

Unbiased Photon Gathering for Light Transport Simulation

Supplementary Material

1 Computing the Angular Bound

In this section, we will explain the mathematical details behind the angular bound of tentative ray tracing, which we have briefly described in Section 3.4 of the paper. We need to choose an angular bound so that Eq. (16) in the paper is analytically integrable. In practical implementation, the direction of a tentative ray is generated by uniform sampling a random number in the 2D space of $[0, 1] \times [0, 1]$. Therefore, we choose an angular bound within which the random numbers generating tentative rays cover an axis-aligned bounding box (AABB) in the 2D space. Then the probability $p^b(\bar{\mathbf{x}}_{s', t'-1})$ is the AABB area.

Here we take two common BSDFs as examples, including the Lambertian model [Lambert 1760] and the Phong model [Phong 1975], and develop an AABB bound in a uniform random number space. The bound is both used to sample the tentative rays and compute the probability density integration. We also show how to handle BSDFs that are a linear combination of multiple components.

1.1 Hemisphere Angular Bound

Given a reflective material, the tentative ray $z_{t'-1} \rightarrow z$ is confined within the upward hemisphere defined by the surface normal at $z_{t'-1}$. The neighborhood we wish to sample is a sphere located at $y_{s'}$ with radius d .

For brevity, we use the local frame $\{X, Y, Z\}$ at $z_{t'-1}$, where Z is the surface normal. The direction $z_{t'-1} \rightarrow z$ can be represented in spherical coordinates as $\{\theta, \phi\}$, where θ and ϕ are the polar and azimuthal angles respectively. When sampling rays, θ and ϕ are typically computed from two independent uniform random numbers $\{r_\theta, r_\phi\}$

$$\theta = f_\theta(r_\theta), \quad \phi = f_\phi(r_\phi), \quad (1)$$

where f_θ and f_ϕ are mapping functions for importance sampling. Since f_θ and f_ϕ are typically monotonic, an AABB bound $\{\theta, \phi\} \in \Theta \times \Phi$ in the angular space $\theta \times \phi$ can be directly converted to an AABB bound $\{r_\theta, r_\phi\} \in R_\theta \times R_\phi$ in the random number space $r_\theta \times r_\phi$

$$\begin{aligned} R_\theta &= [f_\theta^{-1}(\theta|_\Theta)_{\inf}, f_\theta^{-1}(\theta|_\Theta)_{\sup}], \\ R_\phi &= [f_\phi^{-1}(\phi|_\Phi)_{\inf}, f_\phi^{-1}(\phi|_\Phi)_{\sup}], \end{aligned} \quad (2)$$

where the inf and sup subscripts refer to the lower and upper bound respectively.

By definition, the mappings f_θ and f_ϕ are isometric. Therefore, the probability density integration p^b is simply the AABB area in the random number space

$$p^b(\bar{\mathbf{x}}_{s', t'-1}) = \|R_\theta \times R_\phi\|. \quad (3)$$

A simple conservative angular bound of the spherical neighborhood

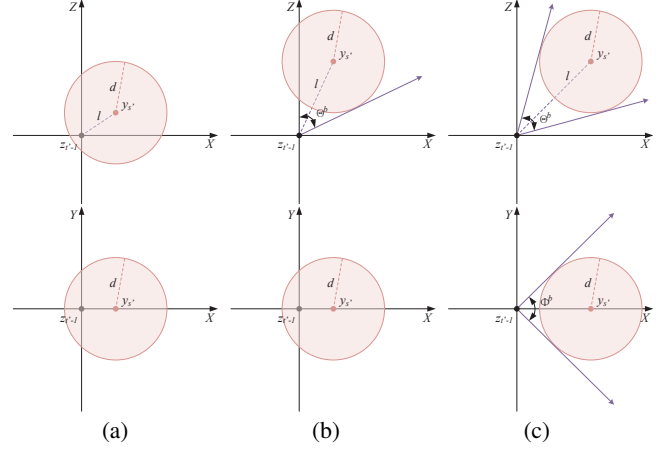


Figure 1: Three different cases of angular bounds. (a) direction is not bounded, as $z_{t'-1}$ is inside the neighborhood. (b) only θ is bounded, as Z -axis intersects the neighborhood. (c) both θ and ϕ are bounded.

can be written for the three cases in Fig. 1

$$\begin{cases} \Theta = [0, \frac{\pi}{2}] \\ \Phi = [0, 2\pi] \end{cases}, \quad \text{if } l \leq d, \\ \begin{cases} \Theta = [0, \min(\frac{\pi}{2}, \theta_c + \arcsin \frac{d}{l})] \\ \Phi = [0, 2\pi] \end{cases}, \quad \text{else if } \frac{d}{l} \geq \sin(\theta_c), \\ \begin{cases} \Theta = [\theta_c - \arcsin \frac{d}{l}, \min(\frac{\pi}{2}, \theta_c + \arcsin \frac{d}{l})] \\ \Phi = [\phi_c - \arcsin \frac{d}{l \sin \theta_c}, \phi_c + \arcsin \frac{d}{l \sin \theta_c}] \end{cases}, \quad \text{otherwise,} \end{cases} \quad (4)$$

where $\{\theta_c, \phi_c\}$ are the spherical coordinates of the direction vector $z_{t'-1} \rightarrow y_{s'}$, and l is the distance to the neighborhood center $l = \|z_{t'-1} - y_{s'}\|$. In the first case shown in Fig. 1(a), $z_{t'-1}$ is inside the neighborhood sphere, and a bound cannot be placed on $\{\theta, \phi\}$. In the second case shown in Fig. 1(b), the neighborhood sphere intersects the Z -axis, and only θ is bounded. In the third case shown in Fig. 1(c), the neighborhood sphere is well separated from the Z -axis, and we can bound both θ and ϕ . The AABB bound in random number space can be computed using Eq. (2), which depends on the BSDF importance sampling functions f_θ and f_ϕ .

1.2 Bounds for Lambertian and Phong BSDFs

1.2.1 Lambertian BSDF

The Lambertian diffuse BSDF is a simple constant function

$$f_s^{\text{Lambertian}}(\omega_i, \omega_o) = \kappa, \quad (5)$$

where we parameterize the BSDF f_s over directions. ω_i corresponds to the direction $z_{t'-1} \rightarrow z_{t'-2}$, and ω_o corresponds to the direction $z_{t'-1} \rightarrow z$ or $\{\theta_o, \phi_o\}$.

As the constant f_s does not affect sampling, the outgoing ray direction $z_{t'-1} \rightarrow z$ for a Lambertian BSDF is typically importance-sampled from the cosine part of the geometric term

$$p_x(z_{t'-2} \rightarrow z_{t'-1} \rightarrow z) \propto \cos \theta_o. \quad (6)$$

The importance function f_θ can be found by inverting the cumulative distribution function

$$\begin{aligned} f_\theta^{-1}(\theta_o) &= 2 \int_{\theta_o}^{\frac{\pi}{2}} \cos \theta \sin \theta d\theta = \cos^2 \theta_o, \\ f_\theta(r_\theta) &= \arccos \sqrt{r_\theta}, \end{aligned} \quad (7)$$

and f_ϕ is simply:

$$f_\phi(r_\phi) = 2\pi r_\phi. \quad (8)$$

The bounds and p^b can be computed using Eq. (2) and Eq. (3):

$$\begin{aligned} r_\theta &\in [\cos^2 \Theta_{\text{sup}}, \cos^2 \Theta_{\text{inf}}], \\ r_\phi &\in \left[\frac{\Phi_{\text{inf}}}{2\pi}, \frac{\Phi_{\text{sup}}}{2\pi} \right], \\ p^b &= (\cos^2 \Theta_{\text{inf}} - \cos^2 \Theta_{\text{sup}}) \left(\frac{\Phi_{\text{sup}}}{2\pi} - \frac{\Phi_{\text{inf}}}{2\pi} \right). \end{aligned}$$

1.2.2 Phong BSDF

The Phong BSDF [1975] is a perceptually based model for glossy reflectance, which is symmetric around the mirror reflectance direction ω_r

$$f_s^{\text{Phong}}(\omega_i, \omega_o) = \kappa \cos^k \theta_{o-r}, \quad (9)$$

where θ_{o-r} is the angle between ω_o and ω_r . As k is typically a large value which makes f_s dominate p_x , we directly use a normalized version of f_s^{Phong} as the importance function and leave the cosine term mentioned in the previous subsection out of the sampling.

For convenience, we first sample the angles θ_{o-r} and ϕ_{o-r} , then compute ω_o from it. Similar to the Lambertian case, the importance function f_θ can be found by inverting the cumulative distribution function:

$$\begin{aligned} f_\theta^{-1}(\theta_{o-r}) &= 2 \int_{\theta_{o-r}}^{\frac{\pi}{2}} \cos^k \theta \sin \theta d\theta = \cos^{k+1} \theta_{o-r}, \\ f_\theta(r_\theta) &= \arccos r_\theta^{\frac{1}{k+1}}, \end{aligned} \quad (10)$$

and f_ϕ is the same as Eq. (8). The bounds and p^b are:

$$\begin{aligned} r_\theta &\in [\cos^{k+1} \Theta_{\text{sup}}, \cos^{k+1} \Theta_{\text{inf}}], \\ r_\phi &\in \left[\frac{\Phi_{\text{inf}}}{2\pi}, \frac{\Phi_{\text{sup}}}{2\pi} \right], \\ p^b &= (\cos^{k+1} \Theta_{\text{inf}} - \cos^{k+1} \Theta_{\text{sup}}) \left(\frac{\Phi_{\text{sup}}}{2\pi} - \frac{\Phi_{\text{inf}}}{2\pi} \right). \end{aligned}$$

Note that the direction ω_o is sampled in the upper hemisphere defined by ω_r , not the surface normal. Consequently, the direction may point into the surface, resulting in the measurement contribution being constantly 0. While such rays are wasted, in practice this does not occur very frequently because for highly glossy surfaces ω_o typically stays near ω_r . As long as this case is tested and the measurement contribution is properly zeroed, the final Monte Carlo estimation remains unbiased.

1.3 Bounds for Multiple Component BSDF

Importance Sampling in General. A complex BSDF is frequently defined as a linear combination of multiple components, such as a diffuse one and a specular one. The formal definition is

$$f_s = \sum_{i=1}^h f_{s,i}, \quad (11)$$

where h is the number of components. The corresponding probability density of importance sampling without angular bound can be also formulated with respect to the multiple components

$$p = \frac{f_s}{\int_{\Omega} f_s d\omega_o} = \frac{\sum_{i=1}^h \kappa_i p_i}{\sum_{i=1}^h \kappa_i}, \quad (12)$$

$$\kappa_i = \int_{\Omega} f_{s,i} d\omega_o, \quad p_i = \frac{f_{s,i}}{\kappa_i}, \quad (13)$$

where Ω is the domain of outgoing directions. κ_i and p_i represent the reflectivity and sampling probability density for an individual component $f_{s,i}$ respectively.

Importance sampling of a multiple-component BSDF is typically achieved by first selecting an individual component i to sample with a probability proportional to the respective reflectivity values κ_i . Then the sample is generated for component i using its importance function p_i . Formally, the importance sampling takes two steps.

1. Randomly select one component based on reflectivity. The i -th component is chosen with probability

$$p_{\text{select},i} = \frac{\kappa_i}{\sum_{j=1}^h \kappa_j}. \quad (14)$$

2. Sample a direction using the chosen component.

What we are interested in is the acceptance probability p^c for the hypothetical Russian roulette event (Section 3.2 of the paper). Note that the Russian roulette event itself only tests whether a ray hits a fixed spatial neighborhood, which does not depend on the BSDF component used to sample the ray. Consequently, we can simply combine the per-component probability values p_i^c using Eq. (12)

$$p^c = \frac{\sum_{i=1}^h \kappa_i p_i^c}{\sum_{i=1}^h \kappa_i}. \quad (15)$$

Angular Bound. For efficiency and implementation convenience, we choose to apply an independent angular bound for each component, as shown in Fig. 2. Mathematically speaking, this corresponds to applying the same importance sampling process to an updated set of BSDF components $\bar{f}_{s,i}$. Assuming the angular bound for each component i has already been computed as $\{\Theta_i, \Phi_i\}$, one can define $\bar{f}_{s,i}$ explicitly as

$$\bar{f}_s = \sum_{i=1}^h \bar{f}_{s,i}, \quad (16)$$

$$\bar{f}_{s,i} = \begin{cases} f_{s,i}, & \text{if } \omega_o \in \{\Theta_i, \Phi_i\}, \\ 0, & \text{else.} \end{cases} \quad (17)$$




Figure 2: Angular bound with two component BSDF. The bounds of the two components are different. But the neighborhood S is included in their intersection.

Substituting Eq. (17) into the general multi-component sampler in Eq. (12), we formulate the importance sampling probability for \bar{f}_s

$$p = \frac{\bar{f}_s}{\int_{\Omega} \bar{f}_s d\omega_o} = \frac{\sum_{i=1}^h \kappa_i p_i^b \bar{p}_i}{\sum_{i=1}^h \kappa_i p_i^b}, \quad (18)$$

$$p_i^b = \int_{\Theta_i \times \Phi_i} p_i d\omega_o, \quad \bar{p}_i = \frac{\bar{f}_{s,i}}{\kappa_i p_i^b}, \quad (19)$$

where p_i^b is the probability density integration of each component i inside its own bound, and $\kappa_i p_i^b$ can be thought of as an effective reflectivity for the angularly bounded BSDF component \bar{f}_s . \bar{p}_i is the normalized probability density of an individual component $\bar{f}_{s,i}$. The importance sampling consists of the same two steps as we introduced for multiple components without angular bounds. However, Eq. (14) has to be updated accordingly. Replacing the original reflectivity κ_i with the effective reflectivity $\kappa_i p_i^b$, we get

$$p_{select,i}^b = \frac{\kappa_i p_i^b}{\sum_{j=1}^h \kappa_j p_j^b}. \quad (20)$$

Applying the same procedure to the Russian roulette probability p^c defined in Eq. (15), we get

$$\bar{p}^c = \frac{\sum_{i=1}^h \kappa_i p_i^b \bar{p}_i^c}{\sum_{i=1}^h \kappa_i p_i^b}, \quad (21)$$

where \bar{p}_i^c is the per-component Russian roulette acceptance probability after applying the corresponding angular bound. By definition, it differs from the original p_i^c by the normalization factor p_i^b

$$\bar{p}_i^c = \frac{p_i^c}{p_i^b}. \quad (22)$$

Substituting Eq. (22) into Eq. (21), we get

$$\bar{p}^c = \frac{\sum_{i=1}^h \kappa_i p_i^c}{\sum_{i=1}^h \kappa_i p_i^b}. \quad (23)$$

Finally, we can compare p^c in Eq. (15) with \bar{p}^c in Eq. (23) to define the multi-component normalization factor p^b

$$p^b = \frac{p^c}{\bar{p}^c} = \frac{\sum_{i=1}^h \kappa_i p_i^b}{\sum_{i=1}^h \kappa_i}. \quad (24)$$

Note that p^b only depends on the analytically computed p_i^b and the input constants κ_i . Consequently, p^b itself can also be evaluated analytically, which is a requirement for UPG to remain unbiased.

To conclude, when applying UPG to a multi-component BSDF, we first randomly choose a bounded component $\bar{f}_{s,i}$ with probability $p_{select,i}^b$ defined in Eq. (20). We then sample a tentative ray from the per-component probability density \bar{p}_i , which is achieved by confining the ray direction within the corresponding angular bound $\{\Theta_i, \Phi_i\}$. The resulting Bernoulli trial count N_b is finally scaled by a factor $1/p^b$ to compute an unbiased estimation of the probability reciprocal r , where p^b is computed using Eq. (24).

Note that the sampling cost is $O(h)$. While h is typically a small constant in the original scene description, it can become very large when UPG is combined with particle-guided BDPT, where the importance sampling function is represented as a Gaussian mixture model potentially reaching hundreds of components. This currently prevents us from combining the two algorithms efficiently.

2 Reciprocal of an Integral

In this section, we briefly introduce the unbiased algorithm for estimating the reciprocal of a general integral proposed by Booth [2007], and discuss its relationship with our method proposed in Section 3.4 of the paper.

Given a non-negative function $f(x)$ defined in a domain Ω , Booth [2007] aimed to evaluate the reciprocal of its integral

$$I = \frac{1}{\int_{\Omega} f(x) dx}. \quad (25)$$

Assuming $\int_{\Omega} f(x) dx \in (0, 1)$, Eq. (25) can be evaluated as a series expansion:

$$I = \frac{1}{1-g} = \sum_{i=0}^{+\infty} g^i, \quad (26)$$

$$g = 1 - \int_{\Omega} f(x) dx.$$

Reorganizing the infinite series and taking an independent Monte Carlo estimation for each occurrence of g , we have

$$\hat{I} = 1 + \sum_{i=1}^{+\infty} \prod_{j=1}^i \hat{g}_j, \quad (27)$$

$$\hat{g}_j = 1 - f(x_j),$$

where $\{x_j\}$ is an infinite set of stochastic samples in Ω , and we assume $\int_{\Omega} 1 dx = 1$.

The series can be terminated using Russian roulette once $\prod_{j=1}^i \hat{g}_j$ falls below a pre-specified threshold. However, in our case, f is the binary outcome of a Bernoulli trial, and \hat{g}_j always takes a value of 0 or 1. Therefore, the series in Eq. (27) degenerates into a counting process. Assuming the series of Bernoulli trials succeeds for the first time after N tries, we have:

$$\hat{g}_j = \begin{cases} 1, & j < N \\ 0, & j \geq N \end{cases}, \quad (28)$$

$$\hat{I} = 1 + \sum_{i=1}^N \prod_{j=1}^i \hat{g}_j = N,$$

which is identical to our algorithm in Section 3.4 of the paper.

References

- BOOTH, T. E. 2007. Unbiased Monte Carlo estimation of the reciprocal of an integral. *Nuclear Science and Engineering* 156, 3, 403–407.
- LAMBERT, J. H. 1760. *Photometria Sive de Mensure de Gratibus Luminis, Colorum Umbrae*. Eberhard Klett.
- PHONG, B. T. 1975. Illumination for computer generated pictures. *Commun. ACM* 18, 6 (June), 311–317.