

Importance Sampling of the Phong Reflectance Model

Jason Lawrence

We first describe the Phong reflectance model and its associated *brdf*. We then offer practical advice regarding the implementation of importance sampling in the context of this reflectance model for a Monte Carlo path tracer.

Phong Reflectance Model

We are interested in global illumination and, consequently, desire a model for describing the reflectance of a surface compatible with the definition of the *brdf* that appears in the rendering equation. Although the very famous Phong shading model offers a relatively simple (and somewhat accurate) description of how light illuminates a surface, it lacks a physical basis and, as a result, does not directly provide the *brdf* we require for accurate global illumination schemes. However, a physically plausible *brdf* based on the Phong shading model does exist and is often used in rendering algorithms. This *brdf*, as one would expect, separates the reflectance distribution of a surface into a diffuse and specular component:

$$f_r(x, \theta_i, \theta_o) = f_{r,d}(x, \theta_i, \theta_o) + f_{r,s}(x, \theta_i, \theta_o) = k_d \frac{1}{\pi} + k_s \frac{n+2}{2\pi} \cos^n \alpha$$

where:

- α = the angle between the perfect specular reflective direction and the outgoing ray direction. Values larger than $\frac{\pi}{2}$ are clamped to $\frac{\pi}{2}$ in order to prevent any negative values of the cosine factor.
- k_d = the diffuse reflectivity (the fraction of the incoming energy that is reflected diffusely)
- k_s = the specular reflectivity (the fraction of the *perpendicularly* incoming energy that is reflected specularly)
- n = the specular exponent (higher values give a sharper specular reflection)

In order for this Phong *brdf* to have the property of energy conservation we give the restriction that $k_d + k_s \leq 1$.

Practical Considerations for a Monte Carlo Path Tracer


In terms of our Monte Carlo path tracer, we are then interested in incorporating this *brdf* into our **estimation** of the integral in the rendering equation. Monte Carlo path tracing amounts to computing, for each intersection, a random incoming path of light that is used to estimate this integral. Before we proceed with an algorithmic description of how to “importantly” sample the set of all possible incoming paths of light, we first introduce some useful notation:

At the start of each iteration of the recursive step in our path tracer (i.e. after we compute the intersection of a ray with the scene) we have the following known quantities:

- x = the point of intersection with the current ray and the scene
- \vec{n} = the vector normal to the differential surface patch at the point of intersection
- \vec{w}_o = a unit vector in the direction of the current outgoing ray of light (recall that we are tracing paths backwards from the direction of light propagation: from the eye to the light sources)
- $f_r(x, \theta_i, \theta_o)$ = the Phong reflectance *brdf* described above with diffuse/specular parameters given by the material of the intersected surface

We need to compute:

- \vec{w}_i = an incoming path of light whose radiance we recursively compute. This radiance is then scaled by both the geometry term and the *brdf* (i.e. the integrand of the rendering equation) to give us an estimate of the total amount of light radiated from x towards the direction \vec{w}_o .



Importance sampling is a method by which we intelligently select \vec{w}_i so as to maximize the convergence of our estimate to the actual value (i.e. fewer instances of the random variable \vec{w}_i provide a quicker convergence to the actual value of the integral in the rendering equation than would a uniformly random selection of incoming rays of light).

We accomplish this by noticing that, although we lack any knowledge about $L_i(x, \vec{w}_i)$ (otherwise we would not need to trace any additional rays), we do have information about the local reflectance properties of the surface patch at x . More importantly, if **our random selection of \vec{w}_i favors paths of light that would potentially contribute the most possible energy to the total energy reflected towards the direction \vec{w}_o** it makes intuitive sense that we will arrive at our answer quicker. This is exactly how importance sampling works: we sample the space of all incoming rays of light in a way that favors the reflection distribution function, resulting in a provably unbiased estimate of the integral in the rendering equation that converges toward the actual answer quicker. It helps to think of the case of perfect specular reflection to convince yourself this is true.

We accomplish this non-uniform sampling of light rays by warping uniform random variables (easy to generate) to paths of light. I'll skip the mathematical details of generating this warping beyond noting that the inverse of the cumulative distribution function of a random variable (F^{-1}) that follows some distribution P gives a transformation from a uniform random variable u to a random variable $F^{-1}(u)$ which has a distribution of P as well. Now we are ready to describe the algorithm:

We first randomly select whether we'll compute a diffuse sample or a specular sample. We accomplish this by generating a uniform random variable $u \in [0,1]$ and evoking the rule:

- $u < k_d$: take a diffuse sample and compute its contribution
- $k_d \leq u < k_d + k_s$: take a specular sample and compute its contribution
- $k_d + k_s \leq u$: the contribution is 0

For a diffuse sample we want to generate a random path that favors the maximum contribution of energy as prescribed by the diffuse component of the Phong *brdf* scaled by the cosine of the incoming angle. Intuitively, this translates to generating rays of incoming light which would potentially have the largest impact on the integral in the rendering equation. This gives the following probability distribution function:


$$pdf(\theta_i) = \frac{1}{\pi} \cos(\theta_i)$$

Therefore, to generate a mapping from two uniform random variables $u_1, u_2 \in [0,1]$ that follow this distribution, we compute the separable marginal cumulative distribution functions of the pdf. We then take their inverses and obtain the following transformation (in spherical coordinates):

$$\vec{w}_i = (\theta, \phi) = (\arccos(\sqrt{u_1}), 2\pi u_2)$$

For a specular sample we wish to generate random paths that favor the maximum contribution of energy as prescribed by the specular component of the Phong *brdf*. Clearly, this demands a sampling that favors the region of maximum specular reflection (i.e. the “specular lobe” that exists along the direction of perfect specular reflective direction w.r.t. *the outgoing ray of light*):

$$\theta_i = \pi/6$$



This gives the following desired distribution:

$$pdf(\theta_i) = \frac{n+1}{2\pi} \cos^n(\alpha)$$

Where α is the angle between \vec{w}_i (the incoming path of light) and the direction of perfect specular reflective direction w.r.t. \vec{w}_o (the outgoing path of light). Notice that we have evoked the rule of reciprocity in the Phong *brdf*!

Again, in order to generate a mapping from two uniform random variables $u_1, u_2 \in [0,1]$ to a ray that follows this distribution, we compute the inverse of the cumulative distribution function, obtaining the following transformation (again in spherical coordinates):

$$\vec{w}_i = (\alpha, \phi) = (\arccos(u_1^{\frac{1}{n+1}}), 2\pi u_2)$$

Notice, \vec{w}_i is given in terms of the direction of perfect specular direction. The conversion from this local coordinate system to world coordinates is omitted.

References

- [1] E. Lafortune and Y. Willems. Using the modified Phong reflectance model for physically based rendering. Technical Report CW197, Dept. Comp. Sci., K.U. Leuven, 1994.