

Radiometry

Steve Marschner
Cornell University
CS 5625 Spring 2020

Radiometry is a system for describing the flow of radiant energy through space. Understanding the basics is key to understanding the integrals we compute in rendering. Radiometry is essentially a geometric topic—all the subtleties have to do with geometry, not the physics of light at all.

There are a number of assumptions that let us use the model of *geometric optics* in radiometry:

- visible light (or thereabouts)
- incoherent light
- unpolarized light (random polarization)
- macroscopic scale

Together these assumptions pretty much guarantee that we will not see any effects of the wave nature of light—no interference fringes, diffraction, etc. The restriction to unpolarized light is perhaps the least valid, and it's made mainly for convenience. Including polarization in everything we do in this course is not fundamentally difficult; it's just a bunch of extra bookkeeping, and polarization is uniformly disregarded in the interactive context.

In geometric optics a convenient mental model is to think of light as being literally a flow of particles, like a bunch of BBs flying around the scene. Each particle (which we may loosely call a photon) is a point that has a direction and a wavelength associated with it. The particles stream through space in all directions, all at the same speed. Each one carries a small fixed amount of energy, and they travel in straight lines until they hit something. They do not interact with one another (they are non-interfering).

You can think of the photodetectors in a digital camera or in the human eye quite literally as photon counters. So ultimately the quantity we want to simulate is simply a number of particles that make their way to a particular pixel.

This particle count represents an amount of energy. What radiometry is about is measuring the energy in different regions of space, direction, and time.


Aside: solid angles

One unit of measurement that figures prominently in radiometry is solid angle. This is a generalization of the notion of an angle from 2D to 3D space. In 2D one can think of directions (unit vectors) as being points on the unit circle, and in 3D one can think of directions as being points on the unit sphere. The definition of solid angle is by analogy between these two ideas:

An angle is a section of the unit circle; the magnitude (size) of the angle is its arc length:



A solid angle is a section of the unit sphere; the magnitude of the solid angle is its area:



The idea of an angle as a set of directions may be new. In 2D the distinction between an angle (the set) and the angle's magnitude (the number) is a fine one, because as long as the set is connected it is just an interval on the circle. In 3D, the question of the shape of a solid angle is much more interesting, because there can be all sorts of regions on the sphere that have the same area.

In 2D geometry one talks about the angle *subtended* by something at another point: for example, the angle at a vertex of a triangle is the angle subtended by the opposite side of the triangle. Similarly in 3D the solid angle subtended by an object is the set of directions that point toward that object.

1. Radiometric units, top down

Energy Q , measured in Joules (J)

- this is what you can actually measure with some detector: how many particles hit it.

Power, or flux: P or Φ , watts ($W = J/s$)

- energy per unit time $P = dQ/dt$
- synonyms: radiant flux, Φ

In steady state (which we normally assume in graphics) energy and power are basically interchangeable, so we are sometimes a bit sloppy in distinguishing them. When we make a power measurement with some real detector, we get energy captured over some time interval (the exposure). This is an estimate of the instantaneous power: it's the average power over the exposure.

Power can be thought of as a derivative of energy: dQ/dt . Power is the rate of increase of energy as you sit and let light fall on your detector. You often see this written as $dQ = P dt$, though I try to avoid this.

Really, radiant flux is the starting point for radiometry; all other radiometric quantities are densities of power.

Area density of power: irradiance E

- power per unit area $E = d\Phi/dA$ (measured at a point on a surface)

The term irradiance is used for light arriving at a surface; “radiant exitance” is the less often used term for the power-per-unit-area leaving a surface. The surface can be real or imagined, but one needs a surface to talk about irradiance.

Solid angle density of power: intensity I (W/sr)


- power per unit solid angle $I = d\Phi/d\omega$ (measured in a direction)

This concept, in itself, is mostly used with point sources—that is, emitting objects so small that we don’t care about *where* on the object light is coming from; just what direction it goes in. However, it can also be used to describe light arriving at a small object, when we care only about the total amount of light arriving from a given direction, not about *where* it arrives on the object.

Density of power wrt. solid angle and area: Radiance L (W/m²sr)

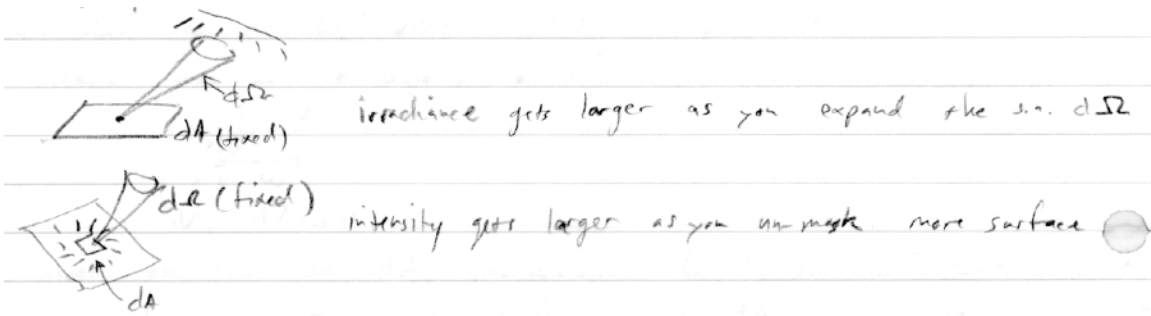
- flux per unit solid angle per unit area $L = d^2\Phi/(dAd\omega)$ (measured in a particular direction at a particular point on a surface or in space).

$L(\mathbf{x}, \omega)$ is a measure of the density of photons passing near \mathbf{x} and traveling in directions near ω . The details of “near” are the key to the definition—this is a geometric question. Radiance measures the radiant flux that would be collected if we look at photons whose directions lie within a solid angle $d\omega$ around ω and land on an area of surface dA at \mathbf{x} that is perpendicular to the direction ω .



The derivative interpretation: if you enlarge dA or widen $d\omega$ you will collect more light, and L is the constant of proportionality.

Radiance is a second derivative: It is irradiance per unit solid angle, or it is intensity per unit area.



Radiance invariance

Radiance is very useful because it is *conserved along lines* through empty space. That is, if we look at the radiance at two points \mathbf{x}_1 and \mathbf{x}_2 , in the direction ω that points from one to the other, we will measure the same radiance at both points: $L(\mathbf{x}_1, \omega) = L(\mathbf{x}_2, \omega)$. To see why this is, we can measure radiance at each point. We can make this measurement using whatever solid angles and areas we want, so let's choose two areas dA_1 and dA_2 that are perpendicular to ω :



Think of the set of lines connecting these two patches—that is, the set of lines that pass through both rectangles. Along these lines travel all the photons that pass first through dA_1 and then through dA_2 . To measure radiance at each end, we need to choose solid angles. The trick is to choose the solid angle that exactly corresponds to all directions that pass through the other surface. We define $d\omega_1$ as the solid angle subtended by dA_2 from dA_1 , and likewise $d\omega_2$ is the solid angle subtended by dA_1 from dA_2 . (You may notice that these solid angles are different for different points on the surfaces. This is OK because in the limit of small areas the solid angles are all the same.) The sizes of these solid angles are

$$|d\omega_1| = dA_2/r^2$$

$$|d\omega_2| = dA_1/r^2$$

Now let $d^2\Phi$ be the flux between the surfaces—that is, the energy carried by all the particles that pass first through dA_1 and then through dA_2 . Because of how we arranged the solid angles for measuring radiance, $d^2\Phi$ is the power to be used for measuring radiance on both ends:

$$L_1 = \frac{d^2\Phi}{dA_1 |d\omega_1|} \quad L_2 = \frac{d^2\Phi}{dA_2 |d\omega_2|}$$

If we then substitute in the expressions for the sizes of the two solid angles, we end up with


$$L_1 = L_2 = \frac{r^2 d^2 \Phi}{dA_1 dA_2}$$

This conservation of radiance means that, in the absence of anything that gets in the way, radiance is really more a property of *lines* in space than a property of (point, direction) pairs. Finding that $dA_1 |d\omega_1| = dA_2 |d\omega_2|$ is really just saying that solid angle times area is a well-defined measure on the space of lines: we can measure the same set of lines in two different places and get the same answer (see camera example at end).

Projected area and projected solid angle

Our definitions of radiance are for surface areas perpendicular to the direction in question. What about when we need to work with some other surface, one that's not perpendicular?

The simple answer: a surface that's inclined at an angle θ to ω has the same particles passing through it in the direction ω as a perpendicular surface whose area is smaller by a factor of $\cos \theta$.



If we call the inclined surface area dA' and continue to call the perpendicular surface area dA , then $dA = dA' \cos \theta$. So the radiance can be written

$$L(\mathbf{x}, \omega) = \frac{d^2 \Phi}{dA d\omega} = \frac{d^2 \Phi}{dA' \cos \theta d\omega}$$

The presence of the cosine is quite an inconvenience, and hiding it inside a unit of measure is appealing. The two ways of doing this are to group $(dA' \cos \theta)$ and call it the “projected area” of dA' in the direction ω or to group $(d\omega \cos \theta)$ and call it the “projected solid angle.” One often hears radiance defined as “flux per unit solid angle per unit projected area” (or with the “projected” in front of “solid angle”), but really, this “projected” business is just saying that the area for radiance is measured on a surface perpendicular to the direction of propagation.

Projected area is a fairly intuitive idea: it's the area of the projection of a surface perpendicular to a direction. Projected solid angle is a little less familiar, but it's similarly a way of measuring something (a solid angle) after projecting it perpendicular to a direction. (See the Nusselt analog below). Plain old solid angle is a measure for regions of the sphere which assigns measure by area; projected solid angle is the same only it counts areas near the equator less, and areas near the poles more.

I find projected solid angle to be a very handy notational tool, though you might prefer to keep things explicit by leaving the $\cos \theta$ where you can see it..

2. Radiometric units, bottom up

One can also define all these quantities by starting with radiance as the basic concept and thinking of everything else as integrals of radiance. From its definition it is clear that we can find power from radiance by integrating over an area and a solid angle. For instance, to find the total power exiting from a surface S :

$$\Phi = \int_S \int_{\Omega} L(\mathbf{x}, \omega) \langle \omega, \mathbf{n} \rangle d\sigma(\omega) dA(\mathbf{x})$$


The term $\langle \omega, \mathbf{n} \rangle$ is the cosine factor to allow for surfaces that are not perpendicular to ω . The notation $d\sigma(\omega)$ means that the integration variable is ω and the measure we are using is σ , my name for the solid angle measure on the direction sphere. Similarly $dA(\mathbf{x})$ means that \mathbf{x} is the integration variable and area is the measure. The integration domain Ω is the hemisphere of directions that have positive dot product with the surface normal \mathbf{n} , and we integrate radiance in the ω direction because we're interested in radiance leaving the surface. Integrating $L(\mathbf{x}, -\omega)$ would give the total power falling on the surface.

If the surface S is not flat, then \mathbf{n} is a function of \mathbf{x} , and so is the hemisphere Ω .

The notation of this integral can be simplified by using projected solid angle. I define a new measure μ , called “projected solid angle measure,” by giving its relationship to σ for infinitesimal sets:

$$\mu(d\omega) = \sigma(d\omega) \langle \omega, \mathbf{n} \rangle$$

There is a nifty geometric way of thinking about this, known as the “Nusselt analog”: for any set D in a hemisphere, μD is the area of the shadow of D on the hemisphere’s base plane:



Using projected solid angle measure we can write

$$\Phi = \int_S \int_{\Omega} L(\mathbf{x}, \omega) d\mu(\omega) dA(\mathbf{x})$$

or even just

$$\Phi = \int_{S \times \Omega} L d\mu dA$$

Doing only the solid angle integral gives irradiance (with $-\omega$) or radiant exitance (with $+\omega$):

$$E(\mathbf{x}) = \int_{\Omega} L(\mathbf{x}, -\omega) d\mu(\omega)$$

$$M(\mathbf{x}) = \int_{\Omega} L(\mathbf{x}, \omega) d\mu(\omega)$$

Doing only the area integral gives intensity:


$$I(\omega) = \int_S L(\mathbf{x}, \omega) \langle \omega, \mathbf{n} \rangle dA(\mathbf{x})$$

We could invent a “projected area” measure in which to hide the cosine factor, but since this computation is not nearly so common as irradiance, I won’t bother.


3. Examples

Detectors for various quantities

For some of these quantities it is helpful to think of a detector. For flux, think of a large-area detector (like a photodiode) that just collects all the photons that land on it:




For irradiance, think of a piece of photographic film. At each point it records a signal that has to do with the density of photons landing near that point, but it does not care about where they come from:



Digital image sensors achieve the same sort of measurement with an array of very small flux meters.


For radiance, the canonical example is a camera or the eye. If we look at a little area in the image (a pixel, say), that collects light arriving at the lens from some solid angle. The aperture of the

camera defines a small area. So the pixel value is an estimate of flux per unit solid angle (in the pixel's ray direction) per unit area (at the aperture's position):



Reflection from perfect diffuser

A perfect diffusing reflector, known as a Lambertian surface, is a surface that reflects a fraction $0 \leq R < 1$ of the incoming irradiance (that is, the radiant exitance M_r is R times the irradiance E_i), with uniform radiance L_r in all directions:



What is the radiance? We know it's related to the radiant exitance M :

$$M_r = \int_{H^2} L_r(\omega) d\mu(\omega)$$

Since radiance is constant this is just $M_r = L_r \mu(H^2) = \pi L_r$ (use Nusselt analog to see that $\mu(H^2) = \pi$). This means that


$$L_r = \frac{R}{\pi} E_i$$

This relationship between reflected radiance and incident irradiance will come up again later when we discuss the BRDF.


Area and solid angle in a camera

Here is an example of radiance invariance. Suppose I have a camera pointed at a wall 1 meter away, and the camera's aperture has an area $A_1 = 1 \text{ cm}^2$. Let's say a single pixel maps to an area $A_2 = 1 \text{ mm}^2$ on the wall. From the camera's point of view the pixel's solid angle is 10^{-6} sr (remember area over squared distance). From the wall, the camera's aperture subtends a solid angle of 10^{-4} sr . If the wall's radiance is L , the power from A_2 that hits the camera's aperture is $\Phi = L$

$\times 10^{-6} \text{ m}^2 \times 10^{-4} \text{ sr}$. Interpreting this as a radiance measurement at the camera, we get $\Phi / 10^{-4} \text{ m}^2 / 10^{-6} \text{ sr}$ which is L again.




If we point the camera out the window at a mountainside 1 km away, the pixel maps to an area of 1 m^2 but the solid angle of the camera from the mountain is only 10^{-10} sr and if the mountain has the same radiance as the wall we'll end up with the same power coming into the camera.



Radiance falloff in a pinhole camera

Suppose I make a box with a small hole in one face and a piece of film on the opposite face. This is a pinhole camera. If we point this camera at a uniform radiance L (say, the overcast sky), what is the irradiance on the film?



By thinking of a very flat box (say a shirt box with the pinhole in the center of the top) it's obvious that the irradiance is not uniform: it's bright near the pinhole and dimmer farther away. Furthermore it has to be radially symmetric, so the irradiance is a function of the angle at which the light strikes the film; call that α . Call the distance from the pinhole to the film d and the area of the aperture dA_a .

If we stand on the film, we see uniform radiance coming from the solid angle subtended by the aperture. This solid angle is

$$\sigma(d\omega_f) = \frac{dA_a \cos \alpha}{r^2} = \frac{dA_a \cos^3 \alpha}{d^2}$$


The irradiance is L times $\mu(d\omega_f)$, which is

$$E_f = \frac{L dA_a \cos^4 \alpha}{d^2}$$


So the irradiance is proportional to the square of the aperture (no surprise there), drops off with the square of the pinhole distance (no surprise there), and also drops off toward the corners of the image as the fourth power of $\cos \alpha$. This turns out also to be true of a camera made with an ideal lens.

Resolved and unresolved objects in images

When we point a camera at an object, we think of the camera as recording that object's radiance. This means that if we move the camera closer to or farther from the object the pixel value stays the same. This is only true if the radiance is constant over the pixel—that is, if the object is big enough to be *resolved* by the pixel grid:




If the object is small and/or far away (a droplet of water, a star, etc.) so that it falls entirely within one pixel, then it behaves a bit differently. It's no longer possible to tell the radiance from the pixel value, and the pixel value doesn't stay the same with distance. Instead, the value is related to the intensity of the object, not the radiance, and it falls off with distance squared:



Solid angle of a spherical cap

Consider a spherical cap of angular radius α .



What is its solid angle? This is an easy integration problem in spherical coordinates if we align the pole with the center of the cap:

$$\int_0^{2\pi} \int_0^\alpha \sin \theta \, d\theta \, d\phi = 2\pi(1 - \cos \alpha).$$

Here I am using the conventional spherical coordinates for graphics: θ counts up from zero at the pole; ϕ counts up counterclockwise around the equator.


It's interesting to note that $1 - \cos \alpha$ is the length of the projection of the cap onto the vertical axis in the drawing. The same thing works for any band $\alpha < \theta < \beta$ on the sphere.

What is the projected solid angle of the spherical cap? You can do this by integrating $(\sin \theta \cos \theta \, d\theta \, d\phi)$, which is not hard.

Or you can do it in your head with the Nusselt analog. The projection of the spherical cap onto the equatorial plane is a circle of radius $\sin \theta$. The area of that circle is $\pi \sin^2 \theta$, which is the projected solid angle of the spherical cap.

Irradiance of a disc-shaped source

Suppose we have a disc-shaped light source, of radius R and constant radiance L . What is the irradiance due to that source on a surface facing the source at a distance r ?



The answer is the integral of the source's radiance over its subtense on the directional hemisphere at \mathbf{x} , with respect to projected solid angle. Since the radiance is constant, this is L times the projected solid angle of the spherical cap subtended by the source. Using the result of the previous example, the irradiance is

$$E(\mathbf{x}, \mathbf{n}) = \pi L \frac{R^2}{r^2 + R^2}$$