



Volume light transport

Steve Marschner
Cornell University
CS 6630 Spring 2012, 1–6 March

So far we have always assumed that, other than surfaces, there is nothing that impedes the flow of light. This means radiance is conserved along lines—what we called the radiance invariance law. As you recall, this was proved using an argument about this picture:



and it hinged on the fluxes through the two areas, restricted to the mutually subtended solid angles, being the same. If there is some stuff in the way, this may not be true—particles that were headed from dA_1 to the second might not make it to dA_2 , or particles that did not come from dA_1 might nevertheless show up at dA_2 . There are four ways this can happen, illustrated here:



1. A particle may disappear, leaving no trace: *absorption*.
2. A particle may spontaneously appear out of thin air: *emission*.
3. A particle that was headed for dA_1 may be deflected somewhere else: *outscattering*.
4. A particle that was just passing by might be deflected towards dA_2 : *inscattering*.

These four kinds of events lead to four terms in the *Radiative Transfer Equation*, a widely used model for the behavior of light in an interacting medium. The equations proceed from arguments about what happens to radiance as we move along a ray—in what way the radiance fails to be


conserved when it interacts with the medium. You can think of this as being like a flashlight beam or laser beam, shining through the fog.

The RTE is a macroscale, empirical model that has proved extremely useful for a variety of problems, though it is not a physical theory derived by first principles from Maxwell's equations. Rather, it describes the behavior of incoherent light on a large scale, and is derived from intuitive arguments about the interaction of light with small particles suspended in the medium—for instance, water droplets in a cloud or particles of ash and soot in smoke.

Although the whole theory rests on the idea of small, well-separated scattering particles, it nonetheless proves to be a good model in many other situations: in tissue, for instance, where scattering is caused by tightly packed semi-transparent anatomical structures; or in turbulent heated air that has varying refractive index but no particles. Generally, as long as the structures causing scattering are small and randomly dispersed, we can usually expect the RTE to be a good model, but if scattering is not by particles then we should be on the lookout for discrepancies between the model and reality.

Emission

Beginning with the simplest case: a medium that just emits light (think of gas in a neon tube, or the liquid in a glow stick). What is the effect on the radiance traveling along a ray through the medium?



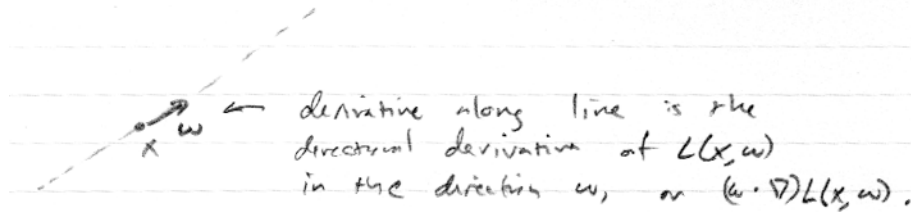
If we look perpendicularly at a unit-thickness slab of this medium, we will see some emitted radiance ϵ , in addition to the light L_i that travels, unimpeded, through the slab. If we stack two slabs together, or alternatively make the slab 2 units thick, we see the transmitted radiance $L_i + 2\epsilon$. It's easy to see the general rule: $L_t = L_i + d\epsilon$, where d is the distance the ray travels through the medium.

Another way to write this is as a differential equation:

$$\frac{dL}{ds}(s) = \epsilon(s), \quad \text{or,} \quad L'(s) = \epsilon(s)$$

The nice thing is that this generalizes to non-constant ϵ , by letting ϵ depend on s . This is a statement about radiance as a function of distance along a particular ray. To turn this into a statement about the 3D radiance distribution in the volume, we just have to recognize $L(s)$ as a slice of the broader distribution:


$$L(s) = L(\mathbf{x}(s), \omega)$$




The derivative of $L(s)$ is a directional derivative of $L(\mathbf{x}, \omega)$, so the effect of emission on the volume's radiance distribution can be written:

$$(\omega \cdot \nabla)L(\mathbf{x}, \omega) = \varepsilon(\mathbf{x})$$

The radiance in any direction is increasing as you move in that direction, at a rate proportional to ε .



How to interpret the quantity ε , physically? It seems logical to measure emission in terms of power per unit volume. So suppose our medium emits a power density Q_P (W/m³). If we look at our beam, with cross sectional area A and solid angle $d\omega$, traveling through the medium we can chop out a cylindrical chunk and see how much radiance it contributes to the beam.




The indicated volume ΔV emits power ($Q_P A \Delta s$), but only a fraction $d\omega/(4\pi)$ winds up in the solid angle $d\omega$. So the power increases by $\Delta\Phi = \rho_P A \Delta s d\omega / 4\pi$, and if we measure radiance using the throughput $A d\omega$, then $L' = Q_P / 4\pi = \varepsilon$. The units of ε , which as the derivative of radiance with respect to length have to be radiance per meter, or W / (m³ sr), can also be thought of as a distribution of power over *volume* and solid angle (hence the 4π in the denominator).

Absorption

Now suppose we have a medium that contains randomly distributed absorbing particles (think of soot—they are completely black). Any photons that hit these particles will disappear from the

beam. An absorbing-only medium is transparent (not cloudy) but dims images seen through it. Examples: sunglasses, wine, gemstones.

Using the same kind of argument as with emission, let's look at how absorption changes radiance with distance.



Absorption decreases radiance by an amount proportional to the radiance that is already there. Writing this as a differential equation,

$$L'(s) = -\sigma_a(s)L(s)$$

The coefficient σ_a is known as the absorption coefficient, or the absorption cross-section. Its units are inverse distance (because multiplying by length is supposed to give a unitless quantity).

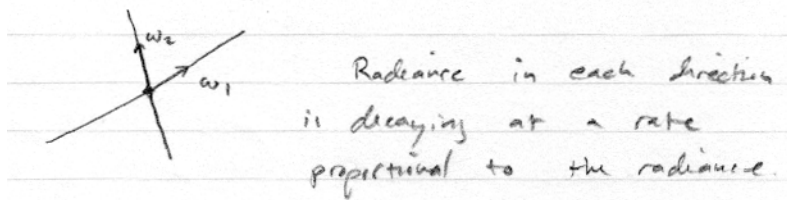
If σ_a is constant, then the solution to this equation is

$$L(s) = e^{-\sigma_a s} L_0$$


where L_0 is the radiance at $s = 0$. As with emission we can reinterpret this single-variable equation as a statement about the 3D distribution:

$$(\omega \cdot \nabla)L(x, \omega) = -\sigma_a L(x, \omega)$$


So the radiance along any line is decaying exponentially—but the radiance in different directions at a single point is still completely separate.



A laser beam:



Looking more closely at the physical meaning of absorption cross-section: Just as with emission think of a slice of a cylindrical beam, this time filled with randomly scattered absorbing particles at a density of ρ particles per unit volume. Let the size of the particles, measured as a cross-sectional area (that is, projected area), be A_p . As viewed from the direction of propagation, these particles block a fraction of the area of the beam.



There are $n = \rho A \Delta s$ particles, each blocking a projected area A_p , so the fraction of area blocked is

$$\frac{\rho A \Delta s A_p}{A}$$

and the radiance transmitted through the interval Δs is

$$L(s + \Delta s) = (1 - \rho \Delta s A_p) L(s), \quad \text{hence} \quad L'(s) = -\rho A_p L(s).$$

The product ρA_p , which has units of (1 / volume) (area), or (1 / length), is the absorption coefficient σ_a . Some authors like to keep the density and size of particles separate, but they only enter into the equations as a product, so I like to just have the coefficient. Also, remember that even though we derive this theory based on particles, we like to apply it to substances (e.g. tissue) that are not made of particles. In these cases ρ and A_p make no sense but σ_a still does.

Integral form of the equation so far

Before we move on to scattering, let's solve this differential equation to find the radiance distribution in integral form. That is, we will write the distribution resulting from given emission and absorption distributions, and a given incident light distribution from outside the medium, as an integral involving these known quantities. You could call this the "emission-absorption rendering integral." It's interesting to do this case first because it is exactly the same analysis that leads to the full equation of transfer—an integral equation, not just an integral—but it looks simpler.


If we look at the 1D equation for emission along a beam:

$$L'(s) = \varepsilon(s)$$

the solution is just an integral of the emission:

$$L(s) = \int_0^s \varepsilon(s') ds' + C$$

and in the 3D context it looks like this:



$$L(\mathbf{x}, \omega) = \int_{\mathbf{y}}^{\mathbf{x}} \varepsilon(\mathbf{x}') d\mathbf{x}' + L(\mathbf{y}, \omega)$$

Here I am using a slight abuse of integral notation: the integral from \mathbf{y} to \mathbf{x} is a line integral along the line segment connecting \mathbf{y} and \mathbf{x} , and \mathbf{x}' is the integration variable. The point \mathbf{y} is the place along the ray leading to (\mathbf{x}, ω) where we know the radiance—probably on a surface, or at the edge of the region that contains the scattering medium.

This equation shows that radiance at any point in an emitting medium is simply the line integral of emission along the straight-line path leading up to (\mathbf{x}, ω) .

For absorption alone, the differential equation is

$$L'(s) = -\sigma_a L(s), \quad \text{or} \quad L'(s)/L(s) = -\sigma_a$$

which is solved by integrating:

$$\int^s \frac{L(s')}{L(s')} ds' = - \int^s \sigma_a(s') ds' + C$$

$$\ln L(s) = - \int^s \sigma_a(s') ds' + C$$

$$L(s) = K e^{-\int^s \sigma_a(s') ds'}$$

Putting this into the context of the 3D distribution and filling in the boundary condition, as before, from the end of the ray:

$$L(\mathbf{x}, \omega) = e^{-\int_{\mathbf{y}}^{\mathbf{x}} \sigma_a(\mathbf{x}') d\mathbf{x}'} L(\mathbf{y}, \omega)$$

So here, as with emission, the radiance depends on the line integral of a volume property along the straight-line path leading up to (\mathbf{x}, ω) . The difference is the $\exp(- \dots)$.


The exponentiated integral above is important enough to be worth making a special name for; I will call it the total absorption between \mathbf{x} and \mathbf{y} :

$$\alpha(\mathbf{x}, \mathbf{y}) = \exp\left(-\int_{\mathbf{x}}^{\mathbf{y}} \sigma_a\right) = \alpha(\mathbf{y}, \mathbf{x})$$

Note that the order does not matter in computing the absorption along a path. This seems wrong at first, but it is true. (It is not true if emission is mixed in.)

Example: emitting blob

Consider a blob of uniformly, isotropically emitting medium. It does not matter what shape the blob has, but its volume is V . What is its intensity, considered as a point emitter?



Answer: it is an isotropic source with intensity ϵV . This can be seen by integrating radiance over the projected area to get intensity. The radiance at each point is a line integral through the blob, and all the line integrals together constitute a volume integral over V .

Example: striped absorber

Consider a unit cube of absorbing medium, with a homogeneous absorption coefficient of 5. Radiance, and hence total intensity, for illumination parallel to any axis is attenuated by $\exp(-5)$, or about $6.7e-3$.

Now consider the same cube but striped with stacked plates of absorber, 1 by 1 by 0.1 units, alternating between $\sigma_a = 1$ and $\sigma_a = 9$. This is the same amount of absorbing material, just distributed differently. It still attenuates by $\exp(-5)$ along the z direction, but along x and y , if we ask about the average attenuation, or attenuation of total intensity across the whole cube, we get the average of $\exp(-1)$ and $\exp(-9)$ which is about 0.18. So just by rearranging the structure we greatly change the absorption and it becomes anisotropic.

Integral for emission and absorption together

By adding the equations for emission and absorption we obtain

$$(\omega \cdot \nabla)L(\mathbf{x}, \omega) = -\sigma_a(\mathbf{x})L(\mathbf{x}, \omega) + \epsilon(\mathbf{x})$$

This is now an inhomogeneous, but still linear, differential equation, which is of the form

$$y' + py = q$$

which is a standard form with well known solution

$$y(x) = \int^x q(t)e^{-\int_t^x p} dt + C_1 e^{-\int^x p}$$


Pattern matching and filling in p and q results in:

$$L(s) = \int_0^s \epsilon(s')e^{-\int_{s'}^s \sigma_a} ds' + L(0)e^{-\int_0^s \sigma_a}$$

Mapping this solution into 3D gives the α - ϵ rendering integral:

$$L(\mathbf{x}, \omega) = \int_{\mathbf{y}}^{\mathbf{x}} \epsilon(\mathbf{x}')\alpha(\mathbf{x}, \mathbf{x}')d\mathbf{x}' + L(\mathbf{y}, \omega)\alpha(\mathbf{x}, \mathbf{y})$$

The solution may have seemed mysterious in the abstract, but mapped into this context it is quite intuitive: light emitted at points \mathbf{x}' along the ray contributes to the radiance at \mathbf{x} , but it is attenuated by absorption along the line segment from \mathbf{x}' to \mathbf{x} . The light from outside the medium contributes in the same way as before.



This second term is sometimes called the “reduced (source) intensity.”

Outscattering

Outscattering is the loss of photons, not by their disappearance, but by their deflection into directions outside the $d\omega$ that we are paying attention to as we follow radiance through the volume. However, the effect on radiance along the line is the same as absorption: decay of radiance at a rate proportional to the radiance:

$$(\boldsymbol{\omega} \cdot \nabla)L(\mathbf{x}, \boldsymbol{\omega}) = -\sigma_s(\mathbf{x})L(\mathbf{x}, \boldsymbol{\omega})$$

The coefficient σ_s is called the scattering cross section, or scattering coefficient. Like absorption it has units of inverse distance, and it can be understood intuitively by thinking of small particles that block part of the light traveling along a beam. The only difference is where the light goes—scattered light contributes to light traveling in other directions, whereas absorbed light just goes away.

In class we saw a demonstration comparing the effects of scattering (due to a dilute solution of whole milk) and absorption (due to a dilute solution of india ink). Although the two beakers appeared very different (one white, the other dingy gray), the (nearly) collimated light that was projected by the overhead projector onto the screen was very similar: both beakers appeared as brown circles that got darker the more scattering material was present.

Inscattering

Light that is removed from one direction by outscattering will show up traveling in other directions. This acts just like emission but it depends on the other light at the same point.

Simple case: isotropic scattering, in which light that is scattered goes equally in all directions. That is, it contributes equally to all paths through the point where it scattered. This means that the inscattering at a point—the scattered power that shows up in our $d\boldsymbol{\omega}$ —depends equally on the light traveling at all directions at a given point: it is a function of the fluence (aka. scalar irradiance).

As with emission and absorption, let's look at a section of beam of cross-section A and length Δs . When we look the particles (of volume density ρ) in this volume that are illuminated, each collects an amount of power equal to its cross sectional area times the fluence: $A_p\phi(\mathbf{x})$. (As before we assume they are sparse enough not to shadow one another locally.) The total power is proportional to the number of particles, which is $\rho A\Delta s$. This power gets distributed uniformly to all directions, so it looks exactly like volume emission with coefficient $\rho A_p\phi(\mathbf{x}) / 4\pi$. The quantity ρA_p is the scattering coefficient, σ_s . Substituting in the definition of fluence, the equation along the ray is

$$(\boldsymbol{\omega} \cdot \nabla)L(\mathbf{x}, \boldsymbol{\omega}) = \frac{\sigma_s}{4\pi} \int_{4\pi} L(\mathbf{x}, \boldsymbol{\omega}') d\sigma(\boldsymbol{\omega}')$$

(I will start writing just $d\boldsymbol{\omega}'$ for integrals over solid angle in the volume context, since projected solid angle is only used with surfaces.) The rate of radiance increase due to scattering depends on all the other light at the point \mathbf{x} , considered as a function of direction. In this simple isotropic case it only depends on the *total* amount of light at \mathbf{x} , the fluence, but in reality some directions will contribute more and some less. Isotropic scattering is the analog of ideal diffuse reflection, and $(1/4\pi)$ is the analog of a constant (Lambertian) BRDF like (R/π) . To describe anisotropic scattering we need a different weighting function for the directional radiance at \mathbf{x} , that can tell us how likely it is that a scattered particle that was traveling in direction $\boldsymbol{\omega}$ will be deflected to the direction $\boldsymbol{\omega}'$. This function is called the *phase function*, f_p , and it is the analog of the BRDF for

surfaces. (The name comes from the function that tell you how bright a celestial object, for example the moon, appears as a function of sun angle—a function of the phase of the moon.)

With the phase function replacing the constant weighting for radiance, we have the standard form:

$$(\boldsymbol{\omega} \cdot \nabla)L(\mathbf{x}, \boldsymbol{\omega}) = \sigma_s \int_{4\pi} f_p(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}')L(\mathbf{x}, \boldsymbol{\omega}')d\boldsymbol{\omega}'$$

Here the phase function is always a normalized probability distribution, given the value of one argument:

$$\int_{4\pi} f_p(\boldsymbol{\omega}, \boldsymbol{\omega}')d\boldsymbol{\omega}' = 1$$

Note that this is different from the BRDF, which is not normalized: a surface that absorbs some light will have a BRDF with an integral less than 1. This is really just a design decision in the math: separating absorption and scattering into two separate terms, so that light that scatters has to show up somewhere, is convenient for analysis. It is perfectly possible to formulate the same equation with a non-normalized phase function that describes absorption and scattering together.

The Radiative Transfer Equation

Now we have derived everything we need to get the main equations of volume scattering; we just need to plug it all together. Adding up the four terms (emission, absorption, out-, and inscattering):

$$(\boldsymbol{\omega} \cdot \nabla)L(\mathbf{x}, \boldsymbol{\omega}) = \varepsilon(\mathbf{x}) - \sigma_a(\mathbf{x})L(\mathbf{x}, \boldsymbol{\omega}) - \sigma_s(\mathbf{x})L(\mathbf{x}, \boldsymbol{\omega}) + \sigma_s(\mathbf{x}) \int_{4\pi} f_p(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}')L(\mathbf{x}, \boldsymbol{\omega}')d\boldsymbol{\omega}'$$

Following convention in moving the negative terms across the equals sign and grouping outscattering and absorption together with the coefficient $\sigma_t = \sigma_a + \sigma_s$, we obtain:

$$(\boldsymbol{\omega} \cdot \nabla)L(\mathbf{x}, \boldsymbol{\omega}) + \sigma_t(\mathbf{x})L(\mathbf{x}, \boldsymbol{\omega}) = \varepsilon(\mathbf{x}) + \sigma_s(\mathbf{x}) \int_{4\pi} f_p(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}')L(\mathbf{x}, \boldsymbol{\omega}')d\boldsymbol{\omega}'$$

This is the Equation of Transfer, or the Radiative Transfer Equation. It is an “integro-differential equation” for the radiance in the volume, so called because it is a differential equation but the terms in the equation involve integrals of the unknown function.

The Volume Rendering Equation

The derivation that led to the rendering integral for absorption and emission is completely agnostic to the difference between absorption and outscattering, or emission and inscattering. The effects on the radiance along a single ray are the same, so the same linear, inhomogeneous differential equation applies. Again we have the form

$$y' + py = q$$

and this time we have

$$p(s) = \sigma_t(s)L(s, \omega)$$

$$q(s) = \varepsilon(\mathbf{x}) + \sigma_s(\mathbf{x}) \int_{4\pi} f_p(\mathbf{x}, \omega, \omega') L(\mathbf{x}, \omega') d\omega'$$

so, following the same solution as before, and proceeding directly to 3D to avoid some lengthy equations:

$$L(\mathbf{x}, \omega) = \int_{\mathbf{y}}^{\mathbf{x}} \tau(\mathbf{x}', \mathbf{x}) \left(\varepsilon(\mathbf{x}') + \sigma_s(\mathbf{x}') \int_{4\pi} f_p(\mathbf{x}', \omega, \omega') L(\mathbf{x}', \omega') d\omega' \right) d\mathbf{x}' + \tau(\mathbf{x}, \mathbf{y}) L(\mathbf{y}, \omega)$$

This is now just an integral equation, no longer an integro-differential one. Its form is actually very similar to the familiar rendering equation for surfaces. Rearranging this equation a bit to reveal its form as an integral equation:

$$L(\mathbf{x}, \omega) = \tau(\mathbf{x}, \mathbf{y}) L(\mathbf{y}, \omega) + \int_{\mathbf{y}}^{\mathbf{x}} \tau(\mathbf{x}', \mathbf{x}) \varepsilon(\mathbf{x}') d\mathbf{x}' + \int_{\mathbf{y}}^{\mathbf{x}} \tau(\mathbf{x}', \mathbf{x}) \sigma_s(\mathbf{x}') \int_{4\pi} f_p(\mathbf{x}', \omega, \omega') L(\mathbf{x}', \omega') d\omega' d\mathbf{x}'$$

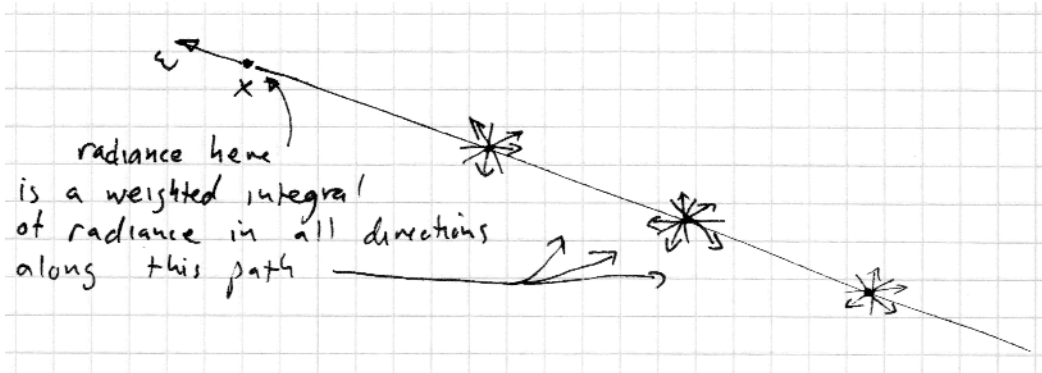
The first two terms we can call the “source term”—they are simply integrals of known quantities, and do not depend on the unknown radiance $L(\mathbf{x}, \omega)$. The other term can be seen as just a weighted integral of the unknown function:

$$L(\mathbf{x}, \omega) = L^0(\mathbf{x}, \omega) + \int_{\mathbf{y}}^{\mathbf{x}} \int_{4\pi} K(\mathbf{x}, \mathbf{x}', \omega, \omega') L(\mathbf{x}', \omega') d\omega' d\mathbf{x}'$$

where

$$L^0(\mathbf{x}, \omega) = \tau(\mathbf{x}, \mathbf{y}) L(\mathbf{y}, \omega) + \int_{\mathbf{y}}^{\mathbf{x}} \tau(\mathbf{x}', \mathbf{x}) \varepsilon(\mathbf{x}') d\mathbf{x}'$$

$$K(\mathbf{x}, \mathbf{x}', \omega, \omega') = \tau(\mathbf{x}', \mathbf{x}) \sigma_s(\mathbf{x}') f_p(\mathbf{x}', \omega, \omega')$$



If we define a linear operator \mathbf{K} to encapsulate the integration:

$$\mathbf{K}L = \int_{\mathbf{y}}^{\mathbf{x}} \int_{4\pi} K(\mathbf{x}, \mathbf{x}', \omega, \omega') L(\mathbf{x}', \omega') d\omega' d\mathbf{x}'$$

then the equation is simply:

$$L = L^0 + \mathbf{K}L$$

which is the same form as the surface rendering equation.