Lighting and reflectance

All of discussion up to now has been about geometry. Today we will talk about *radiometry* which concerns measures of light.[1] We will begin by talking informally about light sources, then we will discuss light in space, then we discuss light arriving at a surface, and finally we discuss light reflection from a surface.

Light sources

You know what light sources are, but you probably have not thought about how to model them. Think about the sun vs. a lamp vs. a spotlight vs. the sky. We informally think of all of these as light sources, but each illuminates a scene in a different way.

- *sunlight*: The key property of sunlight is that all the rays are approximately parallel. Sunlight is sometimes referred to a *parallel source*. According to this model, any light ray from the sun that arrives at a point $\mathbf{X}$ in the scene has direction $\mathbf{l}$. Note that the sun subtends a small angle in the sky, and so the rays from the sun are not all exactly parallel. The sun (and moon) are each about one half a degree of visual angle.

- *candle, light bulb*: Such sources are located in the scene. They have a finite size which is typically small compared to the distance to the objects that they illuminate. For this reason, it is common to approximate the location of the source by a point, which gives us a *point light source* model. According to this model, light is emitted from a point and radiates outwards from that point.

- *spotlight*: Many small light sources in a scene do not illuminate equally in all directions, but rather are designed to send most of their illumination in particular directions. A spotlight or car headlight are two examples. Even a typical desklamp is designed to have non-uniform directional properties – the light is supposed to illuminate the desk, rather than shining directly into one's eyes.

- *sky*: the sky is quite different from the above sources. Light from the sky comes from all directions, and it can be seen at all positions.

[BEGIN ASIDE (not discussed in class): While the above sources seem to define distinct categories, in the real world these categories do not have "hard edges" between them. For example:

- sunlight is accompanied by a blue sky. You never get only sunlight. (If the light is coming from the moon, then you get scattered skylight as well, though this scattered light is typically so weak that the visual system can hardly detect it.)

- if sunlight enters a room through a small window, the window will act as a spotlight. Movie studios take advantage of this, to simulate sunlight coming through windows.

[1] The term "radiometry" actually is more general than this, and includes measurements of electromagnetic waves outside the visible range e.g. into ultraviolet and infrared and beyond. The term *photometry* is usually used for visible light range only.

- what happens if the beam of sunlight through the window hits a white wall on the opposite side of the room? In this case, the bright spot on the wall acts like a small light source, like a panel light. This is an example of *interreflections*, where surfaces in the scene illuminate each other.
- if it is a cloudy day outside, then the light that enters the room through the small window would come from many different directions. Now the window itself would act as a small light source – similar in effect to a panel light on the wall. The light would diverge into the room in many different directions.

If you wish to read more about these categories of light sources and their relationships, have a look at: `http://www.cim.mcgill.ca/~langer/MY_PAPERS/Langer-Zucker-CVIU-97.pdf`
END ASIDE]

We will return to the light problem later in the course when we discuss "shape from shading". For now let's turn to another aspect of lighting. To talk about the lighting in a real scene, we need to be able to make measurements. I'll now introduce the basic terminology one needs.

Radiometry

Light travels in straight lines through space[2]. We would like to be able to talk about how bright a light ray is. (Equivalently, we could talk about how the density of light rays.) Imagine a thin black cylindrical tube and consider the amount of light that passes through the tube. We assume that light gets absorbed if it strikes the side of the tube, and so we are only considering those rays that pass through the tube without striking the sides. We measure the power[3] of light through the tube. I am ignoring color properties for now.

The power of light that gets through the tube, depends on the length of the tube and on the cross-sectional area the tube (πr^2, where r is the tube radius). If you hold the area constant and increase the length of the tube, you decrease the amount of light that passes through the tube because you will have a tighter restriction on the direction of rays that can pass through the tube. Similarly, if you hold the length of the tube constant, and you decrease the area of the tube, you again decrease the number of rays that pass through the tube. (Notice that these arguments are reminiscent of our discussion of f-stop last class.)

For thin tubes, the angle subtended by the diameter at the far end of the tube, when seen from the tube axis at the opening of the near end of the tube is the ratio $\frac{diameter}{tubelength}$. This angle is in units of *radians*, which you will recall is an arclength on a unit circle. That is, if an object is seen from a point, and that object is projected onto a unit circle, then the angle subtended by the object is the length of the circular arc of the projection. One radian is $\frac{1}{2\pi} \times 360 \approx 57$ degrees.

Solid angle

There is a corresponding angular measurement called *solid angle* which specifies the area of a unit sphere subtended by an object (when viewed from a point, namely the center of the sphere). The unit is *steradians*. The maximum value of solid angle is 4π, which is the entire area of the unit

[2] I am ignoring the wave properties of light which become important at spatial scales below about $\frac{1}{1000}mm$.
[3] light energy per second, or Watts

sphere. For example, if you are inside a room, then the solid angle subtended by the room is 4π steradians.

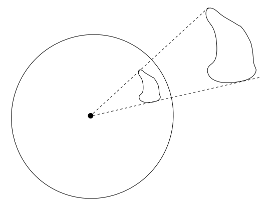

We said earlier that the power of light that passes through the tube depends on the length and width of the tube. Alternatively, the power of light that passes through the tube depends on the area of the tube and on the solid angle subtended by one end of the tube when seen from the other end. (For a thin tube, this solid angle is the area of the cross section of the tube divided by the squared length of the tube.)

Radiance

We are now ready to define radiance. Please then end of hollow black tube (like a drinking straw) at some point $\mathbf{X}$ in space and point it in some direction $\mathbf{l}$. Measure the power of light passing through the tube. Define the *radiance* L of the ray to be the light power[4] per unit cross section area per unit steradian,

$$\text{radiance,} \quad L(\mathbf{X}, \mathbf{l}) \equiv \frac{\text{power}}{\text{cross section area} \, * \, \text{solid angle}}$$

For the radiance of the ray to be well defined, it needs to be approximately constant as the cross section shrinks to zero.[5] Essentially, this requires that the radiance vary sufficiently smoothly as a function of $\mathbf{X}$,$\mathbf{l}$.

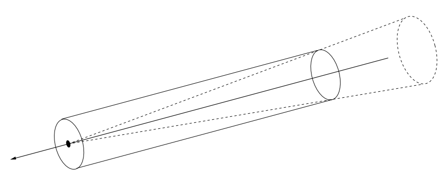

Interesting, if there is no scattering of light in space (e.g. fog), then *radiance is constant along a light ray.* Intuitively, if you place your thin tube at a point $\mathbf{X}$ in space and point it in some direction $\mathbf{l}$, and then move the tube in this direction, the power that you get through the tube will not change (i.e. the brightness of the light you will see through the tube will not change). Things that we look at do not get brighter when we walk toward them. They get bigger (in terms of their solid angle) but they don't get brighter. Note: small light sources *do* seem to get dimmer as you move away from them. But the reason is that the solid angle they subtend gets smaller – not that the radiance decreases.

[4] energy per unit time

[5] Physics tells us that when the cross section diameter shrinks to the wavelength of light, this definition doesn't work because the wave properties of light take over then.

Radiant intensity

There are a number of concepts in radiometry that are related, but are distinct from each other. For example, let's consider the light given off by a small light source, such as a lamp or spotlight. You don't really want to talk about the radiance of each of the rays of light coming from this source. So what do you do?

Typically, you model such a source as a point and you consider a large imaginary sphere centered at that point. You then ask what is the power of light reaching each small area of the sphere. This would tell you how bright is your small light source in each direction. Formally, define the *radiant intensity* as the light power that a small source gives off per unit solid angle in each direction. The units are Watts per steradian. Radiant intensity is a function defined on the unit sphere. (The standard symbol for radiant intensity is I. We will be using the symbol I for an image intensity, though...)

[ASIDE: Lamp manufacturers typically use radiant intensity to characterize how a particular lamp distributes light as a function of direction. This is very important in many applications, for example, you need to space the lights in a tunnel a certain distance apart to provide uniform lighting on the road, and you want to prevent too much glare i.e. light shining directly into people's eyes.]

Irradiance

Another important concept is the power of light that arrives *at a surface*. If you take a point $\mathbf{x}$ on a surface, then light can arrive at $\mathbf{x}$ from a hemisphere of directions. (If the surface is transparent, then light can also arrive from behind the surface, but let's not consider this.)

We can talk about the power of light arriving at a point, but usually one instead talks about the amount of light arriving in a small area containing that point. This allows one to capture the orientation (normal) of the surface which plays an important role.

The *irradiance* E is defined as the light power arriving per unit area on the surface,

$$E(\mathbf{x}) = \int L(\mathbf{x}, \mathbf{l}_{in}) \ \mathbf{n}(\mathbf{x}) \cdot \mathbf{l}_{in} \ d\Omega_{in}$$

where $\mathbf{n}(\mathbf{x})$ is the unit normal vector of the surface at $\mathbf{x}$, and $d\Omega_{in}$ is a small solid angle centered in the direction $\mathbf{l}_{in}$.

The $\mathbf{n} \cdot \mathbf{l}_{in}$ component is needed because the radiance is defined per unit area *on the surface.* Consider a small tube centered on a ray in direction $\mathbf{l}_{in}$ which has one end at the surface point $\mathbf{x}$. If the incoming ray is oblique to the surface, then the surface will intersect the tube at an oblique angle and hence the light will be be spread out over an area proportional to $1/\mathbf{n} \cdot \mathbf{l}_{in} \geq 1$. Thus, to get the power *per unit area* on the surface, we need to weight the radiance by $\mathbf{n} \cdot \mathbf{l}_{in}$.

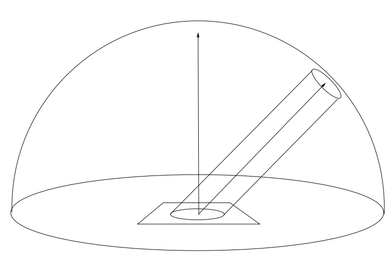

Special case: parallel source (sunlight)

Consider the special case of a "parallel source" such as sunlight. The sun spans a small solid angle and the radiance is roughly approximated by a constant over that angle, so we can write

$$E(\mathbf{x}) = \int L(\mathbf{x}, \mathbf{l}_{in})\mathbf{n}(\mathbf{x}) \cdot \mathbf{l}_{in} d\Omega_{in} \;\approx\; \mathbf{n}(\mathbf{x}) \int L(\mathbf{x}, \mathbf{l}_{in})\mathbf{l}_{in} d\Omega \approx = L_{src}\Omega_{src}\mathbf{n}(\mathbf{x}) \cdot \mathbf{l}_{src}$$

where L_{src} is the radiance of the source, Ω_{src} is the solid angle subtended by the source, and $\mathbf{l}_{src}$ is the direction to the center of the sun.

Surface reflectance

The above discussion was concerned with how much light *arrives* at a surface point $\mathbf{x}$ per unit surface area. To model the intensity of light *reflected* from a surface point $\mathbf{x}$, we must specify not just the irradiance (incoming light), but also say something about the surface material.

There are many different kinds of surface materials, of course, and these differ in visual appearance in many ways – not just color, but also how they reflect light in different directions. For example, some surfaces are shiny and others are not. Surfaces that are shiny can have very different appearances from each other: compare plastic versus metals (and each of these breaks down into many types too). Surfaces like velvet, or leather, or skin all look different from each other. We need a way of describing these.

Bidirectional reflectance distribution function (BRDF)

First suppose that we were to illuminate the surface near a point $\mathbf{x}$ from only a small set of directions centered at $\mathbf{l}_{in}$ which is a unit vector. This determines the component of the irradiance at $\mathbf{x}$ that is due to a small set of directions of incident light. We can then measure the radiance of the surface in some outgoing direction, say $\mathbf{l}_{out}$. The quantity of interest is how this outgoing radiance in direction $\mathbf{l}_{out}$ depends on the incoming radiance $\mathbf{l}_{in}$. If we know this dependence for all pairs $\mathbf{l}_{in}$ and $\mathbf{l}_{out}$, then we have characterized the reflectance properties of the surface.

To be more specific, consider the irradiance that is due rays coming from a (very small) set of directions $d\Omega_{in}$ centered in a particular direction $\mathbf{l}_{in}$,

$$L(\mathbf{x}, \mathbf{l}_{in}) \; \mathbf{n}(\mathbf{x}) \cdot \mathbf{l}_{in} \; d\Omega_{in}.$$

There is a (very small) radiance $d \; L(\mathbf{l}_{out})$ leaving the surface which is due to this incoming component. We define the ratio of these two quantities to be the *bidirectional reflectance distribution function* (BRDF):

$$f(\mathbf{x}, \mathbf{l}_{in}, \mathbf{l}_{out}) \equiv \frac{d \; L(\mathbf{l}_{out})}{L(\mathbf{x}, \mathbf{l}) \; \mathbf{n}(\mathbf{x}) \cdot \mathbf{l} \; d\Omega_{in}}$$

The units of the BRDF are sr^{-1}. Notice that the BRDF is something you can measure. (The device for measuring it is called a goniophotometer.)

It may seem odd to define a ratio of radiance to irradiance but you can see why it is done if you add up all the contributions from the light arriving at the surface from different directions $\mathbf{l}_{in}$:

$$L(x, \mathbf{l}_{out}) = \int d \; L(x, \mathbf{l}_{out}) = \int f(\mathbf{x}, \mathbf{l}_{in}, \mathbf{l}_{out}) \; L(\mathbf{x}, \mathbf{l})\mathbf{n}(\mathbf{x}) \cdot \mathbf{l} \; d\Omega_{in}$$

The form is very similar to what we had earlier for irradiance.

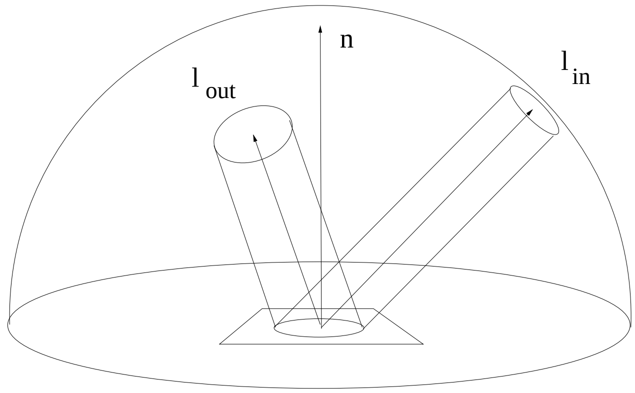

Diffuse (matte, Lambertian) reflectance

For many surfaces, the BRDF $f(\mathbf{x}, \mathbf{l}_{in}, \mathbf{l}_{out})$ is independent of $\mathbf{l}_{out}$ and $\mathbf{l}_{in}$. In particular, for such a surfaces, the radiance leaving a surface point $\mathbf{x}$ depends only on the irradiance and on some constant $f(\mathbf{x})$. It does not depend on the outgoing direction $\mathbf{l}_{out}$. The constant $f(\mathbf{x})$ specifies whether the surface is light or dark colored.

Such surfaces, for which the reflected radiance is the same in all directions, are called *Lambertian*[6], or *matte*. Many computer vision methods assume the surfaces are approximately Lambertian. One reason is that it allows you to find matching points in images taken from different viewpoints. If a pixel has a particular RGB value from from camera position and orientation, and we wish to find a corresponding pixel when the same camera (or another camera of the same "type") moved to another position and orientation, we only need to consider points that have the same RGB values. This Lambertian assumption is only useful when the the surface is indeed approximately Lambertian. Fortunately, many surfaces in the world are approximately Lambertian, at least they are approximately Lambertian over a small range of angles $\mathbf{l}_{out}$.

Of course many surface fail to be Lambertian. Exceptions are surfaces that are shiny or wet or glossy, or surfaces like leather or velvet.

Mirror reflectance

Another common type of surface reflectance is a mirror reflection. For example, a calm water surface behaves this way. Mirror reflection can be modelled as follows. Suppose light arrives at a surface point $\mathbf{x}$ from direction $\mathbf{l} = \mathbf{l}_{in}$. The light is reflected from the mirror obeys the rule: "the angle of incidence equals the angle of reflection." Let $\mathbf{r}$ be the unit vector in the direction of mirror reflection. The vectors $\mathbf{n}$, $\mathbf{l}$, $\mathbf{r}$ all lie in a plane. Using simple geometry, one can easily show that

$$\mathbf{r} = 2(\mathbf{n} \cdot \mathbf{l})\mathbf{n} - \mathbf{l} = (2\mathbf{n}\mathbf{n}^T - \mathbf{I})\mathbf{l}.$$

where $\mathbf{I}$ is the 3×3 identity matrix.

It is possible to express the BRDF of a mirror surface, but the notation would take some time to develop (and involves something called "delta functions"). For most this course, we will be assuming Lambertian surfaces only, which is the usual thing to do in computer vision.

[6] after an 18th century scientist Johann Lambert who wrote a treatise on light

Other BRDF models

"In between" the extremes of Lambertian surfaces and mirror surfaces is a large set of other surface reflectances. Let's briefly discuss what are commonly called glossy or shiny surfaces. Examples include polished surfaces such as floors and apples, oily surfaces such as skin, and plastic surfaces. Such surfaces do not obey the Lambertian reflectance model because the intensity of the light reflected from a surface point depends on where the camera is placed, i.e. from where the surface point is viewed. An example is the "highlight" or "specularity" on an apple or reflected off a polished floor tile. (You will see an example in Assignment 1 Question 2.)

Shiny surfaces are not perfect mirrors. Rather, if the light arriving at these surfaces comes from a parallel source, then the light reflected from the surface is scattered in a *cone* of directions containing the mirror direction $\mathbf{r}$. I The radiance of reflected light will typically be greatest near the mirror reflection direction, and it will fall off continuously for nearby directions. (One popular model of shiny surfaces that is often used in computer graphics is the *Phong model.* Phong's model simply assumes that the reflected radiance in direction $\mathbf{l}_{out}$ depends on the angle between $\mathbf{l}_{out}$ and the mirror reflection direction $\mathbf{r}$.)

If you wish to learn more about other BRDF models, have a look at:
`http://www.cim.mcgill.ca/~langer/PAPERS/SecretVelvetySkin.pdf`