

BSSRDF Notes

Andy Scukanec, Cornell University

July 28, 2004

1 Overview

The 2001 paper by Henrik Jensen, Stephen Marschner, Marc Levoy, and Pat Hanrahan titled “A Practical Model for Subsurface Light Transport” attempts to represent light transport that occurs beneath the surface-air interface in materials which exhibit subsurface activity. There are some assumptions made in the paper which restrict this method to a certain subset of these problems, albeit a widely and naturally occurring subset. The authors of the paper have developed an approximation scheme to estimate the light transport between two points on the surface of an object. As with all approximations, they fail at certain points and only full monte carlo simulation can achieve the true steady state. Regardless, the approximation developed within works in many practical situations and breaks down at a point where the human visual system wouldn’t notice.

1.1 Assumptions

There are two assumptions made in this paper. The first is that the media on the “inside” of the interface scatters far more light than it observes (ie $\sigma_a \gg \sigma_s$ or α , the albedo, is very close to 1). The second restriction asserts that the approximation is only valid for locally flat surfaces. This implies that it will fail near edges and on geometry with high curvature.

2 Basic Setup

Conventions and basic ideas:

- The i subscript will indicate attributes related to incoming light.
- The o subscript will indicate attributes related to outgoing light.
- Given a point and a direction, we want to find the amount of light leaving that point in that direction. Light will be coming in from all over the surface, and some fraction of the light at each

point will make it to our special output point and direction (x_o and ω_o). Thus we will be integrating over the surface.

- For each point of incoming light on the surface, we will want to figure out how much energy is getting to that point from any direction. Thus we will be integrating incoming light over the hemisphere.
- Flux and power are the same thing. Both are represented as Φ and are in units of Watts.
- Any variable labeled x (regardless of subscript etc) is a 3-d point.
- Any variable labeled ω (regardless of subscript etc) is a 2-d direction. These are often represented by a θ, ϕ pair, although we don’t really care how they are represented in this paper.
- Integral subscripts: \mathcal{H}^2 means integrate over the positive hemisphere, \mathcal{M} means integrate over all of the surface area of the manifold, \mathcal{S}^2 means over the whole sphere.
- The measure $d\sigma(\omega)$ refers to the solid angle measure, and the measure $d\mu(\omega)$ refers to the projected solid angle measure (ie - with the cosine term rolled in).
- $\sigma_s + \sigma_a = \sigma_t$. All three attributes are in units of mm-1 or some other inverse distance. NOTE: the absolute scale of these numbers matters very much.
- Albedo, or α , is the fraction of non-absorbed light. $\alpha = \frac{\sigma_s}{\sigma_t}$.
- A mean free path is the average distance a photon of light would travel into the media, and is equal to $\frac{1}{\sigma_t}$.
- The relative index of refraction $\eta = \frac{\eta_-}{\eta_+}$ is the fraction of the internal absolute index of refraction η_- to the external absolute index of refraction η_+ . Hypothetically, either or both absolute

indices could vary with position, so η could as well. For simplicity though, I will drop the function semantics from η .

- The subscripts v and r have nothing to do with the function parameter r . The subscripts v and r indicate a relationship with the virtual and real sources of the dipole. The meaning is not important if you only mean to implement the paper. The meaning is very important if you need to understand the mathematics. The function parameter r simply represents a scalar distance.
- Any scalar quantities with a prime after them are “reduced” quantities meant to adjust for anisotropy in the phase function. Any vector quantities with a prime are inside the surface and have been properly refracted.
- ξ always refers to a scalar chosen uniformly at random.
- In my illustrations, w stands for ω .

The BSSRDF relates incident flux ($\frac{J}{s} = W$) to exitant radiance:

$$(1) \quad dL_o(x_o, \omega_o) = S(x_i, \omega_i; x_o, \omega_o) d\Phi_i(x_i, \omega_i)$$

The integral form of this equation is:

$$(2) \quad L_o(x_o, \omega_o) = \int_{\mathcal{M}} \int_{\mathcal{H}^2} S(x_i, \omega_i; x_o, \omega_o) L_i(x_i, \omega_i) d\mu(\omega_i) dA(x_i)$$

S is really a sum of two terms - the single scattering term and the diffusion approximation:

$$(3) \quad S(x_i, \omega_i; x_o, \omega_o) = S_d(x_i, \omega_i; x_o, \omega_o) + S^{(1)}(x_i, \omega_i; x_o, \omega_o)$$

Both S_d and $S^{(1)}$ can be probabilistically sampled, although they will likely require different sampling methods. I will outline my sampling methods later, but these are hardly the only ones that would work.

2.1 Evaluating S_d

$$(4) \quad S_d(x_i, \omega_i; x_o, \omega_o) = \frac{F}{\pi} R_d(\|x_i - x_o\|)$$

$$(5) \quad R_d(r) = \frac{\alpha}{4\pi} \left[z_r(\sigma_{tr}d_r(r) + 1) \frac{\exp(-\sigma_{tr}d_r(r))}{d_r^3(r)} + z_v(\sigma_{tr}d_v(r) + 1) \frac{\exp(-\sigma_{tr}d_v(r))}{d_v^3(r)} \right]$$

$$\begin{aligned} d_r(r) &= \sqrt{z_r^2 + r^2} & z_r &= z_v + 4AD \\ d_v(r) &= \sqrt{z_v^2 + r^2} & z_v &= \frac{1}{\sigma'_t} \\ A &= \frac{1 + F_{dr}}{1 - F_{dr}} & D &= \frac{1}{3\sigma'_t} \\ \sigma_{tr} &= \sqrt{3\sigma_a\sigma'_t} & \alpha' &= \frac{\sigma'_s}{\sigma'_t} \\ \sigma'_t &= \sigma'_s + \sigma_a & \sigma'_s &= (1 - g)\sigma_s \end{aligned}$$

$$\begin{aligned} F_{dr} &= \int_{\mathcal{H}^2} F_r(\eta, n \cdot \omega) (n \cdot \omega) d\sigma(\omega) \\ &\approx -1.44\eta^{-2} + 0.71\eta^{-1} + 0.668 + 0.0636\eta \\ g &= \int_{S^2} (\omega \cdot \omega) p(\omega \cdot \omega) d\sigma(\omega) \end{aligned}$$

The phase function p and the mean cosine g are usually considered material properties and should be specified on a per material basis, as well as σ_s , σ_t , and η .

...and that should be that. You should now have enough information to evaluate S_d . This is only half of it though, since you will have to sample S_d and evaluate the single scattering term too.


2.2 Evaluating $S^{(1)}$

Rather than writing an actual expression for $S^{(1)}$, it's easier to write an expression for the outgoing radiance due to $S^{(1)}$.

$$(6) \quad L_o^{(1)}(x_o, \omega_o) = \frac{\sigma_s(x_o)}{\sigma_{tc}} F p(\omega_i \cdot \omega_o) E L_i(x_i, \omega_i)$$

$$\begin{aligned} E &= e^{-s'_i\sigma_t(x_i)} e^{-s'_o\sigma_t(x_o)} & F &= F_t(\eta^{-1}, \omega_o) F_t(\eta, \omega_i) \\ \sigma_{tc} &= \sigma_t(x_o) + G\sigma_t(x_i) & G &= \frac{|n_i \cdot \omega'_o|}{|n_i \cdot \omega'_i|} \\ s'_o &= \frac{\log(\xi)}{\sigma_t(x_o)} & s'_i &= s_i \frac{\|\omega_i \cdot n_i\|}{\cos(\theta'_i)} \end{aligned}$$

Here, ξ is chosen uniformly at random in the range $[0, 1]$. s_i is the observed distance and s'_i is the refracted distance. In practice, I found it much easier to shoot rays from x_{inside} towards the light (ignoring the refraction on the way out, sadly), and using the distance from x_{inside} to the intersection of the light ray with the surface as s'_i .



3 Sampling Methods Used

In keeping with the monte carlo way of life, I will define all my sampling methods by using estimators and ensuring that the expected value of the estimator is in fact the integral I want to compute.

3.1 The Diffusion Term

In sampling the diffusion term, I have not been able to find a way to sample such that bias is completely avoided. I have not concerned myself overly with it though, as this method is inherently biased anyways. The basic difficulty in sampling the diffusion term is that R_d is defined on a scalar distance. Given an arbitrary point x_o on a surface of arbitrary geometry, it is difficult to arbitrarily pick another point on the surface.

In order to combat the lack of guarantees about the surface, given x_o , I picked (according to your favorite method) a random light in the scene. Next, I created a disc centered about x_o , with its normal oriented toward the light (or chosen point on that light if it is an area light). The radius of the disc is up to you, but I use $\frac{n}{\sigma_t}$, where n ranges from 1 to 4 or 5. Choosing a point on the disc (which is easy to do), I shoot a ray from the point on the light to the chosen point (called x_{disc}). If that ray intersects the surface \mathcal{M} , then I record the intersection location as x_i . The ray's direction is ω_i , and regardless of your particular raytracing infrastructure, it should be simple to query the light for its exitant radiance on the sampled point in the direction of x_{disc} . Given this setup, it is now possible to evaluate the amount of radiance


leaving x_o in ω_o . This will serve as an estimate of $L_o(x_o, \omega_o)P(\text{picking } x_i)$. To actually estimate L_o , divide the resulting quantity by the probability of picking x_i (I will show how to compute this shortly). Repeatedly taking this measurement and averaging the results will yield a good estimate of L_o .

$$P[x_i] = \frac{P[x_{disc}] P[x_{light}]}{|n_{disc} \cdot \omega_{light}|} \quad P[x_{disc}] = \frac{1}{\pi R_{max}^2}$$

$$r = R_{max} \sqrt{\xi_1} \quad \theta = 2\pi \xi_2$$

$$R_{max} = \frac{n}{\sigma_t} \quad \xi_1, \xi_2 \in (0..1]$$

n is simply the number of mean free paths that you wish the disc to encompass, n_{disc} is the normal of the disc. ω_{light} is the direction from the light to x_{disc} . $P[x_{light}]$ is the probability of picking that specific light and point on that light (determined by your own methods). r and θ are the polar coordinates of x_{disc} , generated in a uniform manner. Note that it is probably better to use a non-uniform distribution in order to importance sample R_d , and that the orientation of the x and y axes for the disc do not matter (just so long as the x , y , and normal axis form an orthonormal basis). In otherwords, rotating the disc about its normal changes nothing.



3.2 The Single Scattering Term

4 Extra Material

There is a confirmed error in equation 4 of the referenced paper. The rewritten equation in this document is correct:

$$R_d(r) = \frac{\alpha}{4\pi} \left[z_r(\sigma_{tr}d_r(r) + 1) \frac{\exp(-\sigma_{tr}d_r(r))}{d_r^3(r)} + z_v(\sigma_{tr}d_v(r) + 1) \frac{\exp(-\sigma_{tr}d_v(r))}{d_v^3(r)} \right]$$