B.K.P. Horn
Robot Vision
MIT-Press, 1986

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 experimentally, 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 ($W \cdot m^{-2}$—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 ($W \cdot m^{-2} \cdot sr^{-1}$—watts per

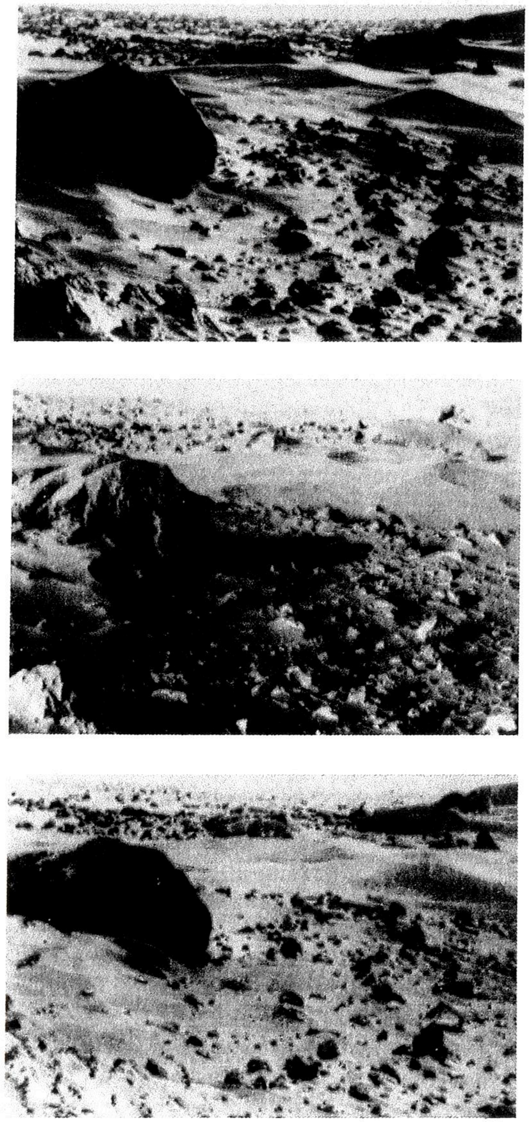

Figure 10-1. The appearance of a surface depends greatly on the lighting conditions. We have to understand how images are formed if we are to recover information about the surface from one or more images. Shown here are three views of the surface of Mars taken by Viking Lander I. (Pictures IPL PIC ID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

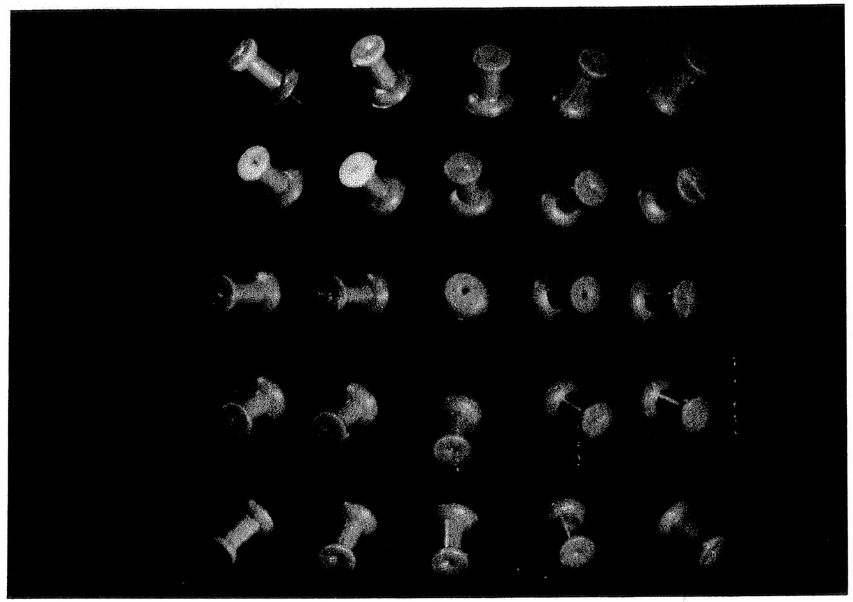

Figure 10-2. The appearance of an object depends greatly on its attitude in space relative to the viewer. Not only does the outline vary, but the brightness pattern within the silhouette changes.

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.

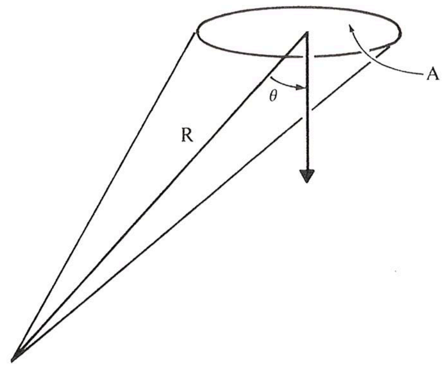

Figure 10-3. The solid angle subtended by a small patch is proportional to its area A and the cosine of the angle of inclination θ; it is inversely proportional to the square of its distance R from the origin.

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

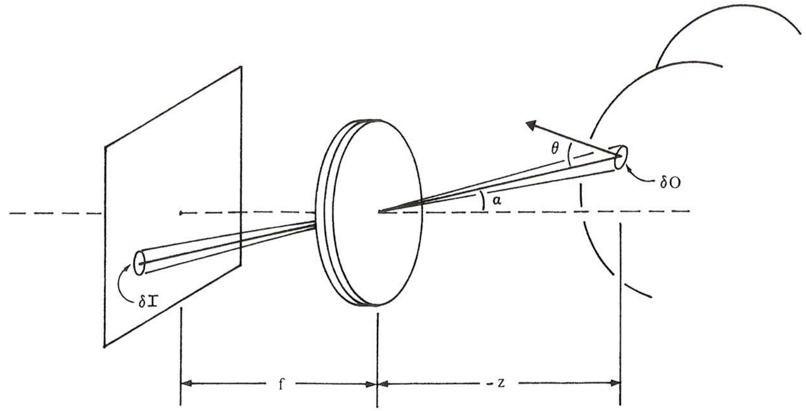

Figure 10-4. To see how image irradiance is related to the radiance of the surface, we must determine the size of the region in the image that corresponds to the patch on the surface.

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

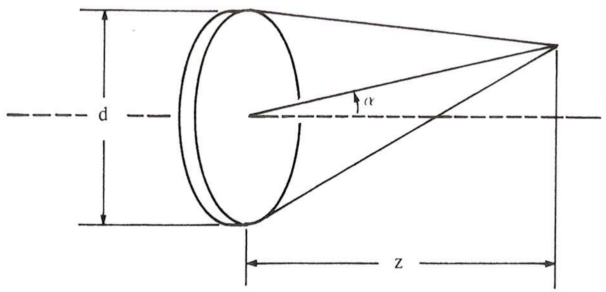

Figure 10-5. To see how image irradiance is related to the radiance of the surface, we must determine how much of the light emitted by the surface is gathered by the lens. This depends on the solid angle that the lens subtends when viewed from the light-emitting surface.

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.

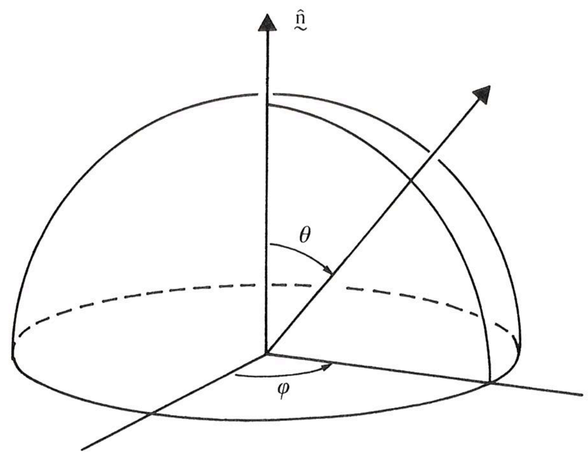

Figure 10-6. The direction of incident and emitted light rays can be specified in a local coordinate system using the polar angle θ and the azimuth ϕ.

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

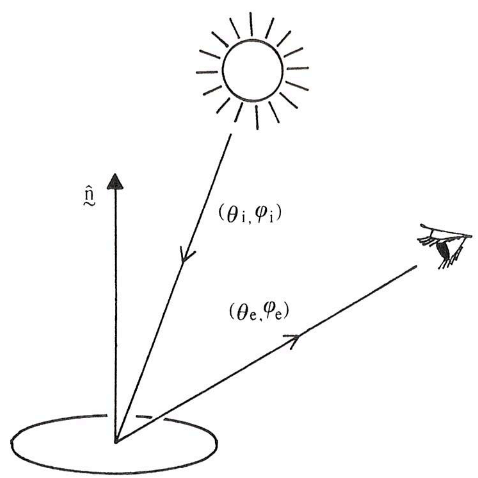

Figure 10-7. The bidirectional reflectance distribution function is the ratio of the radiance of the surface patch as viewed from the direction (θ_e, ϕ_e) to the irradiance resulting from illumination from the direction (θ_i, ϕ_i).

(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

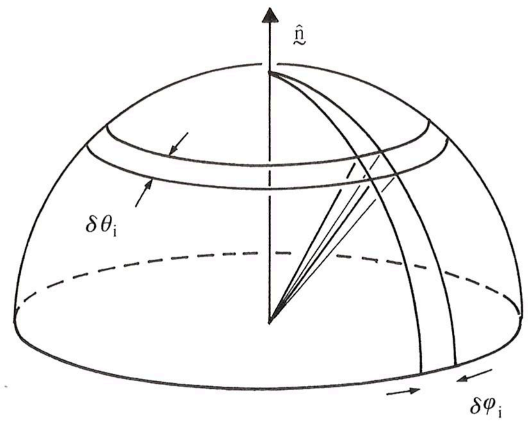

Figure 10-8. In the case of an extended light source, we must integrate the product of the bidirectional reflectance distribution function and the source radiance over all incident directions.

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.

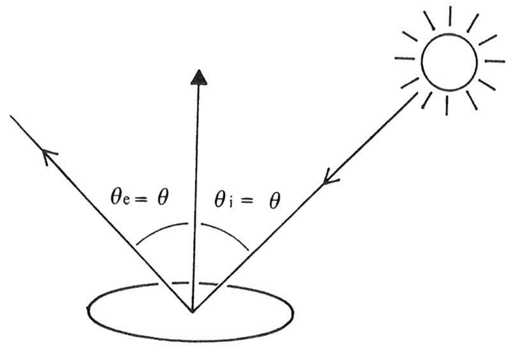

Figure 10-9. A specular surface reflects all incident light in a direction that lies in the same plane as the incident ray and the surface normal. The emittance angle θ_e between the reflected ray and the normal equals the incident angle θ_i between the incident ray and the normal.

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

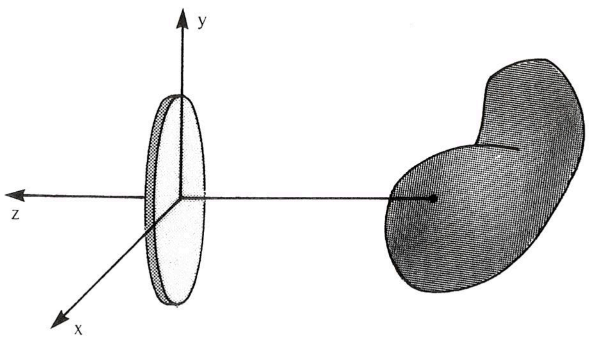

Figure 10-10. A surface can be conveniently described in terms of its perpendicular distance $-z(x, y)$ from some reference plane parallel to the image plane.

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).

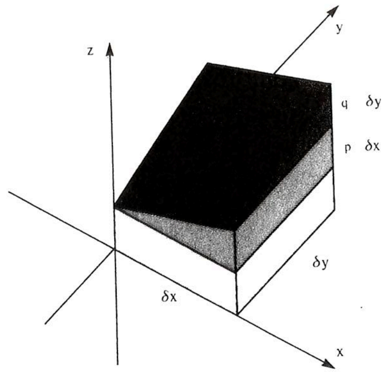

Figure 10-11. Surface orientation can be parameterized by the first partial derivatives p and q of the surface height z.

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,