10

Reflectance Map: Photometric Stereo

We are now ready to look more carefully at how the brightness pattern in an image depends on the shape of the object being imaged. In this chapter we develop the photometric stereo method for recovering the orientation of surface patches from a number of images taken under different lighting conditions. In the next chapter we consider the more difficult problem of recovering surface shape from a single image. The photometric stereo method is simple to implement, but requires control of the lighting. The orientation of a patch of the surface corresponding to a given picture cell is determined by means of a simple lookup table, which is built using a calibration object of known shape. The orientations of the surface patches are conveniently represented as a needle diagram.

In order to understand how these methods work, we need to know something about radiometry. We have to make more precise the term "brightness" and learn how image irradiance depends on scene radiance. The detailed dependence of surface reflection on the geometry of incident and emitted rays is given by the bidirectional reflectance distribution function. The reflectance map can be derived from that function and the distribution of light sources. The reflectance map is useful because it makes explicit the relationship between surface orientation and brightness. It is of critical importance in recovering surface orientation from measurements of brightness. In practice, one can determine the reflectance map experimen-

tally, using a calibration object of known shape, rather than computing it from the bidirectional reflectance distribution function. The reader with less interest in radiometry may want to skip the first few sections and proceed directly to the discussion of the reflectance map.

Finally, we introduce the image irradiance equation, which will play an important role in the next chapter.

10.1 Image Brightness

The image of a three-dimensional object depends on its shape, its reflectance properties, and the distribution of light sources. Three views of a portion of the surface of Mars are shown in figure 10-1. The three images were obtained using a camera in a fixed position. The differences between them are due to differences in lighting conditions. While clearly the same underlying surface is being portrayed, the detailed patterns of brightness are quite different. Some edges that show up in strong contrast in one of the pictures, for example, are not visible in another. This example demonstrates the importance of the position and distribution of the sources of illumination in determining the brightness pattern.

The image of a three-dimensional object also depends on the position of the object relative to the imaging system and on the object's attitude in space. In the case of a solid of revolution, attitude has only two degrees of freedom, so that the attitude can be specified by giving the direction of the axis of revolution. Figure 10-2 shows how the silhouette of a push-pin changes with its attitude. More important, note how the brightness pattern within the silhouette changes, particularly where the glossy surface reflects light from the source directly toward the viewer. These variations in the image of an object become even more complicated when we consider general shapes, for then attitude has three degrees of freedom. It is clear that the binary image-processing methods discussed earlier will be of no avail in that situation.

We have to understand how the image is formed in order to turn the process around and recover information about the permanent properties of the object, such as its shape and surface reflectance. However, to understand how the brightness at a particular point in the image is determined, we must first discuss radiometry.

10.2 Radiometry

The amount of light falling on a surface is called the *irradiance*. It is the power per unit area ($\mathrm{W \cdot m^{-2}}$—watts per square meter) incident on the surface. The amount of light radiated from a surface is called the *radiance*. It is the power per unit area per unit solid angle ($\mathrm{W \cdot m^{-2} \cdot sr^{-1}}$—watts per

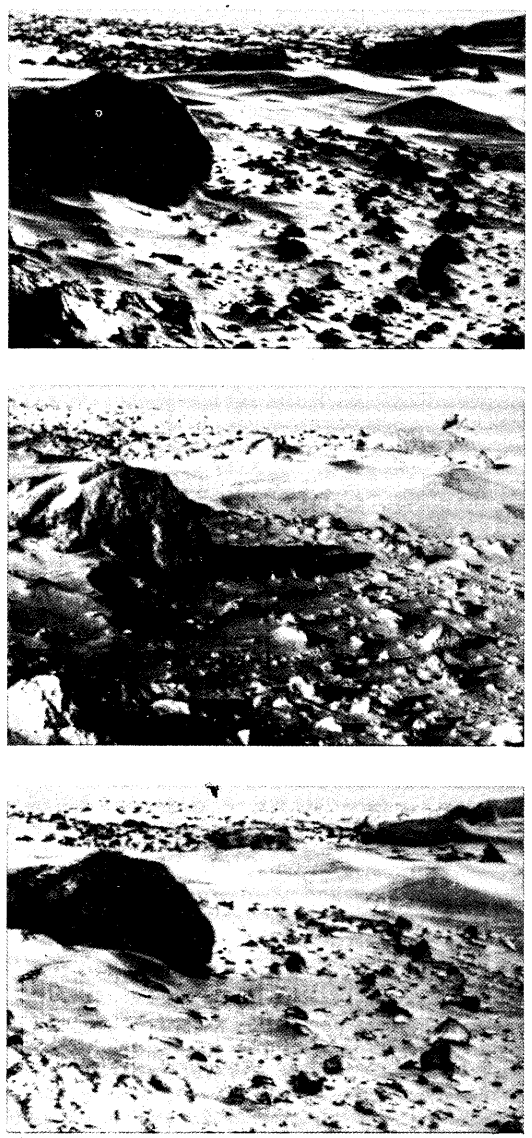

Figure 10-1. The appearance of a surface depends greatly on the lighting conditions. We have to understand how images are formed if we are to recover information about the surface from one or more images. Shown here are three views of the surface of Mars taken by Viking Lander I. (Pictures IPL PIC ID 77/03/06/134059, 77/02/09/172643, and 77/02/05/042935, courtesy of the National Space Science Data Center, Greenbelt, Maryland.)

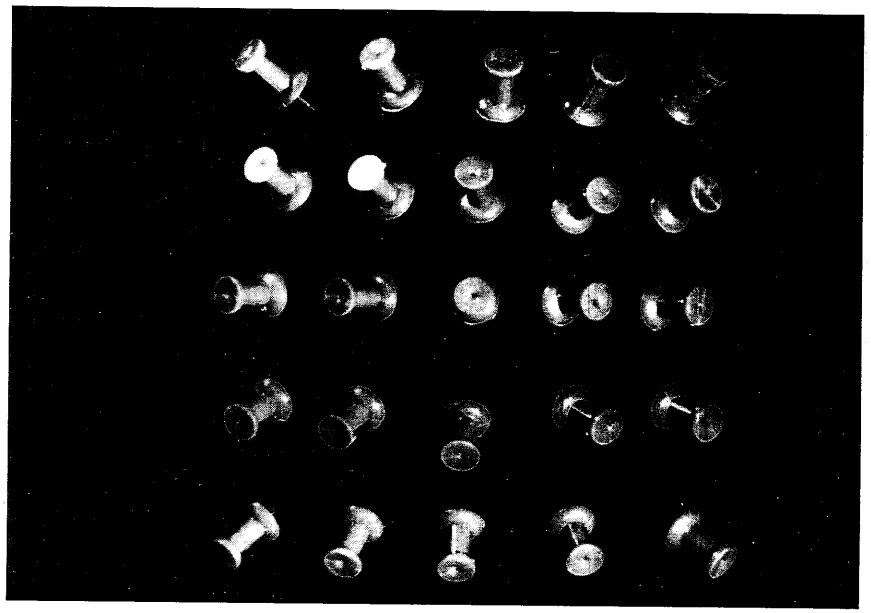

Figure 10-2. The appearance of an object depends greatly on its attitude in space relative to the viewer. Not only does the outline vary, but the brightness pattern within the silhouette changes.

square meter per steradian) emitted from the surface. The complexity of the latter concept stems from the fact that a surface can radiate into a whole hemisphere of possible directions and can radiate different amounts of energy in different directions.

The solid angle of a cone of directions is defined as the area cut out by the cone on the unit sphere. A hemisphere of directions, for example, has a solid angle of 2π. A small planar patch of area A at distance R from the origin (figure 10-3) subtends a solid angle

$$\Omega = \frac{A \cos\theta}{R^2},$$

where θ is the angle between a surface normal and a line connecting the patch to the origin.

Brightness is determined by the amount of energy an imaging system receives per unit apparent area. The unit area mentioned in the definition of radiance is the foreshortened area, the surface area multiplied by the cosine of the angle between a perpendicular to the surface and the specified direction.

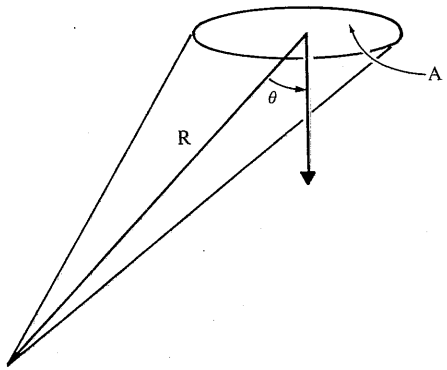

Figure 10-3. The solid angle subtended by a small patch is proportional to its area A and the cosine of the angle of inclination θ; it is inversely proportional to the square of its distance R from the origin.

10.3 Image Formation

We next find the relationship between the radiance at a point on an object (scene radiance) and the irradiance at the corresponding point in the image (image irradiance). Consider a lens of diameter d at a distance f from the image plane (figure 10-4). Let a patch on the surface of the object have area δO, while the corresponding image patch has area δI. Suppose that the ray from the object patch to the center of the lens makes an angle α with the optical axis and that there is an angle θ between this ray and a surface normal. The object patch is at a distance $-z$ from the lens, measured along the optical axis. (The minus sign arises from our convention for the coordinate system, with the z-axis pointing toward the image plane.)

The ratio of the area of the object patch to that of the image patch is determined by the distances of these patches from the lens and by foreshortening. Rays passing through the center of the lens are not deflected. As a result, the solid angle of the cone of rays leading to the patch on the object is equal to the solid angle of the cone of rays leading to the corresponding patch in the image. The apparent area of the image patch as seen from the center of the lens is $\delta I \cos\alpha$, while the distance of this patch from the center of the lens is $f/\cos\alpha$. Thus the solid angle subtended by this patch is just

$$\frac{\delta I \cos\alpha}{(f/\cos\alpha)^2}.$$

Similarly, the solid angle of the patch on the object as seen from the center

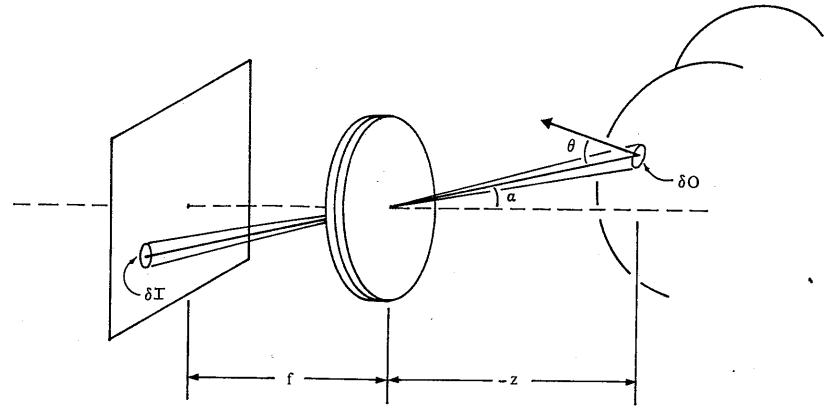

Figure 10-4. To see how image irradiance is related to the radiance of the surface, we must determine the size of the region in the image that corresponds to the patch on the surface.

of the lens is

$$\frac{\delta O \cos\theta}{(z/\cos\alpha)^2}.$$

If these two solid angles are to be equal, we must have

$$\frac{\delta O}{\delta I} = \frac{\cos\alpha}{\cos\theta}\left(\frac{z}{f}\right)^2.$$

Next, we need to determine how much of the light emitted by the surface makes its way through the lens (figure 10-5). The solid angle subtended by the lens, as seen from the object patch, is

$$\Omega = \frac{\pi}{4}\frac{d^2\cos\alpha}{(z/\cos\alpha)^2} = \frac{\pi}{4}\left(\frac{d}{z}\right)^2 \cos^3\alpha.$$

Thus the power of the light originating on the patch and passing through the lens is

$$\delta P = L\,\delta O\,\Omega\cos\theta = L\,\delta O\,\frac{\pi}{4}\left(\frac{d}{z}\right)^2 \cos^3\alpha\,\cos\theta,$$

where L is the radiance of the surface in the direction toward the lens. This power is concentrated in the image (if we ignore losses in the lens). Since

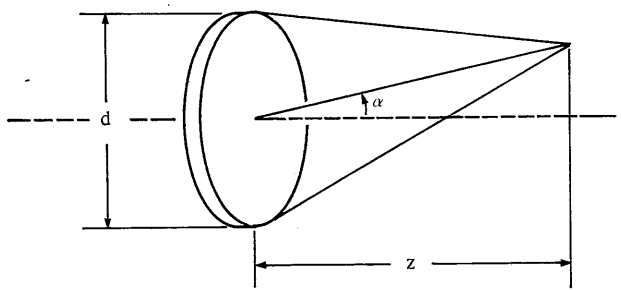

Figure 10-5. To see how image irradiance is related to the radiance of the surface, we must determine how much of the light emitted by the surface is gathered by the lens. This depends on the solid angle that the lens subtends when viewed from the light-emitting surface.

no light from other areas reaches this image patch, we have

$$E = \frac{\delta P}{\delta I} = L \frac{\delta O}{\delta I} \frac{\pi}{4} \left(\frac{d}{z}\right)^2 \cos^3 \alpha \cos \theta,$$

where E is the irradiance of the image at the patch under consideration. Substituting for $\delta O / \delta I$, we finally obtain

$$E = L \frac{\pi}{4} \left(\frac{d}{f}\right)^2 \cos^4 \alpha.$$

Thus image irradiance is proportional to scene radiance. This is the fundamental relationship we shall exploit in recovering information about objects from their images.

The factor of proportionality in the formula above contains the inverse of the square of the effective *f-number*, f/d. It also contains a term that falls off with the cosine to the fourth power of the angle that the ray from the image point to the center of the lens makes with the optical axis. This falloff in sensitivity is not very important when the image covers only a narrow angle, as in a telephoto lens. Moreover, in typical lens systems, multiple apertures, lined up for axial rays, cut off part of the light for ray directions that are inclined to the optical axis. This vignetting effect, discussed in chapter 2, often causes a much more serious falloff in brightness than the cosine term. Multiple apertures are usually introduced to minimize distortions that would otherwise increase in severity as some power of the off-axis angle α. In any case, the dependence of the sensitivity of a given imaging system on the off-axis angle is fixed and can be accounted for.

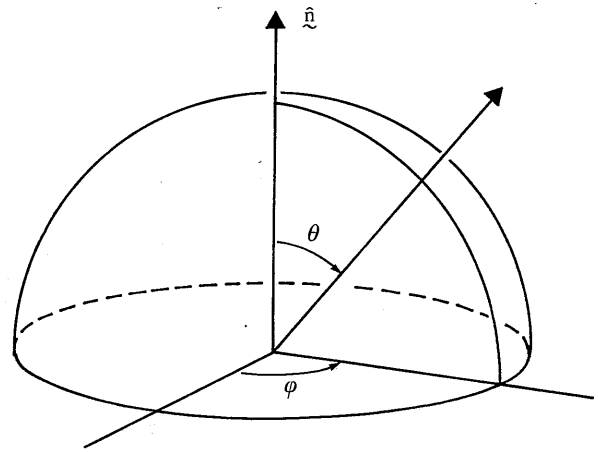

Figure 10-6. The direction of incident and emitted light rays can be specified in a local coordinate system using the polar angle θ and the azimuth ϕ.

To summarize: The important thing is that what we measure—the image irradiance E—is proportional to what we are interested in—the scene radiance L. Looked at another way, we have defined radiance so that it corresponds to our intuitive notion of "brightness," which is, after all, related to image irradiance.

10.4 Bidirectional Reflectance Distribution Function

What determines scene radiance? Scene radiance depends on the amount of light that falls on a surface and the fraction of the incident light that is reflected. It also depends, however, on the geometry of light reflection, as the example of a mirror clearly shows. That is, the radiance of a surface will generally depend on the direction from which it is viewed as well as on the direction from which it is illuminated.

We can describe these directions in terms of a local coordinate system erected on the surface (figure 10-6). Consider a perpendicular to the surface (the normal $\hat{\mathbf{n}}$) and an arbitrary reference line drawn on the surface. Directions can be described by specifying the angle θ between a ray and the normal and the angle ϕ between a perpendicular projection of the ray onto the surface and the reference line on the surface.

These angles are called the *polar angle* and *azimuth*, respectively. They allow us to specify the direction (θ_i, ϕ_i) from which light is falling on the surface and the direction (θ_e, ϕ_e) into which it is emitted toward the viewer

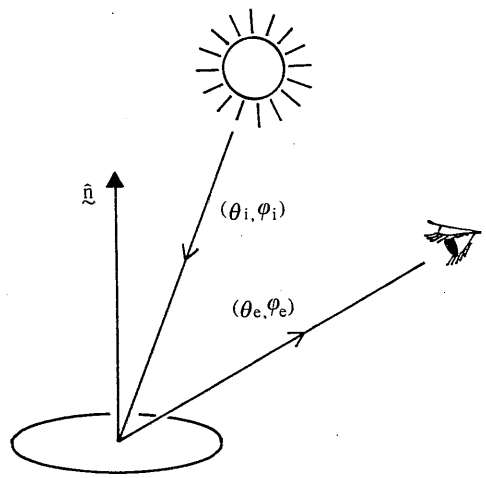

Figure 10-7. The bidirectional reflectance distribution function is the ratio of the radiance of the surface patch as viewed from the direction (θ_e, ϕ_e) to the irradiance resulting from illumination from the direction (θ_i, ϕ_i).

(figure 10-7).

We can now define the *bidirectional reflectance distribution function* (BRDF) $f(\theta_i, \phi_i; \theta_e, \phi_e)$, which tells us how bright a surface appears when viewed from one direction while light falls on it from another. Let the amount of light falling on the surface from the direction (θ_i, ϕ_i)—the irradiance—be $\delta E(\theta_i, \phi_i)$. Let the brightness of the surface as seen from the direction (θ_e, ϕ_e)—the radiance—be $\delta L(\theta_e, \phi_e)$. The BRDF is simply the ratio of radiance to irradiance,

$$f(\theta_i, \phi_i; \theta_e, \phi_e) = \frac{\delta L(\theta_e, \phi_e)}{\delta E(\theta_i, \phi_i)}.$$

We shall study specific examples of BRDFs for several idealized surfaces later in this chapter.

A function of four variables is a bit too cumbersome to use directly in exploring the relationship between image irradiance and surface shape. Fortunately, for many surfaces the radiance is not altered if the surface is rotated about the surface normal. In this case, the BRDF depends only on the difference $\phi_e - \phi_i$, not on ϕ_e and ϕ_i separately. This is certainly true of matte surfaces and specularly reflecting surfaces. It is not true of surfaces with oriented microstructure, as for example the mineral called tiger's eye or the iridescent feathers of some birds.

There is an interesting constraint on the form of the BRDF. If two

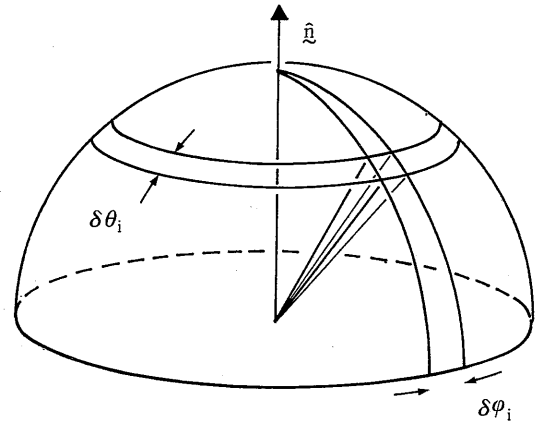

Figure 10-8. In the case of an extended light source, we must integrate the product of the bidirectional reflectance distribution function and the source radiance over all incident directions.

surfaces are in thermal equilibrium, radiation reaching one from the other must be balanced by radiation flowing in the opposite direction. If this were not the case, the surface receiving more radiation would heat up and the other would cool down, thus disturbing the equilibrium. This would violate the second law of thermodynamics. It is easy to show that this implies that the BRDF is constrained by the *Helmholtz reciprocity condition*:

$$f(\theta_i, \phi_i; \theta_e, \phi_e) = f(\theta_e, \phi_e; \theta_i, \phi_i).$$

10.5 Extended Light Sources

So far we have considered the case in which all of the light comes from one direction. In practice, there can be several light sources, or even extended sources, such as the sky. In the case of an extended source, we must consider a nonzero solid angle of directions to obtain a nonzero radiance. Consider an infinitesimal patch of the sky, of size $\delta\theta_i$ in polar angle and $\delta\phi_i$ in azimuth (figure 10-8). This patch subtends a solid angle

$$\delta\omega = \sin\theta_i \, \delta\theta_i \, \delta\phi_i.$$

If we let $E(\theta_i, \phi_i)$ be the radiance per unit solid angle coming from the direction (θ_i, ϕ_i), then the radiance from the patch under consideration is

$$E(\theta_i, \phi_i) \sin\theta_i \, \delta\theta_i \, \delta\phi_i,$$

and the total irradiance of the surface is

$$E_0 = \int_{-\pi}^{\pi} \int_0^{\pi/2} E(\theta_i, \phi_i) \sin\theta_i \cos\theta_i \, d\theta_i \, d\phi_i,$$

where the $\cos\theta_i$ term accounts for the foreshortening of the surface as seen from the direction (θ_i, ϕ_i).

To obtain the radiance of the surface we must integrate the product of the BRDF and the irradiance over the hemisphere of possible directions from which light can fall on a surface. Thus

$$L(\theta_e, \phi_e) = \int_{-\pi}^{\pi} \int_0^{\pi/2} f(\theta_i, \phi_i; \theta_e, \phi_e) E(\theta_i, \phi_i) \sin\theta_i \cos\theta_i \, d\theta_i \, d\phi_i.$$

Again, the $\cos\theta_i$ term in the integrand accounts for foreshortening. The result is a function of two variables only, θ_e and ϕ_e, which specify the direction of the ray emitted toward the viewer.

10.6 Surface Reflectance Properties

An *ideal Lambertian surface* is one that appears equally bright from all viewing directions and reflects all incident light, absorbing none. From this definition we can deduce that the BRDF $f(\theta_i, \phi_i; \theta_e, \phi_e)$ must be a constant for such a surface. What is this constant? To determine it, we integrate the radiance of the surface over all directions and equate the total radiance so obtained to the total irradiance:

$$\int_{-\pi}^{\pi} \int_0^{\pi/2} f\, E \cos\theta_i \sin\theta_e \cos\theta_e \, d\theta_e \, d\phi_e = E \cos\theta_i,$$

or

$$2\pi f \int_0^{\pi/2} \sin\theta_e \cos\theta_e \, d\theta_e = 1.$$

Using the identity $2\sin\theta\cos\theta = \sin 2\theta$, we obtain $\pi f = 1$. Thus, for an ideal Lambertian surface,

$$f(\theta_i, \phi_i; \theta_e, \phi_e) = \frac{1}{\pi}.$$

Note that since the BRDF is constant for a Lambertian surface, we can compute the radiance L from the irradiance E_0 using

$$L = \frac{1}{\pi} E_0.$$

This simple method cannot, of course, be used for surfaces with other reflectance properties.

The other extreme of surface reflectance properties is illustrated by an ideal specular reflector, which reflects all of the light arriving from the direction (θ_i, ϕ_i) into the direction $(\theta_i, \phi_i + \pi)$ (figure 10-9). The BRDF in this case is proportional to the product of two impulses, $\delta(\theta_e - \theta_i)$ and $\delta(\phi_e - \phi_i - \pi)$. But what is the factor of proportionality k? Once again we integrate over all emittance directions to compute the total radiance of the surface and equate this to the irradiance:

$$\int_{-\pi}^{\pi} \int_0^{\pi/2} k\, \delta(\theta_e - \theta_i)\, \delta(\phi_e - \phi_i - \pi) \sin\theta_e \cos\theta_e \, d\theta_e \, d\phi_e = 1,$$

or

$$k \sin\theta_i \cos\theta_i = 1.$$

Thus, in this case,

$$f(\theta_i, \phi_i; \theta_e, \phi_e) = \frac{\delta(\theta_e - \theta_i)\delta(\phi_e - \phi_i - \pi)}{\sin\theta_i \cos\theta_i}.$$

We can use this result immediately to determine the radiance of a specularly reflecting surface under an extended source:

$$\begin{aligned} L(\theta_e, \phi_e) &= \int_{-\pi}^{\pi} \int_0^{\pi/2} \frac{\delta(\theta_e - \theta_i)\delta(\phi_e - \phi_i - \pi)}{\sin\theta_i \cos\theta_i} E(\theta_i, \phi_i) \sin\theta_i \cos\theta_i \, d\theta_i \, d\phi_i \\ &= E(\theta_e, \phi_e - \pi). \end{aligned}$$

The radiance is just equal to the radiance of the reflected piece of the extended source. This makes eminent sense, since we are looking at the virtual image of the extended source. This simple relationship obviously does not hold for surfaces with other reflectance properties.

The BRDF can be determined experimentally by illuminating a flat sample of the material of interest with a lamp mounted on a goniometer and measuring its irradiance using a sensor mounted on another goniometer. (A *goniometer* has two axes of rotation, so that a device mounted on it can be aimed in an accurately known direction.) The experimental determination of the BRDF is quite tedious because of the four variables involved. Fortunately, only three—θ_i, θ_e, and $(\phi_e - \phi_i)$—are typically really significant.

Another way to obtain the BRDF is to model how light is reflected from a surface and to find the corresponding reflectance properties analytically or by numerical simulation. This has been done for some simple models of surface microstructure. Closed-form solutions are often possible if suitable approximations are introduced. We shall not pursue this topic further here.

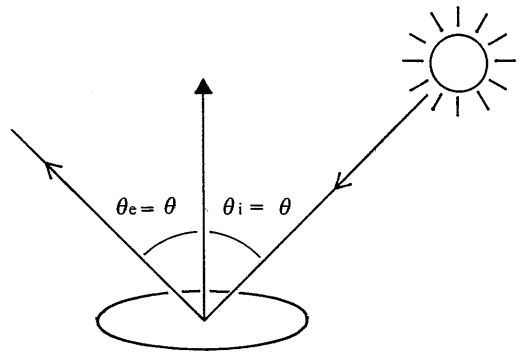

Figure 10-9. A specular surface reflects all incident light in a direction that lies in the same plane as the incident ray and the surface normal. The emittance angle θ_e between the reflected ray and the normal equals the incident angle θ_i between the incident ray and the normal.

10.7 Surface Brightness

How bright will a Lambertian surface be when it is illuminated by a point source of radiance E? A point source located in direction (θ_s, ϕ_s) has radiance

$$E(\theta_i, \phi_i) = E\frac{\delta(\theta_i - \theta_s)\,\delta(\phi_i - \phi_s)}{\sin\theta_s},$$

where the $\sin\theta_i$ term ensures that the integral of this expression is just E. That is, we must have

$$\int_{-\pi}^{\pi}\int_{0}^{\pi/2} E(\theta_i, \phi_i)\sin\theta_i\, d\theta_i\, d\phi_i = E.$$

Using the known BRDF for a Lambertian surface, it is easy to show that in this case

$$L = \frac{1}{\pi}E\cos\theta_i \qquad \text{for } \theta_i \geq 0$$

This is the familiar "cosine" or Lambert's law of reflection from matte surfaces. (Note that the dependence on the cosine of the incident angle comes directly from the dependence of the irradiance on that factor and so can be traced to the foreshortening of the surface as seen from the light source.) Surfaces covered with finely powdered transparent materials, such as barium sulfate or magnesium carbonate, come closest to obeying Lambert's law. It is a reasonable approximation for many other materials, such as paper, snow, and matte paint.

Next, consider the same surface under a "sky" of uniform radiance E. Here

$$L = \int_{-\pi}^{\pi}\int_{0}^{\pi/2} \frac{E}{\pi}\sin\theta_i\cos\theta_i\, d\theta_i\, d\phi_i = E.$$

The radiance of the patch is the same as the radiance of the source!

This leads to an interesting thought experiment. If we were to build a bottle of arbitrary shape, coated the interior with Lambertian material, and then introduced some light through a tiny hole, every surface patch would be equally bright. Peering through another tiny hole, we would not be able to discern the shape of the inner surface, since every portion of it would look equally bright. With an overcast sky, snow fields tend to have very low contrast and make vision difficult for the same reason. This is called a *white-out* condition. The fact that the surface is white is, of course, not the problem, since the shape of the surface is easy to discern under point-source illumination. Recall that a specular surface under a uniform extended source also appears to have the same radiance as the extended source. So Lambertian reflection is not required for this effect to occur.

A more complicated case, but more realistic perhaps, is a hemispherical sky above a tilted Lambertian surface patch. We show in exercise 10-3 that the total radiance in this case is

$$\frac{E}{2}(1 + \cos\alpha) = E\cos^2\frac{\alpha}{2},$$

where α is the inclination of the surface normal with respect to the vertical direction. The variation in brightness due to surface orientation changes is less here than it would be if there were a single point source overhead.

10.8 Surface Orientation

The BRDF is of fundamental importance for understanding reflectance from a surface. It is not exactly what we need, however, to understand image formation. First of all, to factor in the distribution of light sources we need to integrate the BRDF over all possible directions of the incident light. This yields a function that depends on two parameters only. We can relate these two parameters to surface orientation, a very important aspect of the surfaces being imaged. To do this successfully, however, we have to abandon the local coordinate system used in the definition of the BRDF and use a viewer-centered coordinate system instead.

Let us start by developing a reasonable notation for surface orientation. A smooth surface has a tangent plane at every point. The orientation of this tangent plane will be taken to represent the orientation of the surface at that point. The surface normal, a unit vector perpendicular to the

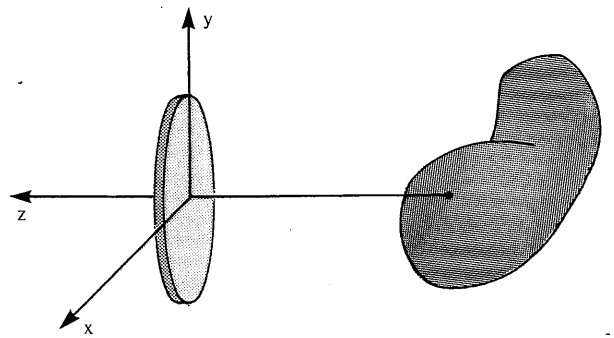

Figure 10-10. A surface can be conveniently described in terms of its perpendicular distance $-z(x, y)$ from some reference plane parallel to the image plane.

tangent plane, is appropriate for specifying the orientation of this plane. The normal vector has two degrees of freedom, since it is a vector with three components and one constraint—that the sum of squares of the components must equal one.

Alternatively, we can imagine placing this vector with its tail at the center of a unit sphere, called the *Gaussian sphere*. The head of the vector touches the sphere at a particular point, which we can use to denote surface orientation. The position of this point on the sphere can be specified by two variables, polar angle and azimuth, say, or latitude and longitude.

We must fix the coordinate system relative to which these measurements are to be made. It is convenient to choose this system so that one axis is lined up with the optical axis of the imaging system. We can place the origin of the system at the center of the lens, with two axes parallel to the image plane. Since we would like to have a right-handed coordinate system, we choose to have the z-axis point toward the image.

A portion of a surface can now be described by its perpendicular distance $-z$ from the lens plane (or any reference plane parallel to it). This distance will depend on the lateral displacement (x, y) (figure 10-10). What we would like to do next is to write the surface normal in terms of z and the partial derivatives of z with respect to x and y.

The surface normal is perpendicular to all lines in the tangent plane of the surface. As a result, it can be found by taking the cross-product of any two (nonparallel) lines in the tangent plane (figure 10-11). Consider taking a small step δx in the x-direction starting from a given point (x, y).

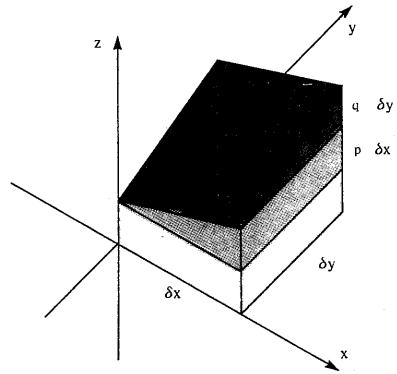

Figure 10-11. Surface orientation can be parameterized by the first partial derivatives p and q of the surface height z.

The change in z can be shown (using Taylor series expansion) to be

$$\delta z = \frac{\partial z}{\partial x}\delta x + e,$$

where e contains higher-order terms. We use the abbreviations p and q for the first partial derivatives of z with respect to x and y, respectively. Thus p is the slope of the surface measured in the x-direction, while q is the slope in the y-direction. The relationship of p and q to the orientation of the surface patch is shown in figure 10-12. If we take a small step of length δx in the x-direction, the height changes by $p\,\delta x$. Similarly, a small step of length δy in the y-direction leads to a change in height of $q\,\delta y$.

We can write the first small step in vector form as $(\delta x, 0, p\,\delta x)^T$. Thus a line parallel to the vector

$$\mathbf{r}_x = (1, 0, p)^T$$

lies in the tangent plane at (x, y). Similarly, a line parallel to

$$\mathbf{r}_y = (0, 1, q)^T$$

lies in the tangent plane also. A surface normal can be found by taking the cross-product of these two lines. It remains for us to decide whether we want the normal to point toward or away from the viewer. If we let it point toward the viewer, we obtain

$$\mathbf{n} = \mathbf{r}_x \times \mathbf{r}_y = (-p, -q, 1)^T.$$

Appropriately enough, (p, q) is called the *gradient* of the surface, since its components, p and q, are the slopes of the surface in the x- and y-directions,

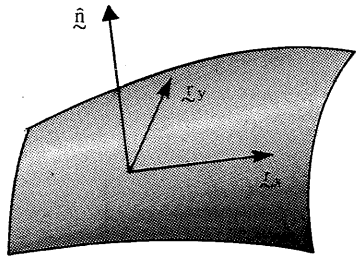

Figure 10-12. The normal is parallel to the cross-product of any two distinct tangents.

respectively.

The unit surface normal is just

$$\hat{\mathbf{n}} = \frac{\mathbf{n}}{|\mathbf{n}|} = \frac{(-p, -q, 1)^T}{\sqrt{1+p^2+q^2}}.$$

We can immediately calculate the angle θ_e between the surface normal and the direction to the lens, provided that the point considered is close to the optical axis relative to the distance from the reference plane. In this case the unit view vector $\hat{\mathbf{v}}$ from the object to the lens is $(0,0,1)^T$, so that

$$\cos\theta_e = \frac{1}{\sqrt{1+p^2+q^2}},$$

a result obtained by taking the dot-product of the two unit vectors.

How do we specify where the light sources are located? Assuming that they are far away from the object, relative to the size of the object, we can specify the direction to each one by a fixed vector. There exists a surface orientation that corresponds to this vector, that is, a surface oriented perpendicularly to the rays arriving from the source. If a normal to this surface is $(-p_s, -q_s, 1)^T$, then the gradient (p_s, q_s) can be used to specify the direction of the source (provided it lies on the same side of the object as the viewer).

From now on, we shall assume that the viewer and the light sources are far from the objects being imaged.

10.9 The Reflectance Map

The reflectance map makes explicit the relationship between surface orientation and brightness. It encodes information about surface reflectance

properties and light-source distributions. It is a representational tool used in developing methods for recovering surface shape from images.

Consider a source of radiance E illuminating a Lambertian surface. The scene radiance is

$$L = \frac{1}{\pi} E \cos\theta_i \qquad \text{for } \theta_i \geq 0,$$

where θ_i is the angle between the surface normal and the direction toward the source. Taking dot-products of the corresponding unit vectors, we obtain

$$\cos\theta_i = \frac{1 + p_s p + q_s q}{\sqrt{1+p^2+q^2}\sqrt{1+p_s^2+q_s^2}}.$$

This gives us a good idea of how brightness depends on surface orientation. The result is called the *reflectance map*, denoted $R(p,q)$. The reflectance map depends on the properties of the surface material of the object and the distribution of light sources. (Note that radiance cannot be negative, so we should, strictly speaking, impose the restriction $0 \leq \theta_i \leq \pi/2$. The radiance will be zero for values of θ_i outside this range.)

Image irradiance is proportional to a number of constants, such as the inverse of the square of the f-number and the fixed brightness of the source. For this reason, the reflectance map is usually normalized in some way, for example, so that its maximum is one. For the Lambertian surface illuminated by a single distant point source we can use

$$R(p,q) = \frac{1 + p_s p + q_s q}{\sqrt{1+p^2+q^2}\sqrt{1+p_s^2+q_s^2}}.$$

Thus, aside from a fixed scale factor, the reflectance map gives the dependence of scene radiance on surface orientation.

It is often convenient to plot the surface $R(p,q)$ as a function of the gradient (p,q). The pq-plane is called *gradient space*, and every point in it corresponds to a particular surface orientation. The point at the origin, for example, respresents the orientation of all planes that are perpendicular to the viewing direction. A contour map in gradient space can be used to depict a reflectance map (figure 10-13). In the case of Lambertian surface material, contours of constant brightness are nested conic sections in the pq-plane, since $R(p,q) = c$ implies

$$(1 + p_s p + q_s q)^2 = c^2(1+p^2+q^2)(1+p_s^2+q_s^2).$$

The maximum of $R(p,q)$ is at $(p,q) = (p_s, q_s)$, as we show in exercise 10-11.

As another example, consider a surface that emits radiation equally in all directions. (This is actually not physically plausible, but we shall describe a feasible modification later.) Such a surface appears brighter

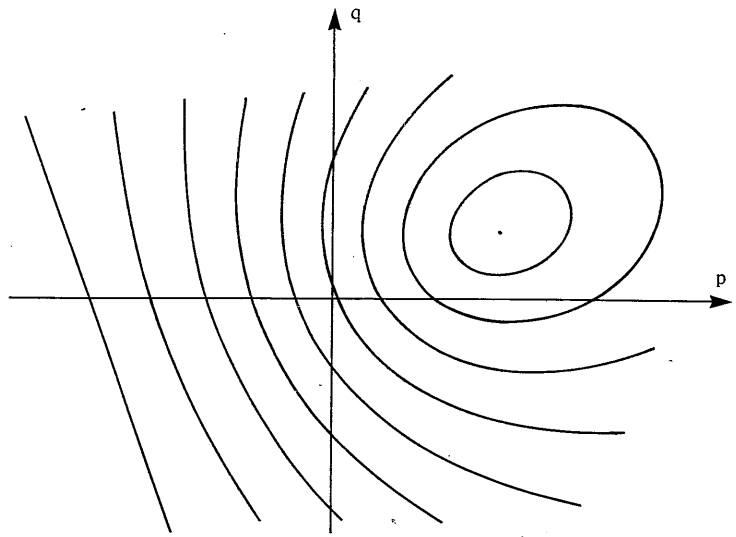

Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p,q)$ occurs at the point $(p,q) = (p_s, q_s)$, found inside the nested conic sections, while $R(p,q) = 0$ all along the line on the left side of the contour map.

when viewed obliquely, since the same power comes from a foreshortened area. This sort of behavior is clearly different from that of a Lambertian surface. Brightness in this case depends on the inverse of the cosine of the emittance angle. Taking into account the foreshortened area as seen from the source, we find that the radiance is proportional to $\cos\theta_i/\cos\theta_e$. Since $\cos\theta_e = 1/\sqrt{1+p^2+q^2}$, we have

$$R(p,q) = \frac{1 + p_s p + q_s q}{\sqrt{1+p_s^2+q_s^2}}.$$

The contours of constant brightness are parallel straight lines (figure 10-14), since $R(p,q) = c$ implies

$$1 + p_s p + q_s q = c\sqrt{1+p_s^2+q_s^2}.$$

These lines are orthogonal to the direction (p_s, q_s).

It turns out that no real surface can have radiance proportional to $\cos\theta_i/\cos\theta_e$, for this expression can be shown to violate the basic constraint

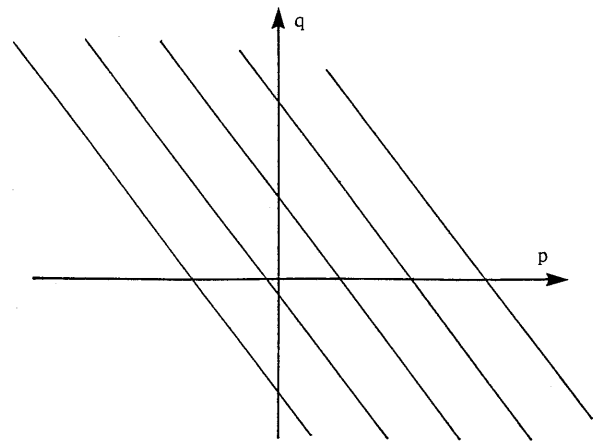

Figure 10-14. In the case of the material in the maria of the moon, the reflectance map can be closely approximated by a function of a linear combination of the components of the gradient. The contours of constant brightness are parallel straight lines in gradient space.

discovered by Helmholtz. The square root of this expression, however, does obey Helmholtz reciprocity, as shown by Minnaert. We are interested here in the shape of the contours in gradient space, and these are straight lines for any function of $\cos\theta_i/\cos\theta_e$, including, of course, its square root. Curiously, the material in the maria of the moon has reflectance properties that can be modeled reasonably well by a function of $\cos\theta_i/\cos\theta_e$.

As a final example, consider a glossy surface. The light reflected from many surfaces has two components, one due to reflection at the interface between air and the material of the surface, the other due to internal scattering of light that has penetrated into the surface layer. If the outer surface is perfectly smooth, the first component of reflection will be essentially mirrorlike or specular. The second component will be diffuse or matte. If the surface is not perfectly smooth, the specular component will tend to be smeared out, so that a point source will give rise to a high peak in the reflectance map, rather than an ideal impulse. This is called *glossy* reflection.

Absorption by the particles in the surface layer will change the brightness of the matte component, while the brightness of the glossy component depends primarily on the refractive index of the material in which these particles are embedded. As a result, the spectral composition of the glossy component is usually fairly close to that of the incident light, while the diffuse component is strongly affected by selective absorption in the surface

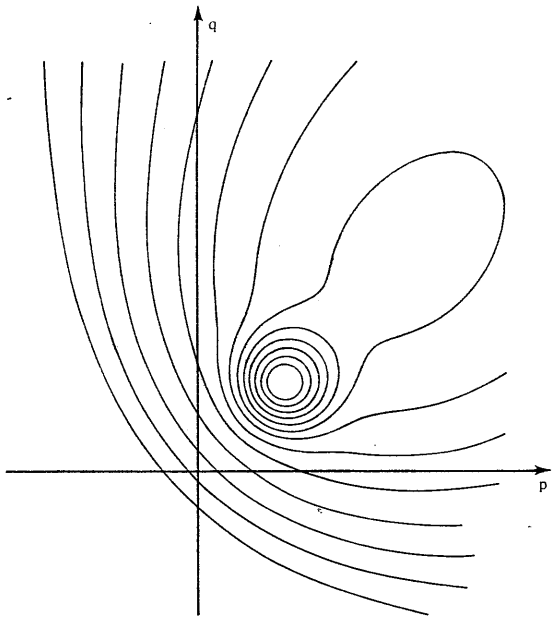

Figure 10-15. Many real surfaces combine diffuse and glossy components of reflection. The glossy component comes from reflection at the interface between air and the surface material, while the diffuse component can be traced to the component of light that has penetrated some distance into the surface, to be refracted and reflected until it reemerges from the surface layer. The reflectance map for such a material illuminated by a point source can have two peaks, corresponding to the surface orientations that maximize each of the two different types of reflection.

layer.

The surface orientation that maximizes the diffuse reflection component is typically one for which the surface normal points at the light source. The surface orientation that maximizes the glossy component, on the other hand, is usually one for which the surface normal points about halfway between the light source and the viewer. Correspondingly, the reflectance map can have two maxima (figure 10-15). Typically, the global maximum is at the glossy peak.

10.10 Shading in Images

How is the brightness pattern recorded in the image affected by the shape of an object? As an example, consider a polyhedron. Ideally, the image

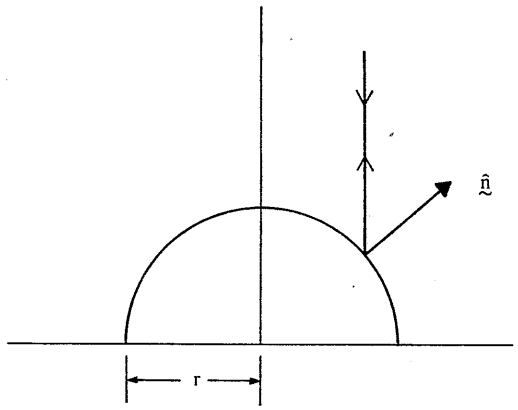

Figure 10-16. The reflectance map is circularly symmetric when the source lies in the same direction as the viewer. In the case of a Lambertian sphere, we obtain a smooth falloff in brightness toward the edge of the circular image of the sphere. This shading effect allows us to estimate the shape of the object.

of a polyhedron will consist of polygonal regions of uniform brightness, since all points on a face of a polyhedron have the same orientation. The brightness of a region in the image depends on the orientation of the corresponding face of the polyhedron. Now consider instead a smoothly curved object. The image of such an object will have spatial variations in brightness due to the fact that surface patches with different orientations appear with different brightness. This variation of brightness is called *shading*. Unfortunately, it depends on more than just the shape of an object. The reflectance properties of the surface and the distribution of light sources are also important.

The reflectance map captures the dependence of brightness on surface orientation. At a particular point in the image we measure the image irradiance $E(x, y)$. It is proportional to the radiance at the corresponding point on the surface imaged, as determined by the projection equation. If the surface gradient at that point is (p, q), then the radiance there is $R(p, q)$. If we normalize by setting the constant of proportionality to one, we obtain

$$E(x, y) = R(p, q).$$

This *image irradiance equation* is fundamental to the methods for recovering surface shape discussed in this and the next chapter.

As an illustration of the shading effect, consider a sphere with a Lambertian surface illuminated by a point source at essentially the same place

as the viewer (figure 10-16). In this case $\theta_e = \theta_i$ and $(p_s, q_s) = (0,0)$, so that

$$R(p,q) = \frac{1}{\sqrt{1+p^2+q^2}}.$$

If the sphere is on the optical axis, we can write the equation for its surface as

$$z = z_0 + \sqrt{r^2 - (x^2+y^2)} \qquad \text{for } x^2+y^2 \le r^2,$$

where r is the radius and $-z_0$ is the distance of its center from the lens. Thus

$$p = -\frac{x}{z-z_0} \qquad \text{and} \qquad q = -\frac{y}{z-z_0},$$

so that

$$\frac{1}{\sqrt{1+p^2+q^2}} = \frac{z-z_0}{r}.$$

Finally, we have

$$E(x,y) = R(p,q) = \sqrt{1 - \frac{x^2+y^2}{r^2}}.$$

The brightness falls off smoothly from its maximum at the center of the image of the sphere to zero at the edges. It is this variation in brightness that allows us to conclude that this is an image of a round, probably spherical, object.

If the sphere had a surface with different reflectance properties, it would give rise to a different image. For example, if

$$R(p,q) = \frac{1+p_s p + q_s q}{\sqrt{1+p_s^2+q_s^2}}$$

and $(p_s, q_s) = (0,0)$, we would obtain a uniformly bright disk in the image. We can see this also by noting that brightness is a function of $\cos\theta_i / \cos\theta_e$ in this case, and $\theta_e = \theta_i$. Thus such a sphere, illuminated from a position near that of the viewer, appears flat to someone used to objects with surface materials that have reflectance properties similar to those of a Lambertian surface. This is why the full moon looks like a flat disk marked with blotches. The blotches are brightness variations due to differences in the surface's reflecting efficiency. Often reflectance properties can be described in terms of the product of two factors: a geometric term expressing the dependence on the angles of light reflection, and another term that is the fraction of the incident light reemitted by the surface. The latter factor is called the *albedo*. We shall discuss methods later for recovering both shape and albedo from images.

q \ p	−.35	−.25	−.15	−.05	+.05	+.15	+.25	+.35
+.35	.58	.71	.81	.86	.91	.98	.99	
+.25	.53	.70	.79	.85	.88	.97	.98	
+.15	.43	.51	.70	.80	.83	.87	.90	
+.05	.32	.42	.51	.71	.80	.84	.85	
−.05	.18	.29	.41	.50	.70	.78	.80	
−.15	.08	.13	.28	.40	.50	.66	.67	
−.25	.01	.06	.11	.26	.39	.49	.51	
−.35								

Figure 10-17. To produce shaded graphics, we can encode the reflectance map in a quantized lookup table.

10.11 Shaded Graphics

The reflectance map $R(p,q)$ comes in handy when we are portraying a shape graphically. To compute the gray-level at a point in the image of a given surface, $z(x,y)$, we find the point of intersection of the corresponding ray from the viewer with that surface. Then we find the gradient (p,q) at that point and use $R(p,q)$ to determine the appropriate brightness.

Reproduction media are limited in the number of gray-levels they can display, and humans can only resolve a relatively small number of brightness gradations. Consequently, we do not need to know $R(p,q)$ with extreme precision. This makes it possible to use a lookup table to estimate $R(p,q)$ based on quantized values of p and q (figure 10-17). A simple computation then allows us to determine the brightness for each point in the image. This information can be loaded into a hardware lookup memory between the frame buffer and the display. The *frame buffer* is a memory normally used to store image gray-levels. In the method discussed here, it stores encoded surface orientation instead. Changes in surface reflectance properties and light-source arrangements require nothing more complicated than changing the lookup table.

A *block diagram* obtained from a digital terrain model of the Dent de Morcles region in the southwestern part of Switzerland is shown in figure 10-18. A block diagram is the oblique projection of a set of profiles

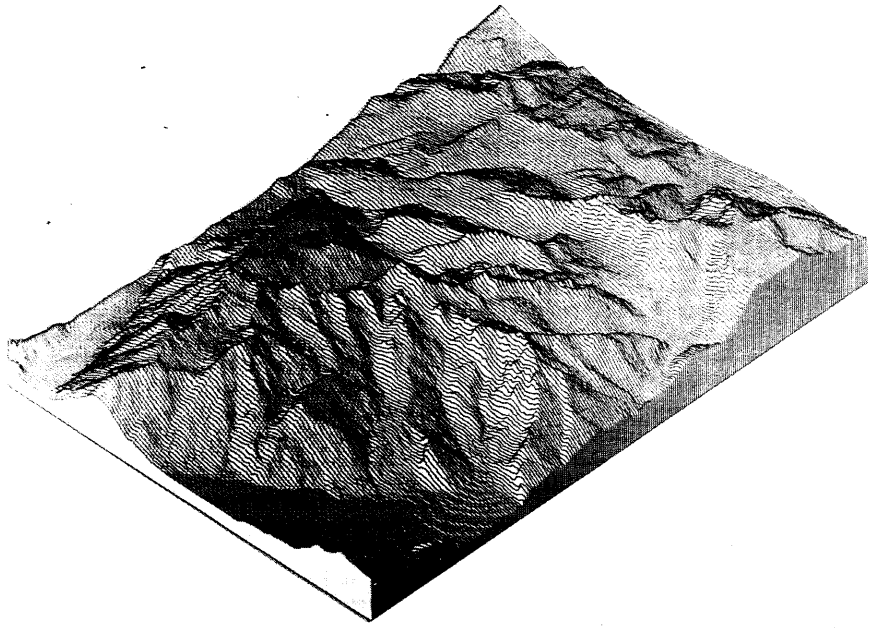

Figure 10-18. A block diagram made from a depth map of the surface of a mountainous region of the earth. (Digital terrain model kindly provided by Kurt Brassel.)

obtained by cutting the surface with a series of parallel vertical planes running from west to east.

Figure 10-19 shows two orthographic shaded views of the same surface, obtained using the shading methods described above, under two different assumed lighting conditions.

10.12 Photometric Stereo

The reflectance map is extremely useful in computer graphics, where an image is created from a description of the shape of an object. But we would like to go in the other direction: Given an image, we would like to be able to recover the shape. There is a unique mapping from surface orientation, specified by p and q, to radiance, given by the reflectance map $R(p, q)$. The inverse mapping is not unique. An infinite number of surface orientations give rise to the same brightness. A contour of constant $R(p, q)$ connects such a set of orientations in the reflectance map.

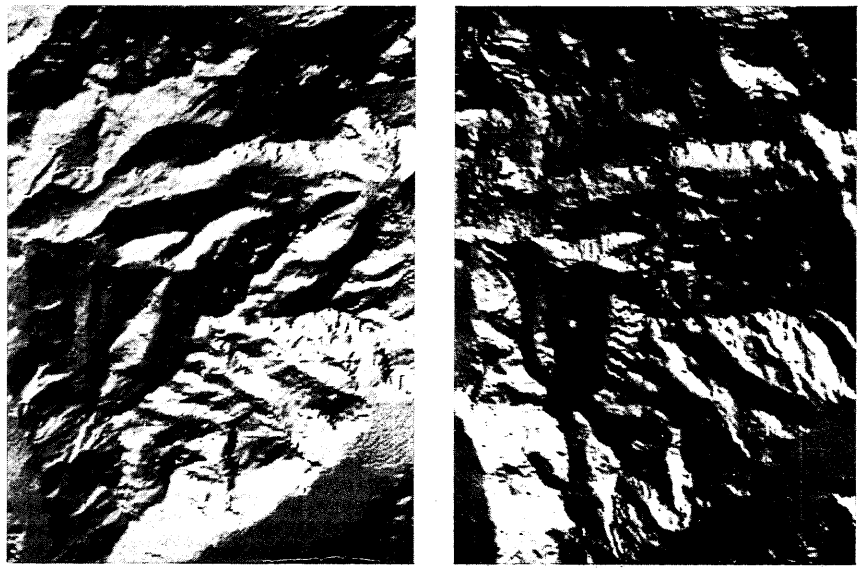

Figure 10-19. Shaded views of the surface shown in the previous figure. Here we are looking down on the surface from above. In the figure on the left, the light is assumed to come from the lower right, as is appropriate for a point on the earth north of the equator in the morning. In the figure on the right, the light is assumed to come from the lower left, corresponding to lighting conditions in the afternoon.

Surface orientation can usually be determined uniquely for some special points, such as those where the brightness is a maximum or minimum of $R(p, q)$. For a Lambertian surface, for example, $R(p, q) = 1$ only when $(p, q) = (p_s, q_s)$. In general, however, the mapping from brightness to surface orientation cannot be unique, since brightness only has one degree of freedom, while orientation has two.

To recover surface orientation locally, we must introduce additional information. To determine two unknowns, p and q, we need two equations. Two images, taken with different lighting, will yield two equations for each image point (figure 10-20):

$$R_1(p, q) = E_1 \qquad \text{and} \qquad R_2(p, q) = E_2.$$

If these equations are linear and independent, there will be a unique solution for p and q.

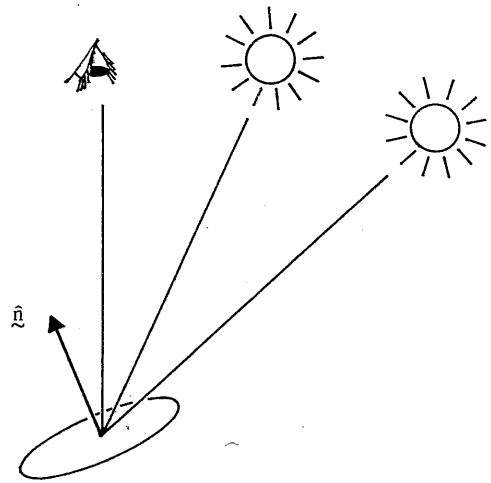

Figure 10-20. In photometric stereo, more than one image is taken from the same viewpoint with different lighting. The brightness of a patch of the surface will generally be different under the two lighting conditions.

Suppose, for example, that

$$R_1(p,q) = \sqrt{\frac{1+p_1p+q_1q}{r_1}} \quad \text{and} \quad R_2(p,q) = \sqrt{\frac{1+p_2p+q_2q}{r_2}},$$

where

$$r_1 = \sqrt{1+p_1^2+q_1^2} \quad \text{and} \quad r_2 = \sqrt{1+p_2^2+q_2^2}.$$

Then

$$p = \frac{(E_1^2r_1-1)q_2-(E_2^2r_2-1)q_1}{p_1q_2-q_1p_2}, \quad q = \frac{(E_2^2r_2-1)p_1-(E_1^2r_1-1)p_2}{p_1q_2-q_1p_2},$$

provided $p_1/q_1 \neq p_2/q_2$. Thus a unique solution can be obtained for surface orientation at each point, given two registered images taken with different lighting conditions. This is an illustration of the method of *photometric stereo.*

Incidentally, the condition $p_1/q_1 \neq p_2/q_2$ precludes the use of this approach on telescopic images of the maria of the moon obtained from earth, because the moon orbits the earth in a plane that is almost the same as the plane in which the earth orbits the sun. The light-source positions corresponding to various phases of the moon all lie along a straight line passing through the origin in gradient space. The form of the equation

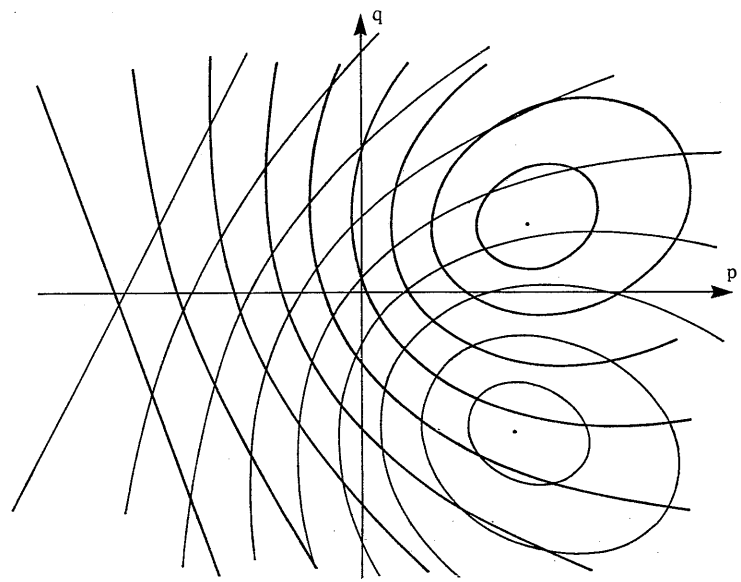

Figure 10-21. In the case of a Lambertian surface illuminated successively by two different point sources, there are at most two surface orientations that produce a particular pair of brightness values. These are found at the intersection of the corresponding contours in two superimposed reflectance maps.

above also suggests that the error in the gradient (p,q) due to the error in brightness (E_1, E_2) depends inversely on $p_1q_2 - q_1p_2$. For best results, the two light-source directions should be far apart in gradient space.

If the equations are nonlinear, we can have either no solutions or several solutions (figure 10-21). Suppose, for example, that

$$R_1(p,q) = \frac{1+p_1p+q_1q}{\sqrt{1+p^2+q^2}\sqrt{1+p_1^2+q_1^2}}$$

and

$$R_2(p,q) = \frac{1+p_2p+q_2q}{\sqrt{1+p^2+q^2}\sqrt{1+p_2^2+q_2^2}}.$$

We show in exercise 10-13 that there can be two solutions, one solution, or none, depending on the particular values of R_1 and R_2. (In pathological cases, such as when $(p_1,q_1) = (p_2,q_2)$, there can even be an infinite number of solutions.)

Often it is better to use three rather than two different illuminating conditions. In some cases this makes the equations linear. More important,

it can improve accuracy and increase the range of surface orientations over which a solution can be obtained. Finally, the third image can allow us to recover another parameter, albedo, as shown in the next section.

10.13 Recovering Albedo

Often a surface is not uniform in its reflectance properties. The easiest situation to deal with is one in which radiance is the product of a reflectance factor, or *albedo*, and some function of orientation only. Here we take albedo to be a number between zero and one that indicates how much light the surface reflects relative to some ideal surface with the same geometric dependence in the BRDF. Suppose, for example, that a matte surface behaves like a Lambertian surface, except that it is "gray," that is, it does not reflect all of the incident light. Its brightness is $\rho \cos\theta_i$, where ρ is the albedo, which can vary from place to place. To recover the albedo and the gradient (p, q), we need three pieces of information, which we can obtain from three image measurements.

As before, we can work out the solution directly in terms of the components of the normal vector, or we can use a more compact notation. If we introduce unit vectors in the directions of three source positions,

$$\hat{\mathbf{s}}_i = \frac{(-p_i, -q_i, 1)^T}{\sqrt{1 + p_i^2 + q_i^2}} \qquad \text{for } i = 1, 2, 3,$$

then

$$E_i = \rho(\hat{\mathbf{s}}_i \cdot \hat{\mathbf{n}}) \qquad \text{for } i = 1, 2, 3,$$

where

$$\hat{\mathbf{n}} = \frac{(-p, -q, 1)^T}{\sqrt{1 + p^2 + q^2}}$$

is the unit surface normal. We thus have three equations

$$E_1 = \rho(\hat{\mathbf{s}}_1 \cdot \hat{\mathbf{n}}), \quad E_2 = \rho(\hat{\mathbf{s}}_2 \cdot \hat{\mathbf{n}}), \quad E_3 = \rho(\hat{\mathbf{s}}_3 \cdot \hat{\mathbf{n}})$$

for the unit vector $\hat{\mathbf{n}}$ and the albedo ρ. We can combine these equations to obtain $\mathbf{E} = \rho \mathbf{S} \hat{\mathbf{n}}$, where the rows of the matrix $\mathbf{S}$ are the source directions $\hat{\mathbf{s}}_1$, $\hat{\mathbf{s}}_2$, and $\hat{\mathbf{s}}_3$, and the components of the vector $\mathbf{E}$ are the three brightness measurements. Thus $\rho \hat{\mathbf{n}} = \mathbf{S}^{-1}\mathbf{E}$, provided that the matrix $\mathbf{S}$ is not singular. We show in exercise 10-16 that in fact

$$\rho \hat{\mathbf{n}} = \frac{1}{[\hat{\mathbf{s}}_1\, \hat{\mathbf{s}}_2\, \hat{\mathbf{s}}_3]} \big(E_1(\hat{\mathbf{s}}_2 \times \hat{\mathbf{s}}_3) + E_2(\hat{\mathbf{s}}_3 \times \hat{\mathbf{s}}_1) + E_3(\hat{\mathbf{s}}_1 \times \hat{\mathbf{s}}_2)\big),$$

where $[\hat{\mathbf{s}}_1\, \hat{\mathbf{s}}_2\, \hat{\mathbf{s}}_3]$ is the triple product $\hat{\mathbf{s}}_1 \cdot (\hat{\mathbf{s}}_2 \times \hat{\mathbf{s}}_3)$. The direction of the surface normal is therefore a constant times a linear combination of three

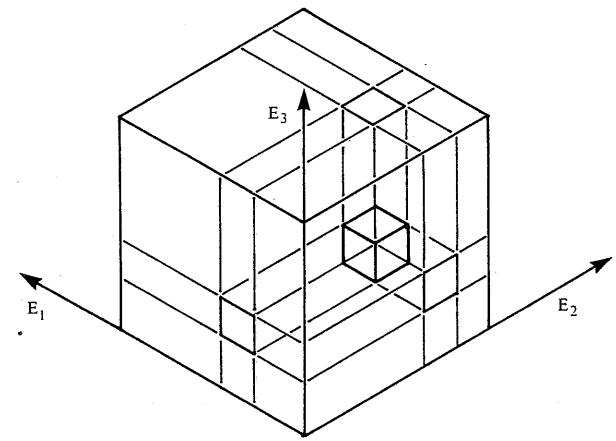

Figure 10-22. Photometric stereo is a fast and robust method for recovering surface shape. The calculation can be encoded in a lookup table based on quantized values of the brightness measurements.

vectors, each of which is perpendicular to two of the light-source directions. Each of these vectors is multiplied by the brightness observed when the third source is used. The albedo is recovered by finding the length of the resulting vector. The computation is straightforward in this case, and a unique result is assured.

10.14 Lookup Tables for Surface Orientation

The reflectance map is usually determined experimentally, in which case it will not be possible to find a closed-form solution to the photometric stereo problem. Moreover, even if a closed-form solution can be found, it usually involves a great deal of computation. Lookup tables indexed on the observed brightness measurements offer a solution to this impasse (figure 10-22). They provide rapid results and adequate accuracy, given that brightness can only be measured with limited accuracy anyway.

The lookup table can be constructed from a reflectance map given as an explicit formula, or it can be filled in experimentally. A calibration object of known shape, such as a sphere, is imaged. The gray-levels obtained at a particular point under various lighting conditions are used to determine an entry in the table. The surface orientation is computed using the known shape of the object and is entered in the appropriate place in the table. Gaps that occur in the table when the sampling of the calibration object does not touch a point with the corresponding orientation can be filled in

by interpolation.

When the photometric stereo technique is applied to an image, a surface orientation estimate is obtained for every picture cell. The result is called a *needle diagram.* It can be displayed graphically by showing a top view of the surface covered with a regular grid of short normal vectors. Each "needle" is the projection of a surface normal vector erected on the patch of the surface corresponding to a particular picture cell. The direction of the needle in the picture is the direction of steepest descent on the surface, and the length is the foreshortened length, proportional to the sine of the angle between the normal and the viewing direction. A needle diagram encodes surface shape and can be used to recognize an object and recover its orientation in space. An example of a needle map is shown in chapter 18, where we discuss picking parts out of a bin. Needle diagrams are used there to provide information that will allow us to determine the attitude in space of an object of known shape.

There are, of course, other ways to encode surface shape, as, for example, in a depth map. A needle diagram can be obtained trivially from a depth map by taking first differences to estimate the gradient, but the reverse problem is overdetermined and requires a least-squares method. We postpone discussion of this problem to the last part of the next chapter, after we introduce the calculus of variations.

10.15 References

A book by Moon & Spencer, *The Photic Field* [1981], discusses radiometry and methods for calculating light flux. It is based in part on Moon's earlier book *The Scientific Basis of Illumination Engineering* [1961], which deals mostly with the flow of light flux in buildings.

The Gaussian sphere and the stereographic projection are described in the book *Geometry and the Imagination* by Hilbert & Cohn-Vossen [1952]. Do Carmo also discusses the Gauss map and Gaussian curvature in *Differential Geometry of Curves and Surfaces* [1976]. Shafer [1985] uses gradient space to analyze shadows and occluding contours in his book *Shadows and Silhouettes in Computer Vision.*

The terms used in radiometry have changed over the years, becoming standardized only relatively recently. This was a problem in particular in the case of the notion of "reflectance." Fortunately, Nicodemus et al. [1977] at the National Bureau of Standards have resolved these issues.

Horn [1977] introduced the reflectance map and showed several of its applications to machine vision. Later, Horn & Sjoberg [1979] related the reflectance map to the bidirectional reflectance distribution function of Nicodemus et al.. Attempts to automate hill shading have a long history; see, for example, Brassel [1974]. Horn [1981] explained these attempts in terms of the unifying notation provided by the reflectance map. Horn & Bachman [1978] registered satellite images with digital elevation models using similar techniques.

A large video lookup table to store the quantized reflectance map is described by Sloan & Brown [1979] and Bass [1981]. Methods developed in computer graphics for portraying shape made use of shading, often based on heuristic approaches. In a few cases, however, fairly realistic physical models of surface structure were employed. See, for example, Tuong-Phong [1975]. Better models for surface reflection are dealt with in the references mentioned at the end of the next chapter.

The photometric stereo method was developed by Woodham [1978b, 1980] and analyzed by Horn, Woodham, & Silver [1978]. Silver [1980] developed ways to apply the basic method to surfaces whose reflectance properties depend on several parameters, while Ikeuchi [1981b] used extended light sources to deal with surfaces that have specular or glossy reflectance properties. Since then, others have used and analyzed the method; see, for example, Coleman & Jain [1982], Ray, Birk, & Kelley [1983], and Horn & Ikeuchi [1984].

Use of the needle diagram for describing surface topography was reported by Horn in 1979. Application of photometric stereo to the bin picking problem is described in chapter 18.

10.16 Exercises

10-1 In this chapter we determined the cosines of the emittance angle θ_e and the incident angle θ_i in terms of the components of the gradient, p and q. In order to calculate an angle accurately we must know both its cosine and its sine. Show that

$$\sin\theta_e = \sqrt{\frac{p^2+q^2}{1+p^2+q^2}}, \quad \sin\theta_i = \sqrt{\frac{(p-p_s)^2+(q-q_s)^2+(q_s p - p_s q)^2}{(1+p^2+q^2)(1+p_s^2+q_s^2)}}.$$

10-2 Here we consider some apparent paradoxes relating to image and scene brightness.

(a) Why is the irradiance of the image of a surface independent of the distance from the surface? After all, when the lens is twice as far from the surface, it collects only one-quarter as much of the light emitted from a given surface patch.

(b) Show that the radiance of the image of a surface in a perfectly specular mirror is equal to the radiance of the surface itself, independent of the shape of the mirror.

(c) The term *intensity*, frequently misused, refers to the power emitted per unit solid angle by a light source ($\mathrm{W{\cdot}sr^{-1}}$—watts per steradian). Show that the intensity of the virtual image of a point source formed by a convex mirror is

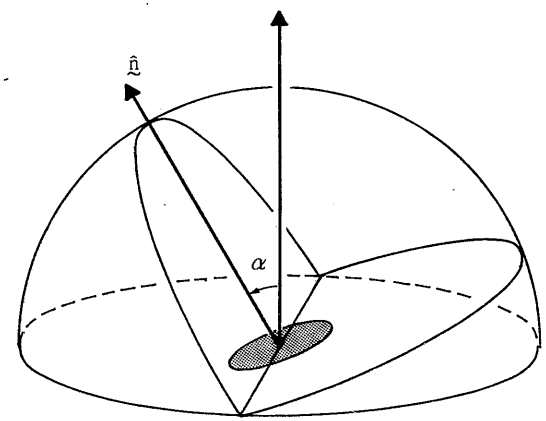

Figure 10-23. A surface patch inclined to the horizontal will receive radiation only from a portion of a hemispherical sky.

less than the intensity of the point source itself. How does the intensity of the virtual image vary with the curvature?

(d) Reconcile the apparently contradictory conclusions of parts (b) and (c).

10-3 Consider a hemispherical sky above a tilted Lambertian surface patch (figure 10-23). Suppose that the surface normal makes an angle $\alpha < \pi/2$ with respect to the zenith. The half of the hemisphere above the surface is unoccluded and thus gives rise to a component of radiance equal to $E/2$, using the result obtained previously for a Lambertian surface. The other half has a dark sector. Only a section from the zenith up to polar angle θ' is visible.

(a) Show that

$$\sin\theta' \sin\alpha \cos\phi = \cos\theta' \cos\alpha,$$

by considering the relevant spherical triangle. Conclude that

$$\tan\left(\frac{\pi}{2} - \theta'\right) = \tan\alpha \cos\phi.$$

Hint: See the appendix for useful formulae from spherical trigonometry.

(b) The contribution from the bright sector of this hemisphere is

$$\int_{-\pi/2}^{\pi/2} \int_0^{\theta'} \frac{E}{\pi} \sin\theta \cos\theta \, d\theta \, d\phi.$$

Show that

$$\frac{E}{2\pi} \int_{-\pi/2}^{\pi/2} \sin^2\theta' \, d\phi = \frac{E}{2\pi} \int_{-\pi/2}^{\pi/2} \frac{1}{1 + \tan^2\alpha \cos^2\phi} \, d\phi.$$

Hint: $\int \sin\theta \cos\theta = (1/2)\sin^2\theta$.

(c) Show that the contribution to the radiance from the portion of the sky considered is

$$\int \frac{1}{1 + a^2\cos^2\phi} \, d\phi = \frac{1}{\sqrt{1+a^2}} \tan^{-1}\left(\frac{\tan\phi}{\sqrt{1+a^2}}\right),$$

or $(E/2)\cos\alpha$. Conclude that the total radiance is

$$\frac{E}{2}(1 + \cos\alpha) = E\cos^2\frac{\alpha}{2}.$$

10-4 Determine the reflectance map for a specular surface illuminated by a point light source. That is, find the radiance from the appropriate bidirectional reflectance distribution function and the light-source distribution. Hint: This may seem a bit tricky because the integrand contains the product of two impulse functions, but it can be dealt with easily using the sifting property of the unit impulse.

10-5 The bidirectional reflectance distribution function for a *Minnaert surface* is

$$f(\theta_i, \phi_i; \theta_e, \phi_e) = \frac{k+1}{2\pi}\left(\cos\theta_i \cos\theta_e\right)^{k-1},$$

where $0 \le k \le 1$. A collimated point source from the direction (θ_s, ϕ_s) with irradiance E_0 is given by

$$E_i(\theta_i, \phi_i) = E_0 \frac{\delta(\theta_i - \theta_s)\,\delta(\phi_i - \phi_s)}{\sin\theta_s}.$$

Show that the radiance of the surface is

$$L_r = E_0 \frac{k+1}{2\pi} \cos^{k-1}\theta_e \cos^k\theta_i,$$

provided that $-\pi/2 \le \theta_i \le \pi/2$. For what values of k is the radiance independent of viewing direction?

10-6 The *gnomonic projection* can be used to map the Gaussian sphere, centered at the origin, onto the plane tangent at the north pole. This is done by projecting lines from the origin through points on the sphere onto the plane $z = 1$. The point $\hat{\mathbf{n}}$ on the Gaussian sphere maps into the point

$$\frac{\hat{\mathbf{n}}}{\hat{\mathbf{n}} \cdot \hat{\mathbf{z}}}$$

in the plane. This plane is just the gradient space we have been using.

(a) The *stereographic projection* provides another way to map the Gaussian sphere onto a plane. Here, a line is constructed from the south pole of the sphere, through points on the sphere, onto the plane $z = 1$. What point in the $z = 1$ plane corresponds to the point $\hat{\mathbf{n}}$ on the Gaussian sphere? Hint: Figure 11.9 might be helpful in visualizing the relationships.

(b) What contour on the sphere corresponds to points on the occluding boundary of an object? Where in the stereographic plane does this contour map to? Assume that the direction from the scene to the viewer is parallel to the north-south axis of the sphere. Hint: On the occluding contour the viewing direction is tangent to the surface.

10-7 The reflectance map can be parameterized in various ways. In this chapter we have concentrated on using the gradient (p, q) as a means of specifying surface orientation. In some cases, the Gaussian sphere is more suitable for this purpose. Each point on the Gaussian sphere corresponds to a particular direction, from the center of the sphere to that point. The orientation of a surface patch can be specified by giving the direction of its surface normal. Thus a given surface orientation can be identified with a particular point on the Gaussian sphere. The reflectance map is merely a means of associating brightness with orientation.

(a) What are the contours of constant brightness on the Gaussian sphere in the case of a Lambertian surface illuminated by a point source? Hint: See figure 10-24a.

(b) Show that there are at most two surface orientations that give rise to a given pair of brightness values when the photometric stereo method is applied to a Lambertian surface. Assume that two different light sources are used. Hint: See figure 10-24b.

10-8 A perfectly specular surface illuminated by a single point source gives rise to a simple reflectance map. It can be represented on the Gaussian sphere by an impulse lying halfway between the direction toward the source and the direction toward the viewer.

(a) Now consider a glossy surface whose microstructure can be modeled in terms of a random variation in surface orientation about a local average. Suppose that there is a limit to the angular deviation of the normal from the local average normal. Describe the reflectance map of this surface on the Gaussian sphere.

(b) Next, consider a perfectly specular surface illuminated by a circularly symmetric source. Suppose that the source has a limited diameter. What is the reflectance map in this case? Hint: It is not the same as the one derived above, although somewhat similar.

(c) Compare the two reflectance maps. How would you distinguish between the two cases in practice?

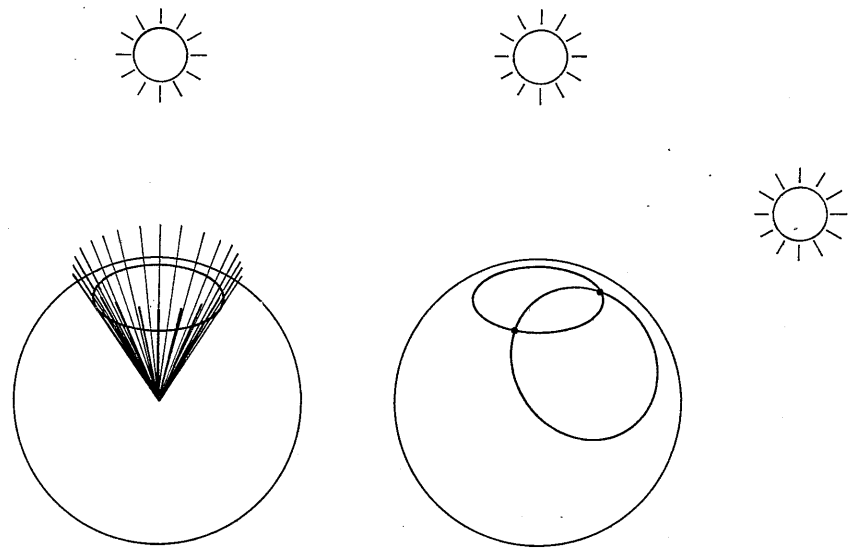

(a) (b)

Figure 10-24. The reflectance map can be plotted on the Gaussian sphere. (a) The contours of constant brightness are particularly simple for a Lambertian surface illuminated by a point source. (b) This makes it easy to prove that there are at most two solutions to the two-source photometric stereo problem in this case.

10-9 Show that in the case of Lambertian reflectance, the average brightness of a region R of a surface is the dot-product of the average surface normal $\overline{\mathbf{n}}$ and the unit vector $\hat{\mathbf{s}}$ in the direction of the light source. Here

$$A\overline{\mathbf{n}} = \iint_R \hat{\mathbf{n}}\, dA,$$

where A is the area of the region in the image. What happens when part of the region is self-shadowed? What does this tell us about the problem of recovering shape from shading? Hint: Note that $\overline{\mathbf{n}}$ is usually not a unit vector.

Suppose we were to model the microstructure of a surface using undulations on a scale too small to resolve. Suppose also that, considered on this small scale, the undulating surface behaved like a Lambertian reflector. Would the surface, considered on a scale at which the undulations cannot be resolved, behave like a Lambertian reflector? Treat separately the two limiting cases in which the incident rays and the direction to the viewer become nearly tangent to the surface.

10-10 Show that in the case of Lambertian reflectance, an extended light source can be replaced by an equivalent point source whose brightness and position can

be determined by integrating the brightness of the extended source over the unit sphere of incident directions,

$$\overline{E}\,\overline{\mathbf{s}} = \iint_S E(\hat{\mathbf{s}})\,\hat{\mathbf{s}}\,dS,$$

where $E(\hat{\mathbf{s}})$ is the radiance of the source in the direction $\hat{\mathbf{s}}$. What does this tell us about the problem of recovering shape from shading?

Apply this idea to the hemispherical sky example. Show that the equivalent light source is at an angle $\alpha/2$ from the nadir, and that its brightness is proportional to $\cos(\alpha/2)$. Hint: Some part of the extended light source cannot be seen from the surface patch.

10-11 A Lambertian surface illuminated by a point source has a reflectance map of the form

$$R(p,q) = \frac{1 + p_s p + q_s q}{\sqrt{1+p^2+q^2}\sqrt{1+p_s^2+q_s^2}},$$

where the light source lies in the direction $(-p_s, -q_s, 1)^T$.

(a) What value of the gradient (p,q) maximizes this expression? Is your maximum a global one? Is it unique?

(b) For what values of the gradient is $R(p,q) = 0$? What is the corresponding contour in gradient space?

10-12 Suppose that a surface has brightness $\sqrt{\cos\theta_i/\cos\theta_e}$, or

$$R(\hat{\mathbf{n}}) = \sqrt{\frac{\hat{\mathbf{n}}\cdot\hat{\mathbf{s}}}{\hat{\mathbf{n}}\cdot\hat{\mathbf{v}}}},$$

where $\hat{\mathbf{n}}$ is the unit surface normal, $\hat{\mathbf{s}}$ is the direction toward the light source, and $\hat{\mathbf{v}}$ is the direction toward the viewer. We measure the brightness at a particular point on the surface under two lighting conditions. We obtain brightness E_1 using light source $\hat{\mathbf{s}}_1$, and brightness E_2 using light source $\hat{\mathbf{s}}_2$. Show that

$$\hat{\mathbf{n}}\cdot\left(E_1^2\,\hat{\mathbf{v}} - \hat{\mathbf{s}}_1\right) = 0 \quad \text{and} \quad \hat{\mathbf{n}}\cdot\left(E_2^2\,\hat{\mathbf{v}} - \hat{\mathbf{s}}_2\right) = 0.$$

Conclude that $\hat{\mathbf{n}}$ is parallel to

$$\mathbf{n} = \pm\left(E_1^2(\hat{\mathbf{s}}_2\times\hat{\mathbf{v}}) - E_2^2(\hat{\mathbf{s}}_1\times\hat{\mathbf{v}}) + (\hat{\mathbf{s}}_1\times\hat{\mathbf{s}}_2)\right).$$

The sign is to be chosen to make $\hat{\mathbf{n}}\cdot\hat{\mathbf{z}} > 0$. Hint: You can check the result by taking dot-products with $\hat{\mathbf{s}}_1$, $\hat{\mathbf{s}}_2$, and $\hat{\mathbf{n}}$.

10-13 An ideal Lambertian surface leads to image irradiance

$$E = \frac{1 + p_s p + q_s q}{\sqrt{1+p^2+q^2}\sqrt{1+p_s^2+q_s^2}}$$

when a point source is located at $(-p_s, -q_s, 1)^T$. A single measurement confines the surface gradient to a conic section (hyperbola, parabola, or ellipse) in pq-space. In photometric stereo, we move the light source to obtain two measurements, E_1 and E_2. This confines the surface gradient to the intersection of two conic sections in gradient space. In general, two conic sections can intersect in as many as four distinct places. Prove that there are at most two solutions for p and q when E has the form given above. Hint: Show that all solutions must lie on a straight line.

10-14 For an ideal Lambertian surface under point-source illumination, the image irradiance is $\cos\theta_i$, which can be written in the form of a dot-product of unit vectors,

$$E = \hat{\mathbf{n}}\cdot\hat{\mathbf{s}},$$

where $\hat{\mathbf{n}}$ is a unit vector normal to the surface and $\hat{\mathbf{s}}$ is a unit vector in the direction of the light source. Suppose, as in the previous exercise, that two measurements, E_1 and E_2, are taken with two light-source positions, $\hat{\mathbf{s}}_1$ and $\hat{\mathbf{s}}_2$. Determine the two possible solutions for the unit surface normal. When is there only one solution? Hint: Use the methods for solving vector equations described in the appendix. Alternatively, use the methods for solving spherical triangles found there.

10-15 Suppose that three light sources, which can be turned on and off, are placed far from a Lambertian surface patch. The angles between the directions to the sources, as viewed from the surface patch, are 90° (figure 10-25). Show that the three values of observed brightness are proportional to the direction cosines of the normal vector of the surface patch expressed in a coordinate system defined by the directions to the three light sources. Conclude that

$$\rho\,\hat{\mathbf{n}} = \mathbf{R}\mathbf{e},$$

where ρ is the albedo, $\hat{\mathbf{n}}$ the unit surface normal, $\mathbf{e}$ the vector whose components are the three measurements, and $\mathbf{R}$ the rotation matrix corresponding to the transformation between the coordinate system defined by the three light source directions and the viewer coordinate system. Note: This is much simpler than the general case covered in this chapter.

10-16 We saw that in the case of a Lambertian surface and three light sources, the photometric stereo problem reduces to solving the three equations,

$$E_1 = \rho(\hat{\mathbf{s}}_1\cdot\hat{\mathbf{n}}), \quad E_2 = \rho(\hat{\mathbf{s}}_2\cdot\hat{\mathbf{n}}), \quad E_3 = \rho(\hat{\mathbf{s}}_3\cdot\hat{\mathbf{n}}),$$

for the unit vector $\hat{\mathbf{n}}$ and the albedo ρ.

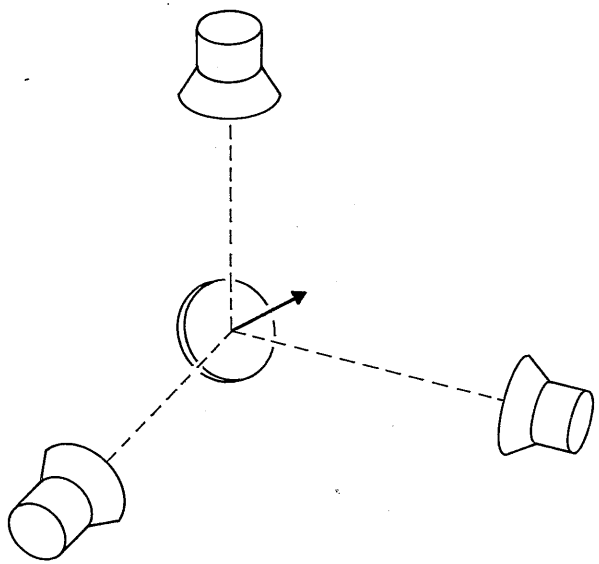

Figure 10-25. If three light sources in directions at right angles to one another are used to illuminate a Lambertian surface, the brightness measurements are proportional to the direction cosines of the surface normal, independent of the viewing direction.

(a) By subtracting equations pairwise, show that

$$\rho\,\hat{\mathbf{n}}\cdot(E_1\hat{\mathbf{s}}_2 - E_2\hat{\mathbf{s}}_1) = 0 \qquad \text{and} \qquad \rho\,\hat{\mathbf{n}}\cdot(E_2\hat{\mathbf{s}}_3 - E_3\hat{\mathbf{s}}_2) = 0.$$

(b) Now use the fact that $\hat{\mathbf{n}}$ is perpendicular to both $(E_1\hat{\mathbf{s}}_2 - E_2\hat{\mathbf{s}}_1)$ and $(E_2\hat{\mathbf{s}}_3 - E_3\hat{\mathbf{s}}_2)$ to show that it must be parallel to

$$(E_1\hat{\mathbf{s}}_2 - E_2\hat{\mathbf{s}}_1)\times(E_2\hat{\mathbf{s}}_3 - E_3\hat{\mathbf{s}}_2) = E_2\big(E_1(\hat{\mathbf{s}}_2\times\hat{\mathbf{s}}_3) + E_2(\hat{\mathbf{s}}_3\times\hat{\mathbf{s}}_1) + E_3(\hat{\mathbf{s}}_1\times\hat{\mathbf{s}}_2)\big).$$

(c) Conclude that

$$\rho\,\hat{\mathbf{n}} = k\big(E_1(\hat{\mathbf{s}}_2\times\hat{\mathbf{s}}_3) + E_2(\hat{\mathbf{s}}_3\times\hat{\mathbf{s}}_1) + E_3(\hat{\mathbf{s}}_1\times\hat{\mathbf{s}}_2)\big)$$

for some constant k.

(d) By taking the dot-product of this relationship with $\hat{\mathbf{s}}_1$, and remembering that $E_1 = \rho(\hat{\mathbf{n}}\cdot\hat{\mathbf{s}}_1)$, show that

$$\rho\,(\hat{\mathbf{n}}\cdot\hat{\mathbf{s}}_1) = k\,E_1[\hat{\mathbf{s}}_1\,\hat{\mathbf{s}}_2\,\hat{\mathbf{s}}_3],$$

where $[\hat{\mathbf{s}}_1\,\hat{\mathbf{s}}_2\,\hat{\mathbf{s}}_3]$ is the triple product $\hat{\mathbf{s}}_1\cdot(\hat{\mathbf{s}}_2\times\hat{\mathbf{s}}_3)$.

(e) Show that $k = 1/[\hat{\mathbf{s}}_1\,\hat{\mathbf{s}}_2\,\hat{\mathbf{s}}_3]$, so that

$$\rho\,\hat{\mathbf{n}} = \frac{\big(E_1(\hat{\mathbf{s}}_2\times\hat{\mathbf{s}}_3) + E_2(\hat{\mathbf{s}}_3\times\hat{\mathbf{s}}_1) + E_3(\hat{\mathbf{s}}_1\times\hat{\mathbf{s}}_2)\big)}{[\hat{\mathbf{s}}_1\,\hat{\mathbf{s}}_2\,\hat{\mathbf{s}}_3]}.$$

10-17 In this chapter we discussed a method for recovering albedo and orientation of a Lambertian surface using three light sources. This method works only for surface patches oriented so that they receive light from all three sources.

(a) Here we study the range of orientations that can be recovered successfully. Show that this range can be thought of as a spherical triangle on the unit sphere. Demonstrate that the corners of this triangle lie at $\hat{\mathbf{s}}_3\times\hat{\mathbf{s}}_1$, $\hat{\mathbf{s}}_2\times\hat{\mathbf{s}}_3$, and $\hat{\mathbf{s}}_1\times\hat{\mathbf{s}}_2$. Hint: A distant point source will illuminate surface patches with orientations lying in a hemisphere.

(b) What are the angles at the corners of this triangle? What is the solid angle of directions of the surface normal for which the photometric stereo method works? Hint: The area of a spherical triangle with angles α, β, and γ is equal to the radius squared times the *spherical excess*, the difference between the sum of the angles and π:

$$A = R^2(\alpha+\beta+\gamma-\pi).$$

(c) How can we maximize the range of orientations? Hint: The spherical triangle under consideration cannot be made larger then a hemisphere.

(d) There is an additional consideration. It is wasteful to have part of the spherical triangle lie in the hemisphere that corresponds to orientations of surface patches that are turned away from the viewer. Show that this will not happen if the view vector lies inside the spherical triangle formed by $\hat{\mathbf{s}}_1$, $\hat{\mathbf{s}}_2$, and $\hat{\mathbf{s}}_3$, that is, if

$$\hat{\mathbf{v}} = a\hat{\mathbf{s}}_1 + b\hat{\mathbf{s}}_2 + c\hat{\mathbf{s}}_3 \qquad \text{for some } a\geq 0,\ b\geq 0, \text{ and } c\geq 0.$$

We are not entirely free to place the light sources anywhere we want. We know, for example, that if two of them are in the same place, no new information is obtained when the second image is taken. Also, when two of them lie in nearly the same direction, there will be only a very small difference in the measurements they provide. It is likely that noise will have a bad influence on the result in this case.

(e) Consider small variations in the brightness measurements. How do they affect the resulting estimates of the surface normal? Hint: Differentiate the equation for $\rho\,\hat{\mathbf{n}}$ with respect to the vector $\mathbf{E} = (E_1, E_2, E_3)^T$.

(f) The standard deviation in the result equals the derivative so obtained times the standard deviation in the measurements. How can the sensitivity to noise be minimized? Hint: The triple product $[\hat{\mathbf{s}}_1\,\hat{\mathbf{s}}_2\,\hat{\mathbf{s}}_3]$ is the volume of the parallelepiped that can be built by extending the corner formed by the three unit vectors.

(g) Discuss the conflict between the source placements suggested by the above considerations. Consider ways to overcome these limitations using extended light sources.

10-18 We might want to take more than two brightness measurements in order to improve accuracy when using the photometric stereo method. Imagine that n light sources are used in turn to obtain n images. Suppose that the surface under consideration is Lambertian and that the direction to the ith source is given by the unit vector $\hat{\mathbf{s}}_i$. This time we assume that the surface can have an albedo ρ different from one. At each point in the image, we wish to find the unit surface normal $\hat{\mathbf{n}}$ that minimizes

$$\sum_{i=1}^{n} (\rho\,\hat{\mathbf{n}} \cdot \hat{\mathbf{s}}_i - E_i)^2 ,$$

where E_i is the ith measurement of brightness at that point.

(a) Show that the vector that minimizes the error sum is

$$\rho\,\hat{\mathbf{n}} = \left[\sum_{i=1}^{n} \hat{\mathbf{s}}_i \hat{\mathbf{s}}_i^T\right]^{-1} \sum_{i=1}^{n} E_i \hat{\mathbf{s}}_i ,$$

where $\mathbf{a}\mathbf{b}^T$ is the *dyadic product* of the vectors $\mathbf{a}$ and $\mathbf{b}$,

$$\begin{pmatrix} a_x b_x & a_x b_y & a_x b_z \\ a_y b_x & a_y b_y & a_y b_z \\ a_z b_x & a_z b_y & a_z b_z \end{pmatrix} ,$$

and $[\,]^{-1}$ indicates the inverse of the matrix.

(b) What is the smallest number n of measurements needed to guarantee that the indicated matrix inverse exists? Warning: This part of the problem is nontrivial.

(c) How does all this change if we assume that the albedo is one? Hint: You may need to introduce a Lagrange multiplier to enforce this constraint, as explained in the appendix.