Chapter 1

Reflection and Shading

In order to render realistic images of three dimensional objects, one needs to predict how actual material surfaces will appear when illuminated by light sources. Shading is one of the most important ways that humans detect the material of objects that we see. A metal spoon looks very differently from a plastic spoon. It is not simply a question of color. A gold colored plastic spoon still looks quite different from a spoon made out of gold.

The shading properties of a material describe how incoming light from various directions reflects out in various direction. For example, a point on a diffuse surface is brightest when the surface normal faces directly at the light and is dimmenst when the surface normal is perpendicular to the direction towards the light. Moreover, a fixed point on a diffuse surface lit by a fixed set of lights appears the same when viewed from any direction. In contrast to this, glossy surfaces have a highlight at surface points whose normals face in special directions relative to the viewer and the light directions. As a viewer moves, the highlights can disappear from some surface points and appear on other surface points.

Because shading is very affected by the surface normal at each point, it is also a feature that helps us determine the shape of an observed object. We know that the surface of an orange is bumpy because the shading patterns we see suggest that. We do not need to feel the orange, or use stereo vision (two eyes) to tell its surface shape.

In this section we will first look at simple shading algorithms that are typically used in computer graphics to simulate various materials. Then we will study how physically accurate materials and reflection can be modeled. In order to understand these physical models, we will need to carfully define the various units that are used to measure the "intensity" of light.

1.1 Shading of meshes

1.2 Bump mapping

Chapter 2

Photometry and Reflection

In the last chapter, we dealt with some simple shading models that are commonly used on computer graphics. These methods though are crude approximations to the ways that actual physical materials reflect light in the real world. When trying to make very high quality photorealistic images, it can be useful to simulate more accurately how materials behave. For example, if one is making a simulation of an unbuilt archictectural model, it may be important to see how the lit space will actually look like.

The way that an actual material property is defined is using a bidirectional reflection distribution function, or BRDF for short. A BRDF encodes how incoming light from each direction contributes to outgoing light in each direction.

In order to understand how BRDFs are defined we will need to understand how light is measured and what units are used. The concepts described in this chapter will also be essential when we discuss global illumination computations.

In this chapter, we will discuss the case of monochromatic light of a single wavelength λ. For actual lighting situations one should measure light and reflectance for each wavelenth separately. In computer graphics, this is approximated by making three measurements R, G, and B.

2.1 Units of light

There are many ways to measure light. We will discuss a variety of measurements one could do. Our goal will be to use measurements that tell us as much as we can about the actual brightness of a light without needing to convey much information about the measurement process.

The simplest thing one could do is place a sensor out in the scene and count the number of photons that hit the sensor. Such a measurement is a measurement of "energy" and is measured in Joules (J).

If I tell you how many Joules I measured, you don't know much about how bright the light was since I didn't tell you how long I kept the sensor out. There are two solutions. I could just tell you how many seconds the sensor was on, or I could divide the energy by the amount of time used in the measurement and report that number. If I divide out the time, I obtain a measurement of "power" which is measured in Watts (W= J/S). We often use the symbol Φ to denote such a measurement.

Given a mesurement of power Φ and time interval Δt , one can of course reobtain an energy measurement.

$$\text{enrgy} = \Phi \Delta t$$

Suppose that the power of the light varys over time, $\Phi(t)$, then to obtain an energy measuremnt, I simply integrate over time

$$\text{enrgy} = \int_{t_0}^{t_1} \Phi(s) dt$$

If I tell you how many Watts I measured, you don't know much about how bright the light was since I didn't tell you how large the sensor was. There are two solutions. I could just tell you how many large the sensor was, or I could divide the power by the size of the sensor and report that number. If I divide out the area, I obtain a measurement of "irradiance" which is measured in (W/m^2). We often use the symbol E to denote such a measurement.

Given a mesurement of irradiance E and the size of the sensor ΔA one can of course reobtain a power measurement.

$$\Phi = E \Delta A$$

Suppose that the irradiance of the light varys over the surface of measurement, $E(\tilde{x})$, then to obtain an power measuremnt, I simply integrate over the area

$$\Phi = \int_{\Omega} E(\tilde{x}) dA$$

Light that comes onto a surface or a sensor may be coming in from a variety of directions. I may be interested in only measuring the brightness of the light coming in from some direction $\vec{w}$. Suppose my sensor is only sensitive to light coming in from a wedge of directions centered at $\vec{w}$ with an "angular area" of Δw, then following our previous recipie, we should divide out the size of the wedge. Doing so will give us a measurement of "Irradiance" that is gives us brightness information without needing to transmit wedge size information. We often use the symbol L to denote such an irradiance measurement.

In order to make this notion precice, we need a way of measuring 3D wedges of directions just like angles measure 2D wedges. In 2D, we can measure angles in radians by intersecting the wedge with a circle, and measuring the arc length l of intersection of the circle curve and the 2D wedge. We can obtain a measurement of radians by

$$\text{radians} = l/r$$

where r is the radius of the circle.

This idea can be extended to measure 3d wedges in units of steradians. One looks at the measures the area a of the intersection of a sphere and a 3D wedge, and computes

$$\mathrm{steradians} = a/r^2$$

Using steradians, we may try to measure irradiance in units of watts per area per steradians $\frac{W}{m^s \mathrm{st}}$ by starting with an irradiance measurement and dividing out the wedge's steradians. This would give us

$$L = \frac{\Phi}{\Delta A \Delta \vec{w}}$$

and

$$\Phi = L \Delta \vec{w} \Delta A$$

This is almost correct, but we have neglected one factor. We would like the irradiance measurement L to tell us the brightness of a light traveling in some direction $\vec{w}$ to be independent of the particular orientation of the sensor. We have not yet accomplished that.

Suppose we have a beam of light carrying some power Φ. Suppose the sensor is perpendicular to the beam let; in this case let us call the area of the sensor intersected by the beam ΔA. If the same sensor is tilted by an angle of θ away rom the beam, then the same power will be measured over a sensor with the larger area $B = \Delta A / \cos(\theta)$. Where θ is the angle between the sensor surface normal and the negative of the incoming light direction $-\vec{v}$. When we divide out this larger area to obtain an irradiance measurement, we will end up with a smaller number.

The way to solve this is to use projected area

$$A_p = A \cos(\theta)$$

If we use projected area then measurement

$$L = \frac{\Phi}{\Delta A_p \Delta \vec{w}} = \frac{\Phi}{\cos(\theta) \Delta A \Delta \vec{w}}$$

will be invariant to sensor orientation. And we can obtain

$$\Phi = L \cos(\theta) \Delta A \Delta \vec{w}$$

Suppose that the radiance of the light varys over the surface of measurement and over the incoming directions , $L(\tilde{x}, \vec{w})$, then to obtain an power measuremnt, I simply integrate over the area and the incoming directions

$$\Phi = \int_{\mathrm{wedge}} \int_{\mathrm{surface}} L(\tilde{x}, \vec{w}) \cos(\theta) dw dA$$

2.1.1 Using radiance

Radiance measurements $L(\tilde{x}, \vec{w})$, are the most sensor independant way of describing the brightness of light coming it at a point from a fixed direction. As such, it is the basic measurement that is used when discussing lighting.

It can be shown that if one measures the radiance at some point $\tilde{x}$ along a direction $\vec{w}$ and obtains a measurement $L(\tilde{x}, \vec{w})$, if one measures the same ray of light at some point further along the ray $\tilde{y} = \tilde{x} + \alpha\vec{w}$, and there is no occlusion between $\tilde{x}$ and $\tilde{y}$, then $L(\tilde{x}, \vec{w}) = L(\tilde{y}, \vec{w})$. Because of this, radiance should be the way to measure the amount of light traveling along a ray.

So for example in rendering, to determine the color at a film point $\tilde{p}$, one needs the value radiance of the ray coming into $\tilde{p}$ along the direction from the pinhole $L(\tilde{x}, \vec{w})$. Becuase radiance is invariant along the ray, this is the same value as the outgoing light from the observed surface point, $L_o(\tilde{y}, \vec{w})$ where $\tilde{y}$ is the surface point seen from the pixel along the direction $-\vec{w}$. The purpose then of the shading process is to determine this outgoing radiance at the surface point. This value is a result of the reflections from incoming light L_i. In the next section we describe how L_o is computed from the L_i.

2.2 The BRDF

Now we are in a position to talk about the best way to measure reflection at some point $\tilde{x}$ on a material. We need to understand how this material scatters incoming light. To do this we must relate how a light coming in from wedge of directions centered at some direction $\vec{w}_i$ with angular width of $\Delta\vec{w}_i$ is scattered in each outgoing direction $\vec{w}_o$. How can we measure this in a way that is invariant to the experimental setup.

Since radiance is the basic way to measure light along a ray, we may first try to measure the fraction of incoming radiance to outgoing radiance

$$d(\tilde{x}, \vec{w}_i, \vec{w}_o) = \frac{L_o(\tilde{x}, \vec{w}_o)}{L_i(\tilde{x}, \vec{w}_i)}$$

The problem with this measurement is that it is not invariant to changes in the size of the incoming wedge of light $\Delta\vec{w}_i$. For example suppose I double the size of the incoming wedge of light, experimentally we usually will find that the ougoing measured radiance will double. We said “usually” because this measurement doubling will only happen if $\Delta\vec{w}_i$ is very small; so small that light from any directions in the wedge scatters the same way. By doubling this very small wedge, we will have twice as much light coming in that scatters in the same way; and our outgoing measurement will double. The exception to this rule will be perfect mirrors. With perfect mirrors, no matter how small the wedge we take, the different directions within the wedge will still act quite different. We will see how to deal with mirrors shortly.

We may get sick of radiance, and instead try to measure the ratio of incoming irradiance to outgoing irradiance

$$e(\tilde{x}, \vec{w}_i, \vec{w}_o) \quad = \quad \frac{E_o}{E_i}$$

This solvs the first problem, but creates another. If we double the size of the incoming wedge, the numerator in the measurement will double as we just saw, but so will the denominator. The problem with this measurement is that experiments show that it will usually not be invariant to the size of the outgoing light sensor. In particular, suppose that my first measurement is done over a very small wedge of directions, so small that light in all of these outgoing directions is the same, then by doubling the size of the measured wedge, we will be measuring twice as much scattered light.

The solution to these problems is to measure the ratio of incoming irradiance to outgoin radiance

$$f(\tilde{x}, \vec{w}_i, \vec{w}_o) \quad = \quad \frac{L_o(\tilde{x}, \vec{w}_o)}{E_i(\tilde{x})} = \frac{L_o(\tilde{x}, \vec{w}_o)}{\frac{\Phi_i}{\Delta A}} = \frac{L_o(\tilde{x}, \vec{w}_o)}{\frac{L_i(\tilde{x}, \vec{w}_i) cos\theta_i \Delta A \Delta w_i}{\Delta A}} = \frac{L_o(\tilde{x}, \vec{w}_o)}{L_i(\tilde{x}, \vec{w}_i) cos\theta_i \Delta w_i}$$

This solves both problems. If we double the size of the incoming wedge both the denominator and the numerator will double, and m_3 will remain unchanged. If we double the size of the outgoing light sensor, both the numerator and denominator will remain unchanged.

If we tabulate this measurement over all incoming directions and all outgoing directions, we call the resulting function the BRDF of the material. We often use the symbol $f(\tilde{x}, \vec{w}_i, \vec{w}_o)$ to represent such a function.

The BRDF at a particular point is a function of four variables, two describing the incoming direction, and two describing the outgoing direction. Most materials have somewhat simpler BRDFs that are called “isotropic”. A point on a isotropic material appears the same if one spins the surface using its normal as the rotation axis. Examples of non-isotropic materials include things like brushed metals and CDROMs.

2.3 The reflection equation

I can use a BRDF to predict the outgoing light distribution given an incoming light distributino

$$L_o(\tilde{x}, \vec{w}_o) = f(\tilde{x}, \vec{w}_i, \vec{w}_o) L_i(\tilde{x}, \vec{w}_i) cos\theta_i \Delta w_i$$

If the BRDF or the radiance of the incoming light varies over the incoming directions of light, then we simply integrate

$$L_o(\tilde{x}, \vec{w}_o) = \int_{\text{wedge}} f(\tilde{x}, \vec{w}_i, \vec{w}_o) L_i(\tilde{x}, \vec{w}_i) cos\theta_i dw_i$$

This equation is so important it has a name, the “reflection equation”.

2.3.1 The no inegral cases

We have shown above that according the measurement equation, the correct way to compute the light leaving a surface point $\tilde{p}$ along a ray direction $\vec{w}_o$ towards a camera is to perform an integral over all of the incoming light at that point. This integration is clearly an expensive computation. There are two types of cases where the integral does not need to be performed.

Mirror reflection and refraction

The first type of case that does not need an integral is a perfect mirror. For a perfect mirror, the irradiance coming out in some direction $L(\tilde{p}, \vec{w}_o)$ is completely deterimined by the irradiance of light coming in from the "bounce direction"

$$L_o(\tilde{p}, \vec{v}) = k_m L_i(\tilde{p}, \vec{b}_v)$$

k_m is a mirror coefficient, that is typically a function of the angle θ beteween $\vec{v}$ and $\vec{n}$. The bounce direction is the direction can be computed as

$$\vec{b}_v = \vec{v} - 2\vec{n}(\vec{n} \cdot \vec{v})$$

$\vec{n}$ is the surface normal.

For consistancy, if one wants to include mirror reflection using BRDFS and the integral reflection equation, one must describes the BRDF function of the mirror surface material using a "delta" function. A delta function can be thought of as an infinitely thin and tall function. An ingegral with a delta function in the integrand can be computed with a single sample.

While we are talking about perfect mirrors, we will also talk about perfect refraction. Some materials, like glass and water, let light rays pass through the material. As the ray passes through the boundary of the material, it changes directions. We will call $\vec{r}_v$ the direction of the ray of light that is refracted by the boundary and exits in the direction $\vec{v}$. In this case, the refraction is expressed as

$$L_o(\tilde{x}, \vec{v}) = k_t L_i(\tilde{x}, \vec{r}_v)$$

where k_t is the (angle dependent) refractive coefficient. Like for the perfect mirror, for perfect refraction, no integral is needed.

The ray $\vec{r}_v$ direction is computed using the following reasoning. Let us call θ_v the angle between $\vec{v}$ and $\vec{n}$ and θ_r the angle between $\vec{r}_v$ and $\vec{-n}$ Elementary optics tells us that at the boundary between two materials,

$$\frac{sin\theta_w}{sin\theta_r} = \frac{\eta_o}{\eta_i} = \text{bend}$$

where η is the index of refraction (a number associated with a material). A vacuum has $\eta = 1$ other materials have $\eta > 1$.

It can be shown that one can explicitly compute the incoming light ray direction as

$$-\vec{r}_v \quad = \quad \left(\text{bend } \ (\vec{n}\cdot\vec{v}) - \sqrt{1-\text{bend}^2(1-(\vec{n}\cdot\vec{v})^2)}\right)\vec{n} - \text{bend } \ \vec{v}$$

When a ray of light passes from air (almost a vacuum) to another medium, the ray gets refracted towards the (-) normal direction When a ray of light passes from the medium back into the air, it gets refracted away from the normal direction. Note that η is also wavelenght dependant, which is why prisms produce rainbows. In the equation above is possible for the term inside the sqaure root to be negative. This implies that that the light would be refracted more than ninety degrees away from the normal. In this case, there is no refraction.

Point light sources

Sometimes it is convenient to treat light sources as single points with no areas. This is a fiction, but a convenient one. Suppose light leaves the point and Φ power is measured by some sensor, we want to measure the density of the light, so we divide out the power by the steradians of the ougoing wedge of light. If I divide out the angles, I obtain a measurement of “radiant intensity” which is measured in (W/st). We often use the symbol I to denote such a measurement.

The reflection equation describes the outgoing radiance as proportional to the incoming irradiance. Recall that

$$f(\tilde{x},\vec{w}_i,\vec{w}_o) \quad = \quad \frac{L_o(\tilde{x},\vec{w}_o)}{E_i(\tilde{x})}$$

With a point light source, all of the irradiace comes from a single direction, and so we can avoid the integral to compute E_i, and simply compute the irradiance in closed form.

To do this, we note that from our definitions,

$$E_i(\tilde{x})\Delta A = \Phi = I_o(\vec{w}_i)\Delta w$$

And so

$$E_i(\tilde{x}) = I_o(\vec{w}_i)\frac{\Delta w}{\Delta A}$$

We can relate the steradians of a wedge and the area of a surface as

$$\frac{\Delta A\cos(\theta_i)}{d^2} = \Delta w$$

And so we see that the incoming irradiance from the point light source can be expressed as

$$E_i(\tilde{x}) = \frac{I_o(\vec{w}_i)\cos(\theta_i)}{d^2}$$

In this case, the reflection equation becomes

$$L_o(\tilde{x}, \vec{w}_o) = f(\tilde{x}, \vec{w}_i, \vec{w}_o) \frac{I(\vec{w}_i)\cos(\theta_i)}{d^2}$$

Again, for consistancy, one could express this as an integral form of the reflectance equation by using delta functions to express the incoming radiance distribution.

2.4 Mommy, where do BRDFs come from?

Now we have determined the basic units used to describe light and materials. Light travels along rays and is measured using irradiance. Light is reflected off of surfaces whose material properties are described using a BRDF. The next question is how do we find out BRDFs for various materials, and how do we represent this information.

There are three types of ways to obtain BRDFs. One way is to build a measuring device that aims light at a real surface from a multitude of directions, and measures the outgoin light from a multitude of directions. This approach is still in its infancy, but is an active area of measurement research in computer graphics. When BRDFs are obtained in this manner, they are usually represented in some type of tabular form.

The second way to obtain a BRDF is to start from physics. Physics can predict how some certain very simple material will reflect light. One then makes some assumptions about how a complex material is made up of simple materials. For exapmle, the microfacet model assumes that up close the surface really looks like perfect little mirrors facing in random directions. Given these first principles, one then can try to derive a closed form expression for reflections off certain classes of materials. Usually these expressions include adjustable parameters to describe different materials within the class.

The third way to come up with a BRDF is to work backwards. One simply makes up a closed form expression hoping to get something that acts reasonably like real materials. The phong reflecion model is such a model. This is the most common method used thusfar in computer grahpics, but there are occasions when it is more important to use realistic material models.

2.5 Simple shading models as BRDFs

Given our understanding of how light and relfection are measured, we can express some of simple shading models in this framework.

A diffuse surface has the simplest possible BRDF, namely a constant

$$f(\tilde{x}, \vec{w}_i, \vec{w}_o) = k_d$$

Thus the outgoing radiance from a diffuse surface can be measured as

$$L_o(\tilde{x}, \vec{w}_o) = \int_{\text{hemi}} k_d L_i(\tilde{x}, \vec{w}_i) cos\theta_i dw_i$$

Suppose that a diffuse surface is illuminated by a single point light source, then this equation becomes

$$L_o(\tilde{x}, \vec{w}_o) \quad = \quad k_d \frac{I(\vec{w}_i) \max(0, \cos(\theta_i))}{d^2}$$

When the light is below the surface, the term $\cos(\theta_i)$ would be negative, and L_o is zero.

The phong reflection model uses the BRDF

$$f(\tilde{x}, \vec{w}_i, \vec{w}_o) = \frac{k_p cos^n \alpha}{cos\theta_i}$$

where α is the angle between the viewing direction $\vec{w}_o$ and and the bounced light direction $\vec{b}_{w_i}$ and n is a user set parameter. If this BRDF is used with a point light source, we get

$$L_o(\tilde{x}, \vec{w}_o) \quad = \quad k_p \frac{I(\vec{w}_i) \max(0, \cos(\alpha))^n}{d^2}$$

One must also check to make sure that the light is not below the surface. In that case, L_o is zero.

In the standard computer grahics model, the $\frac{1}{d^2}$ is often excluded. One describes the surface BRDF as a combination of a diffuse and phong term. The incoming light comes in from a finite number of point light sources. Also a bogus ambiant lighting amount is added to the scene to make up for the fact that in the real world, light bounces around the environment and is comes into a surface from all directions.