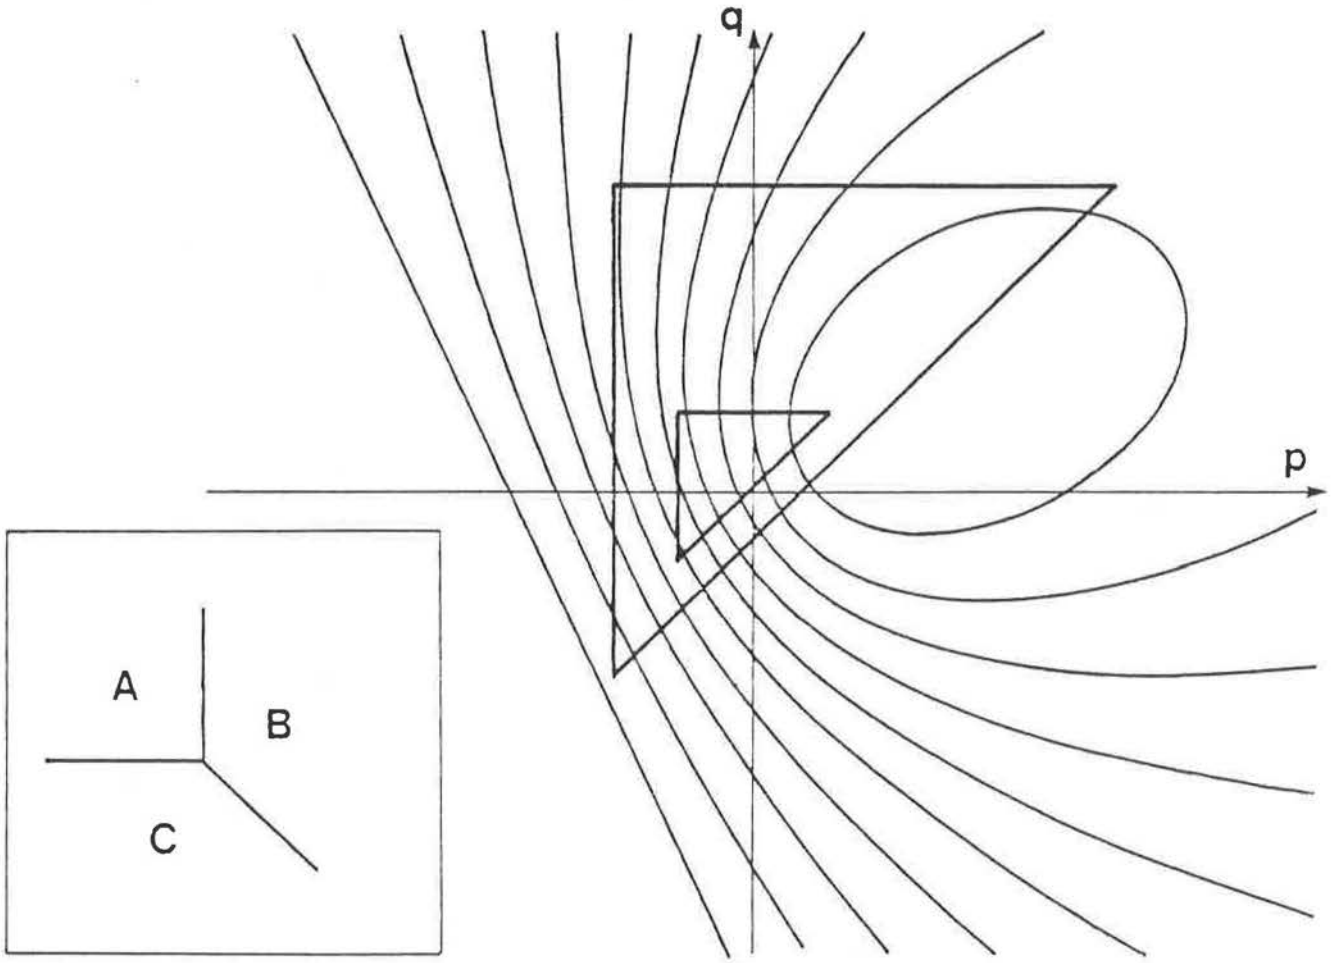

Figure 14. Two possible gradient space triangles corresponding to the trihedral vertex formed by planes A, B and C (shown inset). Intensity measurements on the three planes determine the true position and scale of the triangle in gradient space.

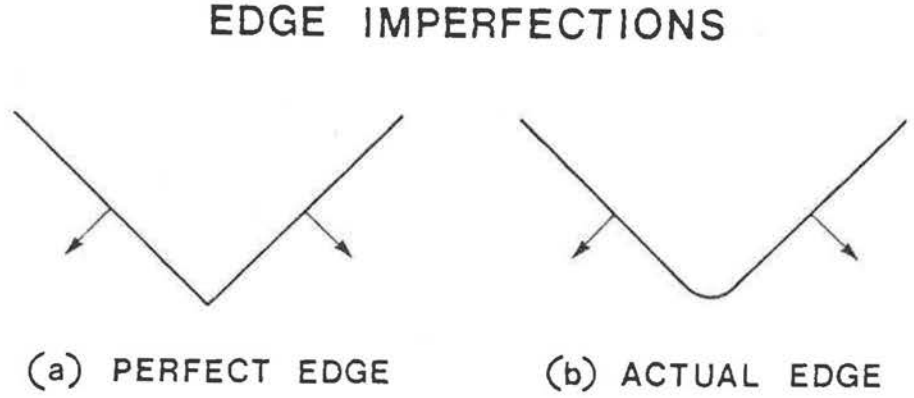

Figure 15. Edge imperfections. In (a), surface orientation is discontinuous where the two planes intersect. Real edges are often rounded. In (b), surface orientation varies continuously from one plane to the other.

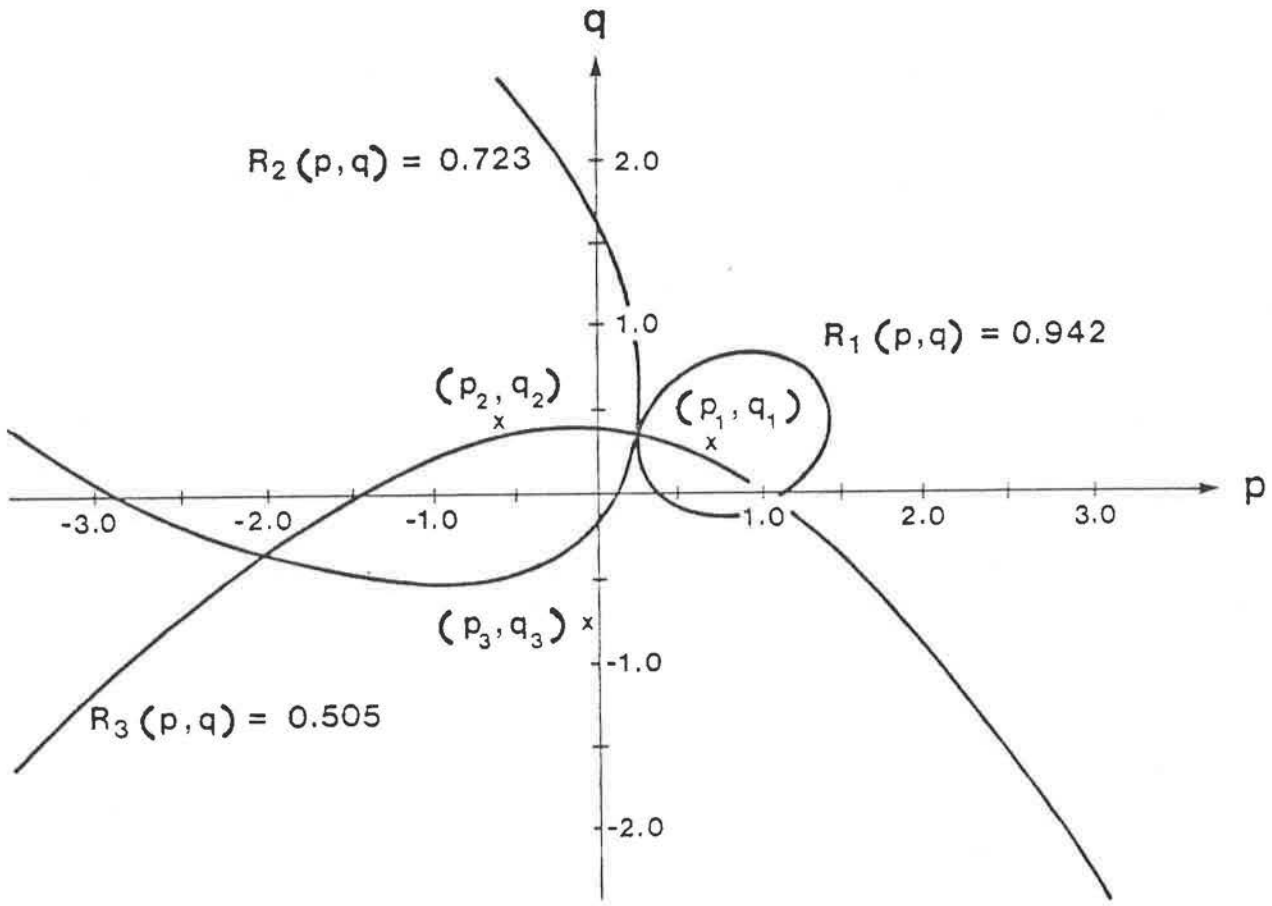

Figure 18. Iso-brightness contours from three reflectance maps are superimposed. Each corresponds to the intensity value at (x, y) obtained from three images obtained under different conditions of illumination. The three intensity measurements were $I_1(x, y) = 0.942$, $I_2(x, y) = 0.723$ and $I_3(x, y) = 0.505$. The common point of intersection determines the gradient (p, q).