

Chapter 2

RADIOMETRY — MEASURING LIGHT

In this chapter, we introduce a vocabulary with which we can describe the behaviour of light. There are no vision algorithms, but definitions and ideas that will be useful later on. Some readers may find more detail here than they really want; for their benefit, sections 2.4, 2.5 and 2.6 give quick definitions of the main terms we use later on.

2.1 Light in Space

The measurement of light is a field in itself, known as **radiometry**. We need a series of units that describe how energy is transferred from light sources to surface patches, and what happens to the energy when it arrives at a surface. The first matter to study is the behaviour of light in space.

2.1.1 Foreshortening

At each point on a piece of surface is a hemisphere of directions, along which light can arrive or leave (figure 2.1). Two sources that generate the same pattern on this input hemisphere must have the same effect on the surface at this point (because an observer at the surface can't tell them apart). This applies to sources, too; two surfaces that generate the same pattern on a source's output hemisphere must receive the same amount of energy from the source.

This means that the orientation of the surface patch with respect to the direction in which the illumination is travelling is important. As a source is tilted with respect to the direction in which the illumination is travelling, it “looks smaller” to a patch of surface. Similarly, as a patch is tilted with respect to the direction in which the illumination is travelling, it “looks smaller” to the source.

The effect is known as **foreshortening**. Foreshortening is important, because *from the point of view of the source* a small patch appears the same as a large patch that is heavily foreshortened, and so must receive the same energy.




Figure 2.1. A point on a surface sees the world along a hemisphere of directions centered at the point; the surface normal is used to orient the hemisphere, to obtain the θ , ϕ coordinate system that we use consistently from now on to describe angular coordinates on this hemisphere. Usually in radiation problems we compute the brightness of the surface by summing effects due to all incoming directions, so that the fact we have given no clear way to determine the direction in which $\phi = 0$ is not a problem.

2.1.2 Solid Angle

The pattern a source generates on an input hemisphere can be described by the **solid angle** that the source subtends. Solid angle is defined by analogy with angle on the plane.

The angle subtended on the plane by an infinitesimal line segment of length dl at a point \mathbf{p} can be obtained by projecting the line segment onto the unit circle whose center is at \mathbf{p} ; the length of the result is the required angle in radians (see Figure 2.2). Because the line segment is infinitesimally short, it subtends an infinitesimally small angle which depends on the distance to the center of the circle and on the orientation of the line:

$$d\phi = \frac{dl \cos \theta}{r}$$

and the angle subtended by a curve can be obtained by breaking it into infinitesimal segments and summing (integration!).

Similarly, the solid angle subtended by a patch of surface at a point \mathbf{x} is obtained by projecting the patch onto the unit sphere whose center is at \mathbf{x} ; the area of the result is the required solid angle, whose unit is now **steradians**. Solid angle is

usually denoted by the symbol ω . Notice that solid angle captures the intuition in foreshortening — patches that “look the same” on the input hemisphere subtend the same solid angle.




Figure 2.2. Top: The angle subtended by a curve segment at a particular point is obtained by projecting the curve onto the unit circle whose center is at that point, and then measuring the length of the projection. For a small segment, the angle is $(1/r)dl \cos \theta$. **Bottom:** A sphere, illustrating the concept of solid angle. The small circles surrounding the coordinate axes are to help you see the drawing as a 3D surface. An infinitesimal patch of surface is projected onto the unit sphere centered at the relevant point; the resulting area is the solid angle of the patch. In this case, the patch is small, so that the angle is $(1/r^2)dA \cos \theta$.

If the area of the patch dA is small (as suggested by the infinitesimal form), then the infinitesimal solid angle it subtends is easily computed in terms of the area of the patch and the distance to it as

$$d\omega = \frac{dA \cos \theta}{r^2}$$

where the terminology is given in Figure 2.2.

Solid angle can be written in terms of the usual angular coordinates on a sphere (illustrated in Figure 2.2). From figure 2.1 and the expression for the length of circular arcs, we have that infinitesimal steps $(d\theta, d\phi)$ in the angles θ and ϕ cut out a region of solid angle on a sphere given by:

$$d\omega = \sin \theta d\theta d\phi$$

Both of these expressions are worth remembering, as they turn out to be useful for a variety of applications.

2.1.3 Radiance

The distribution of light in space is a function of position and direction. For example, consider shining a torch with a narrow beam in an empty room at night — we need to know where the torch is shining from, and in what direction it is shining. The effect of the illumination can be represented in terms of the power an infinitesimal patch of surface would receive if it were inserted into space at a particular point and orientation. We will use this approach to obtain a unit of measurement.

Definition of Radiance

The appropriate unit for measuring the distribution of light in space is **radiance**, which is defined as:

the amount of energy travelling at some point in a specified direction, per unit time, per unit area *perpendicular to the direction of travel*, per unit solid angle (from [Sillion, 1994])

The units of radiance are watts per square meter per steradian ($Wm^{-2}sr^{-1}$). It is important to remember that the square meters in these units are *foreshortened*, i.e. perpendicular to the direction of travel. This means that a small patch viewing a source frontally collects more energy than the same patch viewing a source radiance along a nearly tangent direction — the amount of energy a patch collects from a source depends both on how large the source looks from the patch *and* on how large the patch looks from the source.

Radiance is a function of position and direction (the torch with a narrow beam is a good model to keep in mind — you can move the torch around, and point the beam in different directions). The radiance at a point in space is usually denoted $L(\mathbf{x}, \text{direction})$, where \mathbf{x} is a coordinate for position — which can be a point in free space or a point on a surface — and we use some mechanism for specifying direction.

One way to specify direction is to use (θ, ϕ) coordinates established using some surface normal. Another is to write $\mathbf{x}_1 \rightarrow \mathbf{x}_2$, meaning the direction from point \mathbf{x}_1 to \mathbf{x}_2 . We shall use both, depending on which is convenient for the problem at hand.

Radiance is Constant Along a Straight Line

For the vast majority of important vision problems, it is safe to assume that light does not interact with the medium through which it travels — i.e. that we are in a vacuum. Radiance has the highly desirable property that, for two points \mathbf{p}_1 and \mathbf{p}_2 (which have a line of sight between them), the radiance leaving \mathbf{p}_1 in the direction of \mathbf{p}_2 is the same as the radiance arriving at \mathbf{p}_2 from the direction of \mathbf{p}_1 .

The following proof may look vacuous at first glance; it's worth studying carefully, because it is the key to a number of other computations. Figure 2.3 shows a patch of surface radiating in a particular direction. From the definition, if the




Figure 2.3. Light intensity is best measured in radiance, because radiance does not go down along straight line paths in a vacuum (or, for reasonable distances, in clear air). This is shown by an energy conservation argument in the text, where one computes the energy transferred from a patch dA_1 to a patch dA_2

radiance at the patch is $L(\mathbf{x}_1, \theta, \phi)$, then the energy transmitted by the patch into an infinitesimal region of solid angle $d\omega$ around the direction θ, ϕ in time dt is

$$L(\mathbf{x}_1, \theta, \phi)(\cos \theta_1 dA_1)(d\omega)(dt),$$

(i.e. radiance times the foreshortened area of the patch times the solid angle into which the power is radiated times the time for which the power is radiating).

Now consider two patches, one at \mathbf{x}_1 with area dA_1 and the other at \mathbf{x}_2 with area dA_2 (see Figure 2.3). To avoid confusion with angular coordinate systems, write the angular direction from \mathbf{x}_1 to \mathbf{x}_2 as $\mathbf{x}_1 \rightarrow \mathbf{x}_2$. The angles θ_1 and θ_2 are as defined in figure 2.3.

The radiance leaving \mathbf{x}_1 in the direction of \mathbf{x}_2 is $L(\mathbf{x}_1, \mathbf{x}_1 \rightarrow \mathbf{x}_2)$ and the radiance arriving at \mathbf{x}_2 from the direction of \mathbf{x}_1 is $L(\mathbf{x}_2, \mathbf{x}_1 \rightarrow \mathbf{x}_2)$.

This means that, in time dt , the energy leaving \mathbf{x}_1 towards \mathbf{x}_2 is

$$d^3E_{1 \rightarrow 2} = L(\mathbf{x}_1, \mathbf{x}_1 \rightarrow \mathbf{x}_2) \cos \theta_1 d\omega_{2(1)} dA_1 dt$$

where $d\omega_{2(1)}$ is the solid angle subtended by patch 2 at patch 1 (energy emitted into this solid angle arrives at 2; all the rest disappears into the void). The notation $d^3E_{1 \rightarrow 2}$ implies that there are three infinitesimal terms involved.

From the expression for solid angle above,

$$d\omega_{2(1)} = \frac{\cos \theta_2 dA_2}{r^2}$$

Now the energy leaving 1 for 2 is:

$$\begin{aligned} d^3E_{1 \rightarrow 2} &= L(\mathbf{x}_1, \mathbf{x}_1 \rightarrow \mathbf{x}_2) \cos \theta_1 d\omega_{2(1)} dA_1 dt \\ &= L(\mathbf{x}_1, \mathbf{x}_1 \rightarrow \mathbf{x}_2) \frac{\cos \theta_1 \cos \theta_2 dA_2 dA_1 dt}{r^2} \end{aligned}$$

Because the medium is a vacuum, it does not absorb energy, so that the energy arriving at 2 from 1 is the same as the energy leaving 1 in the direction of 2. The energy arriving at 2 from 1 is:

$$\begin{aligned} d^3E_{1 \rightarrow 2} &= L(\mathbf{x}_2, \mathbf{x}_1 \rightarrow \mathbf{x}_2) \cos \theta_2 d\omega_{1(2)} dA_2 dt \\ &= L(\mathbf{x}_2, \mathbf{x}_1 \rightarrow \mathbf{x}_2) \frac{\cos \theta_2 \cos \theta_1 dA_1 dA_2 dt}{r^2} \end{aligned}$$

which means that $L(\mathbf{x}_2, \mathbf{x}_1 \rightarrow \mathbf{x}_2) = L(\mathbf{x}_1, \theta, \phi)$, so that *radiance is constant along (unoccluded) straight lines*.

2.2 Light at Surfaces

When light strikes a surface, it may be absorbed, transmitted, or scattered; usually, a combination of these effects occur. For example, light arriving at skin can be scattered at various depths into tissue and reflected from blood or from melanin in there; can be absorbed; or can be scattered tangential to the skin within a film of oil and then escape at some distant point.

The picture is complicated further by the willingness of some surfaces to absorb light at one wavelength, and then radiate light at a different wavelength as a result. This effect, known as **fluorescence**, is fairly common: scorpions fluoresce visible light under x-ray illumination; human teeth fluoresce faint blue under ultraviolet light (nylon underwear tends to fluoresce, too, and false teeth generally do not — the resulting embarrassments led to the demise of uv lights in discotheques); and laundry can be made to look bright by washing powders that fluoresce under ultraviolet light. Furthermore, a surface that is warm enough emits light in the visible range.

2.2.1 Simplifying Assumptions

It is common to assume that all effects are local, and can be explained with a macroscopic model with no fluorescence or emission. This is a reasonable model for the kind of surfaces and decisions that are common in vision. In this model:

- the radiance leaving a point on a surface is due only to radiance arriving at this point (although radiance may change *directions* at a point on a surface, we assume that it does not skip from point to point);
- we assume that all light leaving a surface at a given wavelength is due to light arriving at that wavelength;
- we assume that the surfaces do not generate light internally, and treat sources separately.

2.2.2 The Bidirectional Reflectance Distribution Function

We wish to describe the relationship between incoming illumination and reflected light. This will be a function of both the direction in which light arrives at a surface and the direction in which it leaves.

Irradiance

The appropriate unit for representing incoming power which is **irradiance**, defined as:

incident power per unit area *not foreshortened*.

A surface illuminated by radiance $L_i(\mathbf{x}, \theta_i, \phi_i)$ coming in from a differential region of solid angle $d\omega$ at angles (θ_i, ϕ_i) receives irradiance

$$L_i(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega$$

where we have multiplied the radiance by the foreshortening factor and by the solid angle to get irradiance. The main feature of this unit is that we could compute all the power incident on a surface at a point by summing the irradiance over the whole input hemisphere — which makes it the natural unit for *incoming* power.

The BRDF

The most general model of local reflection is the **bidirectional reflectance distribution function**, usually abbreviated BRDF. The BRDF is defined as

the ratio of the radiance in the outgoing direction to the incident irradiance (after [Sillion, 1994])

so that, if the surface of the preceding paragraph was to emit radiance $L_o(\mathbf{x}, \theta_o, \phi_o)$, its BRDF would be:

$$\rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) = \frac{L_o(\mathbf{x}, \theta_o, \phi_o)}{L_i(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega}$$

The BRDF has units of inverse steradians (sr^{-1}), and could vary from 0 (no light reflected in that direction) to infinity (unit radiance in an exit direction resulting from arbitrary small radiance in the incoming direction). The BRDF is symmetric in the incoming and outgoing direction, a fact known as the Helmholtz reciprocity principle.

Properties of the BRDF

The radiance leaving a surface due to irradiance *in a particular direction* is easily obtained from the definition of the BRDF:

$$L_o(\mathbf{x}, \theta_o, \phi_o) = \rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) L_i(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega$$

More interesting is the radiance leaving a surface due to its irradiance (whatever the direction of irradiance). We obtain this by summing over contributions from all incoming directions:

$$L_o(\mathbf{x}, \theta_o, \phi_o) = \int_{\Omega} \rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) L_i(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega$$

where Ω is the incoming hemisphere. From this we obtain the fact that the BRDF is not an arbitrary symmetric function in four variables.

To see this, assume that a surface is subjected to a radiance of $1/\cos \theta_i \text{ W m}^{-2} \text{ sr}^{-1}$. This means that the total energy arriving at the surface is:

$$\begin{aligned} \int_{\Omega} \frac{1}{\cos \theta} \cos \theta d\omega &= \int_0^{2\pi} \int_0^{\frac{\pi}{2}} \sin \theta d\theta d\phi \\ &= 2\pi \end{aligned}$$

We have assumed that any energy leaving at the surface leaves from the same point at which it arrived, and that no energy is generated within the surface. This means that the total energy leaving the surface must be less than or equal to the amount arriving. So we have

$$\begin{aligned} 2\pi &\geq \int_{\Omega_o} L_o(\mathbf{x}, \theta_o, \phi_o) \cos \theta_o d\omega_o \\ &= \int_{\Omega_o} \int_{\Omega_i} \rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) L_i(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega_i d\omega_o \\ &= \int_{\Omega_o} \int_{\Omega_i} \rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) d\omega_i d\omega_o \\ &= \int_0^{2\pi} \int_0^{\frac{\pi}{2}} \int_0^{2\pi} \int_0^{\frac{\pi}{2}} \rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) \sin \theta_i d\theta_i d\phi_i \sin \theta_o d\theta_o d\phi_o \end{aligned}$$

What this tells us is that, although the BRDF can be large for some pairs of incoming and outgoing angles, it can't be large for many.

2.3 Important Special Cases

Radiance is a fairly subtle quantity, because it depends on angle. This generality is sometimes essential — for example, for describing the distribution of light in space in the torch beam example above. As another example, fix a compact disc and illuminate its underside with a torch beam. The intensity and colour of light reflected from the surface depends very strongly on the angle from which the surface is viewed and on the angle from which it is illuminated. The CD example is worth trying, because it illustrates how strange the behaviour of reflecting surfaces can be; it also illustrates how accustomed we are to dealing with surfaces that do not behave in this way. For many surfaces — cotton cloth is one good example — the dependency of reflected light on angle is weak or non-existent, so that a system of units that are independent of angle is useful.

2.3.1 Radiosity

If the radiance leaving a surface is independent of exit angle, there is no point in describing it using a unit that explicitly depends on direction. The appropriate unit is **radiosity**, defined as

the total power leaving a point on a surface per unit area on the surface
(from [Sillion, 1994])

Radiosity, which is usually written as $B(\mathbf{x})$ has units watts per square meter (Wm^{-2}). To obtain the radiosity of a surface at a point, we can sum the radiance leaving the surface at that point over the whole exit hemisphere. Thus, if \mathbf{x} is a point on a surface emitting radiance $L(\mathbf{x}, \theta, \phi)$, the radiosity at that point will be:

$$B(\mathbf{x}) = \int_{\Omega} L(\mathbf{x}, \theta, \phi) \cos \theta d\omega$$

where Ω is the exit hemisphere and the term $\cos \theta$ turns foreshortened area into area (look at the definitions again!); $d\omega$ can be written in terms of θ, ϕ as above.

The Radiosity of a Surface with Constant Radiance

One result to remember is the relationship between the radiosity and the radiance of a surface patch *where the radiance is independent of angle*. In this case $L_o(x, \theta_o, \phi_o) = L_o(\mathbf{x})$. Now the radiosity can be obtained by summing the radiance leaving the surface over all the directions in which it leaves:

$$B(\mathbf{x}) = \int_{\Omega} L_o(\mathbf{x}) \cos \theta d\omega$$

$$\begin{aligned}
&= L_o(\mathbf{x}) \int_0^{\frac{\pi}{2}} \int_0^{2\pi} \cos \theta \sin \theta d\phi d\theta \\
&= \pi L_o(\mathbf{x})
\end{aligned}$$

2.3.2 Directional Hemispheric Reflectance

The BRDF is also a subtle quantity, and BRDF measurements are typically difficult, expensive and not particularly repeatable. This is because surface dirt and aging processes can have significant effects on BRDF measurements; for example, touching a surface will transfer oil to it, typically in little ridges (from the fingertips) which can act as lenses and make significant changes in the directional behaviour of the surface.

The light leaving many surfaces is largely independent of the exit angle. A natural measure of a surface's reflective properties in this case is the **directional-hemispheric reflectance**, usually termed ρ_{dh} , defined as:

the fraction of the incident irradiance in a given direction that is reflected by the surface, whatever the direction of reflection (after [Sillion, 1994])

The directional hemispheric reflectance of a surface is obtained by summing the radiance leaving the surface over all directions, and dividing by the irradiance in the direction of illumination, which gives:

$$\begin{aligned}
\rho_{dh}(\theta_i, \phi_i) &= \frac{\int_{\Omega} L_o(\mathbf{x}, \theta_o, \phi_o) \cos \theta_o d\omega_o}{L_i(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega_i} \\
&= \int_{\Omega} \left\{ \frac{L_o(\mathbf{x}, \theta_o, \phi_o) \cos \theta_o}{L_i(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega_i} \right\} d\omega_o \\
&= \int_{\Omega} \rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) \cos \theta_o d\omega_o
\end{aligned}$$

This property is dimensionless, and its value will lie between 0 and 1.

Directional hemispheric reflectance can be computed for any surface. For some surfaces, it will vary sharply with the direction of illumination. A good example is a surface with fine, symmetric triangular grooves which are black on one face and white on the other. If these grooves are sufficiently fine, it is reasonable to use a macroscopic description of the surface as flat, and with a directional hemispheric reflectance that is large along a direction pointing towards the white faces and small along that pointing towards the black.

2.3.3 Lambertian Surfaces and Albedo

For some surfaces the directional hemispheric reflectance does not depend on illumination direction. Examples of such surfaces include cotton cloth, many carpets, matte paper and matte paints. A formal model is given by a surface whose BRDF is independent of outgoing direction (and, by the reciprocity principle, of incoming direction as well). This means the radiance leaving the surface is independent

of angle. Such surfaces are known as **ideal diffuse surfaces** or **Lambertian surfaces** (after George Lambert, who first formalised the idea).

It is natural to use radiosity as a unit to describe the energy leaving a Lambertian surface. For Lambertian surfaces, the directional hemispheric reflectance is independent of direction. In this case the directional hemispheric reflectance is often called their **diffuse reflectance** or **albedo** and written ρ_d . For a Lambertian surface with BRDF $\rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) = \rho$, we have:

$$\begin{aligned}\rho_d &= \int_{\Omega} \rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) \cos \theta_o d\omega_o \\ &= \int_{\Omega} \rho \cos \theta_o d\omega_o \\ &= \rho \int_0^{\frac{\pi}{2}} \int_0^{2\pi} \cos \theta_o \sin \theta_o d\theta_o d\phi_o \\ &= \pi \rho\end{aligned}$$

This fact is more often used in the form

$$\rho_{brdf} = \frac{\rho_d}{\pi}$$

a fact that is useful, and well worth remembering.

Because our sensations of brightness correspond (roughly!) to measurements of radiance, a Lambertian surface will look equally bright from any direction, whatever the direction along which it is illuminated. This gives a rough test for when a Lambertian approximation is appropriate.

2.3.4 Specular Surfaces

A second important class of surfaces are the glossy or mirror-like surfaces, often known as **specular** surfaces (after the Latin word *speculum*, a mirror). An ideal specular reflector behaves like an ideal mirror. Radiation arriving along a particular direction can leave only along the **specular direction**, obtained by reflecting the direction of incoming radiation about the surface normal. Usually some fraction of incoming radiation is absorbed; on an ideal specular surface, the same fraction of incoming radiation is absorbed for every direction, the rest leaving along the specular direction. The BRDF for an ideal specular surface has a curious form (exercise ??), because radiation arriving in a particular direction can leave in only one direction.

Specular Lobes

Relatively few surfaces can be approximated as ideal specular reflectors. A fair test of whether a flat surface can be approximated as an ideal specular reflector is whether one could safely use it as a mirror. Good mirrors are surprisingly hard to make; up until recently, mirrors were made of polished metal. Typically, unless the




Figure 2.4. Specular surfaces commonly reflect light into a lobe of directions around the specular direction, where the intensity of the reflection depends on the direction, as shown on the left. Phong’s model is used to describe the shape of this lobe, in terms of the offset angle from the specular direction.

metal is extremely highly polished and carefully maintained, radiation arriving in one direction leaves in a small lobe of directions around the specular direction. This results in a typical blurring effect. A good example is the bottom of a flat metal pie dish. If the dish is reasonably new, one can see a distorted image of one’s face in the surface but it would be difficult to use as a mirror; a more battered dish reflects a selection of distorted blobs.

Larger specular lobes mean that the specular image is more heavily distorted and is darker (because the incoming radiance must be shared over a larger range of outgoing directions). Quite commonly it is possible to see only a specular reflection of relatively bright objects, like sources. Thus, in shiny paint or plastic surfaces, one sees a bright blob — often called a **specularity** — along the specular direction from light sources, but few other specular effects. It is not often necessary to model the shape of the specular lobe. When the shape of the lobe is modelled, the most common model is the **Phong** model, which assumes that only point light sources are specularly reflected. In this model, the radiance leaving a specular surface is proportional to $\cos^n(\delta\theta) = \cos^n(\theta_o - \theta_s)$, where θ_o is the exit angle, θ_s is the specular direction and n is a parameter. Large values of n lead to a narrow lobe and small, sharp specularities and small values lead to a broad lobe and large specularities with rather fuzzy boundaries.

2.3.5 The Lambertian + Specular Model

Relatively few surfaces are either ideal diffuse or perfectly specular. Very many surfaces can be approximated as having a surface BRDF which is a combination of a Lambertian component and a specular component, which usually has some form of narrow lobe. Usually, the specular component is weighted by a **specular albedo**. Again, because specularities tend not to be examined in detail, the shape of this lobe is left unspecified. In this case, the surface *radiance* (because it must

now depend on direction) in a given direction is typically approximated as:

$$L(\mathbf{x}, \theta_o, \phi_o) = \rho_d(\mathbf{x}) \int_{\Omega} L(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega + \rho_s(\mathbf{x}) L(\mathbf{x}, \theta_s, \phi_s) \cos^n(\theta_s - \theta_o)$$

where θ_s, ϕ_s give the specular direction and ρ_s is the specular albedo. As we shall see, it is common not to reason about the exact magnitude of the specular radiance term.

Using this model implicitly excludes “too narrow” specular lobes, because most algorithms expect to encounter occasional small, compact specularities from light sources. Surfaces with too narrow specular lobes (mirrors) produce overwhelming quantities of detail in specularities. Similarly, “too broad” lobes are excluded because the specularities would be hard to identify.

2.4 Quick Reference: Radiometric Terminology for Light

Term	Definition	Units	Application
Radiance	the quantity of energy travelling at some point in a specified direction, per unit time, per unit area <i>perpendicular to the direction of travel</i> , per unit solid angle.	Wm^2sr^{-1}	representing light travelling in free space; representing light reflected from a surface when the amount reflected depends strongly on direction
Irradiance	total incident power per unit surface area	Wm^{-2}	representing light arriving at a surface
Radiosity	the total power leaving a point on a surface per unit area on the surface	Wm^{-2}	representing light leaving a diffuse surface

2.5 Quick Reference: Radiometric Properties of Surfaces

Term	Definition	Units	Application
BRDF (Bidirectional Reflectance Distribution Function)	the ratio of the radiance in the outgoing direction to the incident irradiance	sr^{-1}	representing reflection off general surfaces where reflection depends strongly on direction
Directional Hemispheric Reflectance	the fraction of the incident irradiance in a given direction that is reflected by the surface, whatever the direction of reflection	unitless	representing reflection off a surface where direction is unimportant
Albedo	Directional hemispheric reflectance of a diffuse surface	unitless	representing a diffuse surface

2.6 Quick Reference: Important Types of Surface

Term	Definition	Examples
Diffuse surface; Lambertian surface	A surface whose BRDF is constant	Cotton cloth; many rough surfaces; many paints and papers; surfaces whose apparent brightness doesn't change with viewing direction
Specular surface	A surface that behaves like a mirror	Mirrors; polished metal
Specularity	Small bright patches on a surface that result from specular components of the BRDF	

2.7 Notes

We strongly recommend François Sillion’s excellent book [Sillion, 1994], for its very clear account of radiometric calculations. There are a variety of more detailed publications for reference []. Our discussion of reflection is thoroughly superficial. The specular plus diffuse model appears to be originally due to Cook, Torrance and Sparrow. A variety of modifications of this model appear in computer vision and computer graphics; see, for example []. Reflection models can be derived by combining a statistical description of surface roughness with electromagnetic considerations (e.g. []) or by adopting scattering models (e.g. [], where a surface is modelled by colourant particles embedded in a matrix, and a scattering model yields an approximate BRDF).

Top of the list of effects we omitted to discuss is off-specular glints, followed by specular backscatter. Off-specular glints commonly arise in brushed surfaces, where there is a large surface area oriented at a substantial angle to the macroscopic surface normal. This leads to a second specular lobe, due to this region. These effects can confuse algorithms that reason about shape from specularities, if the reasoning is close enough. Specular backscatter occurs when a surface reflects light back in the source direction — usually for a similar reason that off-specular glints occur. Again, the effect is likely to confuse algorithms that reason about shape from specularities.

It is commonly believed that rough surfaces are Lambertian. This belief has a substantial component of wishful thinking, because rough surfaces often have local shadowing effects that make the radiance reflected quite strongly dependent on the illumination angle. For example, a stucco wall illuminated at a near grazing angle shows a clear pattern of light and dark regions where facets of the surface face toward the light or are shadowed. If the same wall is illuminated along the normal, this pattern largely disappears. Similar effects at a finer scale are averaged to endow rough surfaces with measurable departures from a Lambertian model (for details, see [Nayar and Oren, 1995; Oren and Nayar, 1995; Wolff, 1996; Wolff, 1994]). Determining non-Lambertian models for surfaces that appear to be diffuse is a well established line of enquiry.

Another example of an object that does not support a simple macroscopic surface model is a field of flowers. A distant viewer should be able to abstract this field as a “surface”; however, doing so leads to a surface with quite strange properties. If one views such a field along a normal direction, one sees mainly flowers; a tangential view reveals both stalks and flowers, meaning that the colour changes dramatically (the effect is explored in []).

2.8 Assignments

Exercises

1. How many steradians in a hemisphere?
2. We have proven that radiance does not go down along a straight line *in a*

non-absorbing medium, which makes it a useful unit. Show that if we were to use power per square meter of foreshortened area (which is irradiance), the unit must change with distance along a straight line. How significant is this difference?

3. **An absorbing medium:** assume that the world is filled with an isotropic absorbing medium. A good, simple model of such a medium is obtained by considering a line along which radiance travels. If the radiance along the line is N at x , it will be $N - (\alpha dx)N$ at $x + dx$.
 - Write an expression for the radiance transferred from one surface patch to another in the presence of this medium.
 - Now *qualitatively* describe the distribution of light in a room filled with this medium, for α small and large positive numbers. The room is a cube, and the light is a single small patch in the center of the ceiling. Keep in mind that if α is large and positive, very little light will actually reach the walls of the room.
4. Identify common surfaces that are neither Lambertian nor specular, using the underside of a CD as a working example. There are a variety of important biological examples, which are often blue in colour. Give at least two different reasons that it could be advantageous to an organism to have a non-Lambertian surface.
5. Show that for an ideal diffuse surface the directional hemispheric reflectance is constant; now show that if a surface has constant directional hemispheric reflectance, it is ideal diffuse.
6. Show that the BRDF of an ideal specular surface is

$$\rho_{bd}(\theta_o, \phi_o, \theta_i, \phi_i) = \rho_s(\theta_i) \{2\delta(\sin^2 \theta_o - \sin^2 \theta_i)\} \{\delta(\phi_o - \phi_i)\}$$

where $\rho_s(\theta_i)$ is the fraction of radiation that leaves.

7. Why are specularities brighter than diffuse reflection?
8. A surface has constant BRDF. What is the maximum possible value of this constant? Now assume that the surface is known to absorb 20% of the radiation incident on it (the rest is reflected); what is the value of the BRDF?
9. The eye responds to radiance. Explain why Lambertian surfaces are often referred to as having a brightness that is independent of viewing angle.
10. Show that the solid angle subtended by a sphere of radius ϵ at a point a distance r away from the center of the sphere is approximately $\pi(\frac{\epsilon}{r})^2$, for $r \gg \epsilon$.