Volumetric Path Tracing

Steve Marschner
Cornell University
CS 6630 Fall 2009, 24 September

Using Monte Carlo integration is a good, easy way to get correct solutions to the radiative transfer equation. It is not fast (at least not as naïvely implemented), especially for inhomogeneous or highly scattering media, but it is easy to get the right answer (plus random noise) for complex cases where there pretty much exists no other way to get an artifact-free solution.

Deriving the path tracing algorithm pretty much boils down to applying the Monte Carlo integration framework to the integral form of the volume rendering equation that we derived in the radiative transport lectures. As usual, the key to developing correct algorithms is to carefully write down the integrals involved and construct the Monte Carlo estimators systematically.

The integral form of the equation we need to solve, in the notation I've been using in lecture, is:

$$L(x,\omega) = \int_x^y \tau(x,x') \Big[\underbrace{\sigma_s(x') \int_{4\pi} f_p(x',\omega,\omega')\, L(x',\omega')\, d\omega'}_{\text{inscattering}} + \underbrace{\epsilon(x',\omega)}_{\text{emission}} \Big] dx' + \underbrace{\tau(x,y)\, L_e(y,\omega)}_{\text{attenuation}}$$

The MC approach is to identify the integrand, choose a point from a suitable pdf, and compute the ratio:

$$\frac{f}{p} = \tau(x,x') \Big[\sigma_s(x')\, f_p(x',\omega,\omega')\, L(x',\omega') + \frac{1}{4\pi}\, \epsilon(x',\omega) \Big] \Big/ p(x',\omega')$$

Since we will usually choose x' first, then choose ω' conditioned on that choice, we can write this as a nested estimator that estimates the ray integral using an estimate for the scattering integral:

$$\frac{f}{p} = \tau(x,x') \left[\frac{\sigma_s(x')\, f_p(x',\omega,\omega')\, L(x',\omega')}{p(\omega')} + \epsilon(x',\omega) \right] \Big/ p(x')$$

Let's apply this first in some special cases. First we will look at the case for homogeneous, gray (wavelength-independent) media, for which is *much* easier to get a working implementation.

1. Emission only

Here $\tau(x, x') = 1$, and $\sigma_s = \sigma_a = 0$. The integral is just

$$L(x,\omega) = \int_x^y \epsilon(x',\omega)\, dx'$$

If we use uniform sampling with respect to distance along the ray, we have the estimator:

$$g = \epsilon(x',\omega)\, \|x-y\|$$

If the emission is homogeneous, then the answer is simple: $\varepsilon\|x-y\|$. The estimator is constant, because p is proportional to f—this is "perfect" importance sampling, causing a zero-variance estimator. (Of course, it also requires having the answer in closed form in the first place!)

2. Homogeneous, absorption only.

Here $\varepsilon = \sigma_s = 0$, and $\tau(\mathbf{x}, \mathbf{y}) = \exp(-\sigma_a\|\mathbf{x}-\mathbf{y}\|)$.

Since there is no scattering or emission, the integral along the line goes away and we are left with the simple answer:

$$L(\mathbf{x}, \omega) = \tau(\mathbf{x}, \mathbf{y})\, L_e(\mathbf{y}, \omega).$$

There's no integral, so there's no Monte Carlo.

3. Emission + homogeneous absorption.

This is the first case where things become nontrivial. Here $\sigma_s = 0$ so the remaining terms are:

$$L(x,\omega) = \int_x^y \tau(x,x')\, \epsilon(x',\omega)\, dx' + \tau(x,y)\, L_e(y,\omega).$$

For the integral,

$$f = \tau(x,x')\ \epsilon(x',\omega)$$

and for uniform sampling

$$g = \|x-y\|\, e^{-\sigma_a\|x-x'\|}\, \epsilon(x',\omega).$$

This will work, but depending on the density of the medium, it can be mighty inefficient. (Exercise: sketch g as a function of arc length along the ray. What happens to the variance as the medium becomes denser?) Once there is attenuation, uniform sampling is no longer a good strategy, but a very simple importance sampling strategy is available: we can importance sample the attenuation.

Attenuation in the homogeneous medium is $\tau(s) = \exp(-\sigma_t s)$ where s is arc length along the ray. We would like to find a pdf $p(s)$ that is proportional to $\tau(s)$. To make it a pdf we just need to normalize it. The normalization depends upon the domain over which we normalize it—let me start by assuming the ray never exits the medium, so the domain of integration is $[0, \infty]$. Normalizing τ to make a pdf:

$$p(s) = \frac{\tau(s)}{\int_0^\infty \tau(t)dt} = \frac{\exp(-\sigma_t s)}{[-\exp(-\sigma_t t)/\sigma_t]_{t=0}^{\infty}} = \sigma_t \exp(-\sigma_t s)$$

Then, integrating p from 0 to s we obtain $P(s)$, the cumulative distribution function (cdf):

$$P(s) = \int_0^s p(s')ds' = \left[-\exp(-\sigma_t s')\right]_0^s = 1 - \exp(-\sigma_t s)$$

Recall the procedure for choosing a sample according to p is to take a random number ξ that is uniform over the interval [0,1] and solve $P(s) = \xi$. In this case that leads to

$$\xi = 1 - \exp(-\sigma_t s)$$
$$s = -\frac{\ln(1-\xi)}{\sigma_t}$$
or equiv: $$s = -\ln \xi / \sigma_t$$

(since $1 - \xi$ has the same distribution as ξ.)

Using this pdf leads to the estimator:

$$g = \frac{f}{p} = \frac{\tau(x,x')\,\epsilon(x',\omega)}{\sigma_t\,\tau(x,x')} = \frac{\epsilon(x',\omega)}{\sigma_t}$$

This will perform a lot better than uniform random sampling in relatively dense media, and because it makes the estimator so simple to compute, it is an elegant choice. The procedure looks like:

function radiance_estimator($\mathbf{x}$, ω):
    ξ = rand()
    s = –ln ξ / σ_t
    return $\varepsilon(\mathbf{x} - s\omega)$ / σ_t

Of course when the medium doesn't continue indefinitely along the ray, we'd like to choose a point according to $\tau(s)$ but restricted an interval $[0, s_{\max}]$. This is an easy generalization of the previous result (Exercise: The difference is in the normalization factor. Work it out.) but isn't quite so elegant. Since, if the ray ends somewhere, we also need to include the radiance from behind the medium, a simple trick is to generate s and then compare it with $s_{\max}$. The probability that $s > s_{\max}$ is $1 - \mathrm{P}(s_{\max})$, which is $\exp(-\sigma_t s_{\max})$: exactly the attenuation that needs to be applied to the radiance from behind the medium. So if we just return this background radiance whenever we attempt to choose a point past the end, the expected value is correct, leading to the following tweaked procedure:

function radiance_estimator($\mathbf{x}$, ω):
    $s_{\max}$, L_0 = trace_ray()
    ξ = rand()
    $s = -\ln \xi / \sigma_t$
    if $s < s_{\max}$ **then**
        return $\varepsilon(\mathbf{x} - s\omega) / \sigma_t$
    else
        return L_0

Another way to interpret the sampling procedure is that we select a random attenuation ξ and then find the distance s for which $\tau(s) = \xi$.

4. Emission + (inhomogeneous) absorption

In this situation the integral to be computed is still:

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

except this time, when we open the black box called τ, we find:

$$L(\mathbf{x}, \omega) = \int_{\mathbf{y}}^{\mathbf{x}} \exp\left(-\int_{\mathbf{x}'}^{\mathbf{x}} \sigma_t \right) \varepsilon(\mathbf{x}') \, d\mathbf{x}'$$

We can approach this using the same importance sampling scheme as above, but it is no longer so simple to compute the attenuation. Instead of just evaluating an exponential, we have to numerically integrate the spatially varying attenuation coefficient along the ray. Usually the attenuation integral is computed using some very simple quadrature rule—for example, the trapezoid rule, or even just the Riemann sum—on a regularly spaced set of samples along the ray. Establishing a suitable spacing requires knowing something about how fast the attenuation coefficient can vary in the volume.

So assuming we have some scheme for computing the integral, we can make an estimator from it:

$$g = \frac{\tau(x, x') \, \epsilon(x', \omega)}{p(x')}$$

but again, it is crucial to importance sample, which no longer can be done using the simple computation for homogeneous media. Setting p to be proportional to τ as before is possible, but leads to a messy expression for P. However, generalizing the "find where attenuation is ξ" procedure is simple. We need to compute the distance at which $\tau = \xi$ (for a uniform random ξ), which can be done by modifying the numerical integrator so that it marches along the ray from the starting point, accumulating attenuation, until it just crosses ξ, then interpolates to find the sample point.

A simple Riemann integration procedure to find the attenuation between 0 and s is

function attenuation($s_{\max}$, h)
    $T = 0$

for ($s = 0$; $s < s_{\max}$; s += h)
    $f = \sigma_t(s)$
    T += hf
T –= $(s - s_{\max})\, f$
return exp(–T)

To solve for $\tau(s) = \xi$, just alter the stopping condition:

function find_attenuation(τ_0, τ_1, s_0, h)
    $T = \ln(\tau_0)$
    $T_1 = \ln(\tau_1)$
    for ($s = s_0$; $T < T_1$; s += h)
        $f = \sigma_t(s)$
        T += hf
    s –= $(T - T_1) / f$
    return s

Of course these routines can be improved without much extra complexity by using the trapezoid rule.

The pdf that results from this procedure is

$$p(t) = p(\xi)\left|\frac{d\xi}{dt}\right|$$

$$\frac{d\tau}{dt} = \frac{d}{dt}\exp\left(-\int_0^t \sigma_t(x(t'))\,dt'\right)$$

$$= -\exp\left(-\int_0^t \sigma_t(x(t'))\,dt'\right)\sigma_t(x(t))$$

$$= \tau(x, x')\,\sigma_t(x')$$

(Note here that $d\tau/dt = d\xi/dt$.) If we use this pdf for importance sampling, the estimator becomes:

$$g = \frac{\tau(x,x')\,\epsilon(x',\omega)}{p(x')} = \frac{\epsilon(x',\omega)}{\sigma_t(x')}$$

This all leads to the following algorithm for sampling an inhomogeneous medium:

function radiance_estimator($\mathbf{x}$, ω):
    $s_{\max}$, L_0 = trace_ray()
    ξ = rand()
    s = find_attenuation(ξ, 0, h)
    if $s < s_{\max}$ **then**
        return $\varepsilon(\mathbf{x} - s\omega) / \sigma_t(\mathbf{x} - s\omega)$
    else
        return L_0

The disappointing thing is that we can't just plug in a Monte Carlo estimator for the attenuation integral. If it was true that $E\{\exp(X)\} = \exp(E\{X\})$ then we'd be in good shape—but sadly it is not. This means that the only thing we can do is put in a sufficiently accurate numerical estimate of the integral: we can't count on statistical averaging to let us use a very noisy estimate.

We also can't generate an individual sample of s very quickly—we have to integrate along the ray from the eye to the sample point. This means that if we are planning to take many samples along the ray (which, of course, we will need to do), it would be terribly inefficient to handle them independently—in fact, it would be quadratic in the number of samples required.

To generate samples for many values of ξ together, let's assume that we have the list of random numbers $\xi_0, \ldots, \xi_{(n-1)}$ in hand *a priori*, and that they are sorted so that $\xi_0 < \xi_1 < \cdots < \xi_{(n-1)}$. To generate the many samples all at once, we can call the integration procedure multiple times, each time picking up from where it left off; you'll note I added the arguments τ and s_0 so that this is easy to do:

```
function radiance_estimator(x, ω, ξ-list):
    s_max, L_0 = trace_ray()
    S = 0
    n = ξ-list.length
    for (i = 0; i < n; i++)
        ξ = ξ-list[i]
        s = find_attenuation(ξ, s, s_max, h)
        if s < s_max then
            S += ε(x − sω) / σ_t(x − sω)
        else
            break
    return (S + (n − i)L_0) / n
```

(I hope this pseudocode is correct, but as it wasn't actually transcribed from working code there is some chance of bugs.) Here I am giving the samples that fall past $s_{\max}$ (all $n–i$ of them) the value L_0 in the average.

Note that it is easy to perform stratified sampling on the ray integral: you just use stratified random values for the ξ_i, for instance by splitting the interval up into n pieces and choosing one random number in each one:

```
function generate_ξs(n):
    for (i = 0; i < n; i++)
        ξs[i]  = (i + rand()) / n
    return ξs
```

This is easy to do—in fact, easier than generating a sorted list of independent random numbers—and will perform much better.

Another way to look at this, which makes a lot of sense if evaluating ε is cheap, is that we should just use enough samples that the emission is lost.

5. Introducing the scattering term

We now know how to handle absorption and emission in the general case. The remaining issue is how to compute scattering. As we've observed before, outscattering behaves exactly like absorption, as far as attenuation is concerned. Similarly, inscattering behaves exactly like emission, as far as the integral along the ray is concerned. So we can use the method from the previous section but modify the emission by adding a Monte Carlo estimator for the scattering integral.

Estimating the scattering integral in a volume is just like estimating the reflection integral at a surface. It is an integral over the sphere, and we will choose directions to sample, thinking of them as points on the sphere.

$$\epsilon_s(x', \omega) = \sigma_s(x') \int_{S^2} f_p(x', \omega, \omega') \, L(x', \omega') \, d\omega'$$

The first thing to try is importance sampling by the phase function. (This means that whenever we propose a particular phase function, we also need to come up with a procedure for sampling it. There is a very simple procedure available for Henyey-Greenstein, the most popular phase function model out there.) Remember the phase function is normalized to be a probability distribution already, so in this case $p(\omega') = f_p(\mathbf{x}', \omega, \omega')$. Then the estimator is

$$g(\omega') = \frac{f}{p} = \frac{\sigma_s(x') \, f_p(x', \omega, \omega') \, L(x', \omega')}{p(\omega')}$$

$$g(\omega') = \sigma_s(x') \, L(x', \omega')$$

So all we need to do is trace a recursive ray to estimate $L(\mathbf{x}', \omega')$ and the process becomes:

1. Choose ξ in [0, 1] and find the point $\mathbf{x}'$ on the ray for which $\tau(\mathbf{x}, \mathbf{x}') = \xi$.
2. Choose ω' according to the phase function.
3. Recursively estimate $L(\mathbf{x}', \omega')$ and compute the estimator $(\varepsilon(\mathbf{x}', \omega) + \sigma_s(x') \, L(\mathbf{x}', \omega')) \, / \, \sigma_t(\mathbf{x}')$.

That is, we simply replace the statement starting S += … by the following:

```
ω′ = sample_phase_function()
Ls = radiance_estimator(x – sω, ω′, new_ξ_list())
S += (ε(x – sω) + σs Ls) / σt(x – sω)
```

Doing this leads to a path tracing algorithm very similar to a brute-force surface path tracer. Adding luminaire sampling, multiple importance sampling, and Russian Roulette all requires straightforward generalizations of the methods used for surfaces.