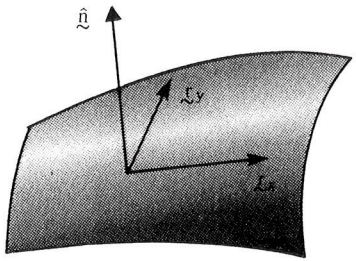

Figure 10-12. The normal is parallel to the cross-product of any two distinct tangents.

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

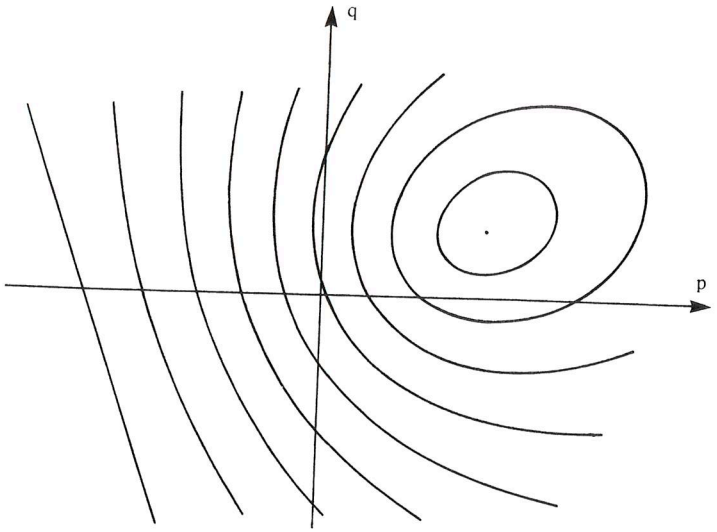

Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p,q)$ occurs at 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

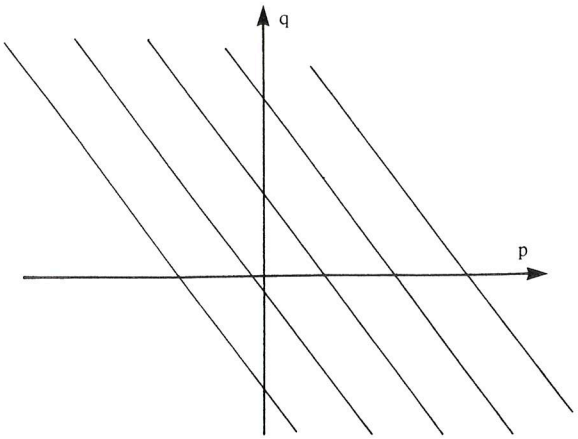

Figure 10-14. In the case of the material in the maria of the moon, the reflectance map can be closely approximated by a function of a linear combination of the components of the gradient. The contours of constant brightness are parallel straight lines in gradient space.

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

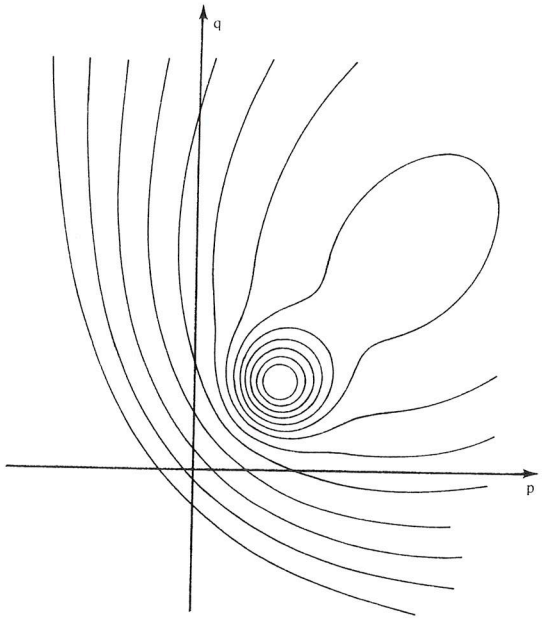

Figure 10-15. Many real surfaces combine diffuse and glossy components of reflection. The glossy component comes from reflection at the interface between air and the surface material, while the diffuse component can be traced to the component of light that has penetrated some distance into the surface, 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

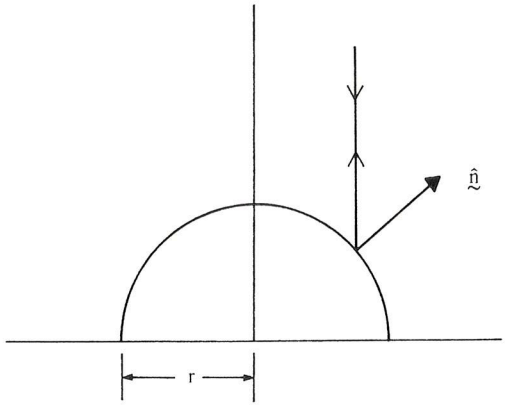

Figure 10-16. The reflectance map is circularly symmetric when the source lies in the same direction as the viewer. In the case of a Lambertian sphere, we obtain a smooth falloff in brightness toward the edge of the circular image of the sphere. This shading effect allows us to estimate the shape of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

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

q

p

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

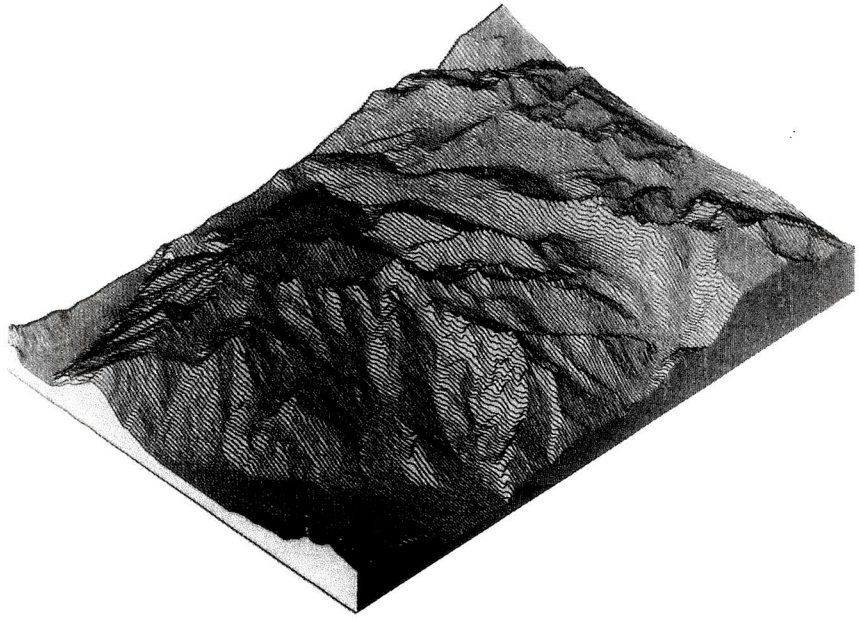

Figure 10-18. A block diagram made from a depth map of the surface of a mountainous region of the earth. (Digital terrain model kindly provided by Kurt Brassel.)

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.

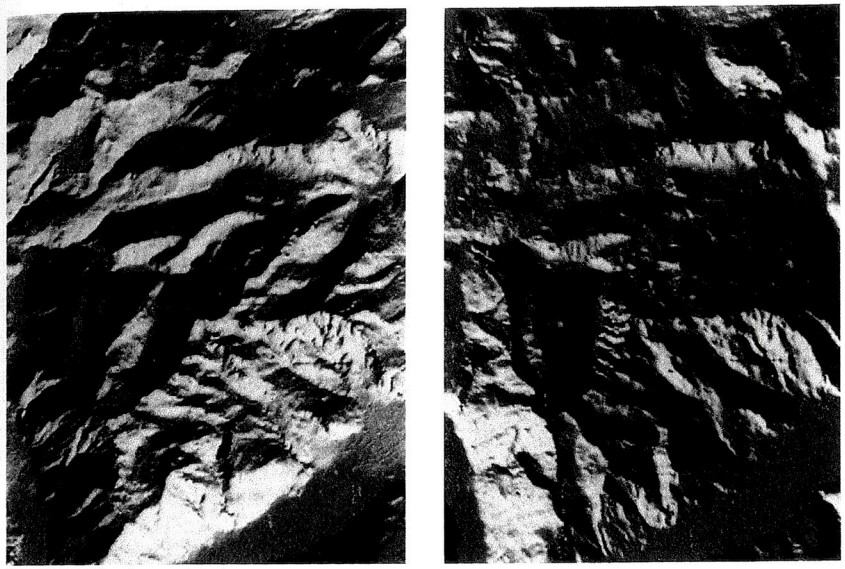

Figure 10-19. Shaded views of the surface shown in the previous figure. Here we are looking down on the surface from above. In the figure on the left, the light is assumed to come from the lower right, as is appropriate for a point on the earth north of the equator in the morning. In the figure on the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

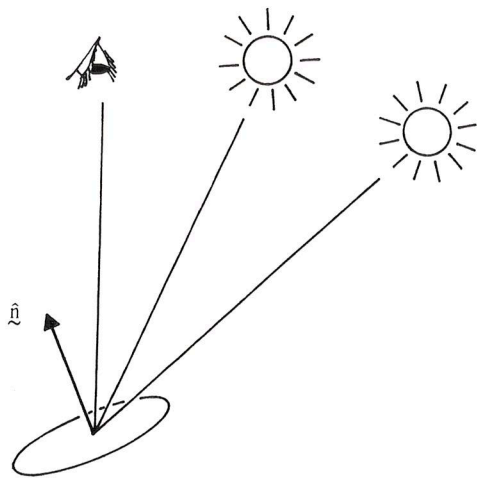

Figure 10-20. In photometric stereo, more than one image is taken from the same viewpoint with different lighting. The brightness of a patch of the surface will generally be different under the two lighting conditions.

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

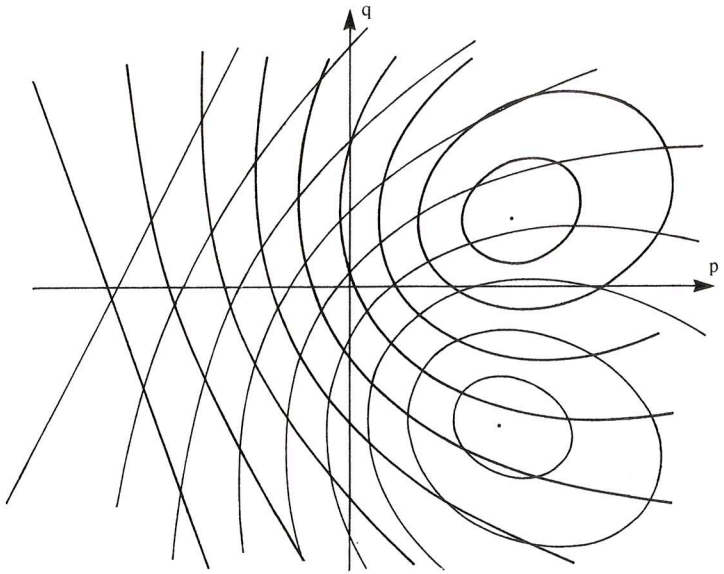

Figure 10-21. In the case of a Lambertian surface illuminated successively by two different point sources, there are at most two surface orientations that produce a particular pair of brightness values. These are found at the intersection of the corresponding contours in two superimposed reflectance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

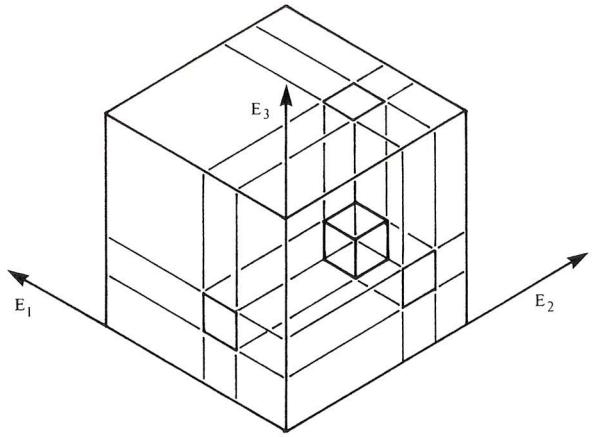

Figure 10-22. Photometric stereo is a fast and robust method for recovering surface shape. The calculation can be encoded in a lookup table based on quantized values of the brightness measurements.

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

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

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