Dave Edwards

Notes on the Ward and Ashikhmin Anisotropic BRDF Models

1 Brief Overview of Path Tracing

Our goal is to evaluate the Rendering Equation for surfaces:

$$L_o(\hat{\omega}_o) = L_e(\hat{\omega}_o) + \int_\Omega \rho(\hat{\omega}_i, \hat{\omega}_o)(\hat{N} \cdot \hat{\omega}_i) L_i(\hat{\omega}_i) d\Omega(\hat{\omega}_i).$$

We can use Monte Carlo integration to do this, which gives us the following approximation:

$$L_o(\hat{\omega}_o) \approx L_e(\hat{\omega}_o) + \frac{1}{N} \sum_{k=1}^{N} \frac{\rho(\hat{\omega}_{i_k}, \hat{\omega}_o)(\hat{N} \cdot \hat{\omega}_{i_k}) L_i(\hat{\omega}_{i_k})}{p(\hat{\omega}_{i_k})}.$$

In path tracing, $N = 1$, so the approxmation reduces to:

$$L_o(\hat{\omega}_o) \approx L_e(\hat{\omega}_o) + \frac{\rho(\hat{\omega}_i, \hat{\omega}_o)(\hat{N} \cdot \hat{\omega}_i) L_i(\hat{\omega}_i)}{p(\hat{\omega}_i)}.$$

By examining the separate parts of this approximation, we find that we need to be able to do five non-trivial operations.

First, we need to be able to generate a sample direction $\hat{\omega}_i$ according to the probability density function (PDF) $p(\hat{\omega}_i)$. This is necessary to keep the Monte Carlo estimate consistent.

Second, we need to be able to evaluate the emitted radiance L_e in the direction $\hat{\omega}_o$. This is usually done by storing a value (or multiple values) for L_e as a material property.

Third, we must be able to evaluate the BRDF ρ for incident and outgoing directions $\hat{\omega}_i$ and $\hat{\omega}_o$. This is also usually stored as a material property, either as a lookup table, or (more commonly) as an analytical formula.

Fourth, we must be able to evaluate the PDF p for a given direction $\hat{\omega}_i$. This is also usually done analytically. Once again, the PDF p must be the same distribution used to generate the sample directions $\hat{\omega}_i$. If this is not true, the Monte Carlo estimate will be inconsistent. There are a number of potential choices for this PDF. Uniform sampling is always an option. Other possibilities include sampling based on incident lighting or sampling based on the shape of the BRDF. Many BRDF models describe a method that can be used to sample directions according to the BRDF, and give the corresponding PDF over directions.

Finally, we must be able to evaluate the incident radiance L_i from the direction $\hat{\omega}_i$. This is done recursively: a ray is traced into the scene, and the incident radiance along the sample direction is estimated using the same approximation. If the ray does not hit an object, we return an environment color for the ray's direction. Also, if the number of rays in the path exceeds a certain limit, we terminate the recursive process.

2 Diffuse / Specular BRDFs

We have already discussed diffuse (Lambertian) BRDFs. The Lambertian BRDF can be written as:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \frac{R(\lambda)}{\pi},$$

where $R(\lambda) : \Re \rightarrow [0, 1]$ gives the reflectance of the material for different wavelengths of incident light.

A useful PDF for sampling Lambertian materials is given by:

$$p(\hat{\omega}_i) = \frac{\hat{N} \cdot \hat{\omega}_i}{\pi}.$$

We generate sample directions according to this PDF by generating a random point (θ, ϕ) in spherical coordinates as follows:

$$\theta = \arcsin(\sqrt{\xi_1})$$

$$\phi = 2\pi\xi_2,$$

where the ξ_i are canonical uniform random variables (i.e., as generated by a call to `drand48()`). We can then use the standard conversion between Cartesian and spherical coordinates to get a direction vector:

$$\hat{\omega}_i = (\cos\phi \sin\theta)\hat{u} + (\sin\phi \sin\theta)\hat{v} + (\cos\theta)\hat{N}.$$

Where $\hat{u}$ and $\hat{v}$ are orthonormal vectors describing the local surface orientation, and $\hat{N}$ is the surface normal vector, which has unit length and is orthogonal to both $\hat{u}$ and $\hat{v}$.

Diffuse BRDFs are useful in describing matte materials. But most materials in the real world are not matte surfaces. Mirrors are not matte at all; they are perfectly shiny, and other materials fall somewhere on the scale between Lambertian and mirrored surfaces. Many BRDF models describe general surfaces using a weighted combination of a diffuse BRDF and a specular BRDF:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = R_d(\lambda)\rho_d(\hat{\omega}_i, \hat{\omega}_o) + R_s(\lambda)\rho_s(\hat{\omega}_i, \hat{\omega}_o).$$

In this equation, R_d describes the spectral diffuse reflectance and R_s describes the spectral specular reflectance. To be physically plausible, these functions must satisfy the following constraint for all λ:

$$R_d(\lambda) + R_s(\lambda) \leq 1,$$

that is, the total amount of reflected light cannot be greater than the total amount of incident light. The terms ρ_d and ρ_s are diffuse and specular BRDFs, respectively. For example, if we use the Lambertian BRDF for ρ_d, we get:

$$\rho_d(\hat{\omega}_i, \hat{\omega}_o) = \frac{1}{\pi}.$$

We now know how to describe diffuse BRDFs and how to combine difffuse and specular BRDFs to represent general-purpose materials. The rest of this document focuses on specular BRDFs; specifically, it describes two specular BRDF models, along with useful PDFs and methods for generating samples according to the BRDFs.

3 Interlude: the Halfway Vector

An important concept in many specular BRDFs is the idea of the *halfway vector*. The normalized halfway vector $\hat{H}$ is written as the normalized average of the incident and outgoing directions:

$$\hat{H} = \frac{\hat{\omega}_i + \hat{\omega}_o}{\|\hat{\omega}_i + \hat{\omega}_o\|}.$$

If we think of the surface as made up of a number of small, perfectly reflective (i.e., mirrored) planar facets, then the halfway vector represents the surface normal of those facets that would reflect light from the given incident direction to the given outgoing direction. In fact, several BRDF models are based directly on this "microfacet" view of a surface.

We assume the microfacets are extremely small, so instead of modelling the geometry for every facet, we describe a surface as a distribution of microfacets. Usually this is done by describing a PDF $p_h(\hat{H})$ over the set of possible halfway vectors that describes how microfacets are distributed across the surface. For example, a smooth surface would have a distribution tightly clumped around the surface normal direction, while a more diffuse surface would have a wider distribution of halfway vectors.

4 The Ward BRDF Model

The Ward BRDF model can represent diffuse / specular surfaces. A Lambertian BRDF is used for the diffuse portion, as discussed earlier. The specular portion of the BRDF is given by the following formula:

$$\rho_s(\hat{\omega}_i, \hat{\omega}_o) = \left[\frac{1}{4\pi\alpha_u\alpha_v\sqrt{(\hat{N}\cdot\hat{\omega}_i)(\hat{N}\cdot\hat{\omega}_o)}} \right] \exp\left(-\frac{[(\hat{H}\cdot\hat{u})/\alpha_u]^2 + [(\hat{H}\cdot\hat{v})/\alpha_v]^2}{(\hat{H}\cdot\hat{N})^2} \right).$$

This specular portion of the Ward BRDF only has two parameters: α_u and α_v, which represent the standard deviations of the surface slope in the $\hat{u}$ and $\hat{v}$ directions, respectively. The values should be non-negative. For smaller values of α_u and α_v, the material should appear shinier, for larger values, it should appear more diffuse.

Ward suggests using the following PDF for sampling incident directions:

$$p(\hat{\omega}_i) = \left[\frac{1}{4\pi\alpha_u\alpha_v(\hat{H}\cdot\hat{\omega}_o)(\hat{H}\cdot\hat{N})^3} \right] \exp\left(-\frac{[(\hat{H}\cdot\hat{u})/\alpha_u]^2 + [(\hat{H}\cdot\hat{v})/\alpha_v]^2}{(\hat{H}\cdot\hat{N})^2} \right).$$

As in most importance sampling schemes, this is very close to the Ward BRDF, but is not exactly the same.

Finally, we generate samples according to the PDF above by generating a (θ_h, ϕ_h) pair for the halfway vector as follows:

$$\phi_h = \arctan\left(\frac{\alpha_v}{\alpha_u} \tan(2\pi\xi_1) \right)$$

$$\theta_h = \arctan\sqrt{\frac{-\log\xi_2}{[\cos(\phi_h)/\alpha_u]^2 + [\sin(\phi_h)/\alpha_v]^2}},$$

where once again, the ξ_i are canonical uniform random variables. We now calculate $\hat{H}$ as:

$$\hat{H} = (\cos\phi_h \sin\theta_h)\hat{u} + (\sin\phi_h \sin\theta_h)\hat{v} + (\cos\theta_h)\hat{N},$$

and finally, we recover the incident sample direction $\hat{\omega}_i$ using the formula:

$$\hat{\omega}_i = 2(\hat{\omega}_o\cdot\hat{H})\hat{H} - \hat{\omega}_o.$$

5 The Ashikhmin BRDF Model

The Ashikhmin BRDF is similar to the Ward BRDF in many ways. It has a diffuse and specular component. Unlike the Ward BRDF, a non-Lambertian BRDF is used for the diffuse component. However, it is also possible to use a Lambertian BRDF, so I will not discuss the diffuse Ashikhmin term further here. The specular portion of the Ashikhmin BRDF is:

$$\rho(\hat{\omega}_i, \hat{\omega}_o) = \frac{\sqrt{(n_u+1)(n_v+1)}}{8\pi(\hat{H}\cdot\hat{\omega}_o)[(\hat{\omega}_i\cdot\hat{N}) + (\hat{\omega}_o\cdot\hat{N}) - (\hat{\omega}_i\cdot\hat{N})(\hat{\omega}_o\cdot\hat{N})]} (\hat{N}\cdot\hat{H})^{\left[\frac{n_u(\hat{H}\cdot\hat{u})^2 + n_v(\hat{H}\cdot\hat{v})^2}{1-(\hat{H}\cdot\hat{N})^2}\right]}.$$

Like the Ward BRDF, the specular portion of the Ashikhmin BRDF has two parameters: n_u and n_v. These parameters control the shininess of the material, but they work in the opposite way to those of Ward. So in this model, increasing the values of n_u and n_v will make the surface shinier, and decreasing the values will make the surface more diffuse.

The PDF for sampling according to the Ashikhmin BRDF is:

$$p(\hat{\omega}_i) = \frac{\sqrt{(n_u+1)(n_v+1)}}{8\pi(\hat{\omega}_i\cdot\hat{H})} (\hat{N}\cdot\hat{H})^{\left[\frac{n_u(\hat{H}\cdot\hat{u})^2 + n_v(\hat{H}\cdot\hat{v})^2}{1-(\hat{H}\cdot\hat{N})^2}\right]}.$$

As with the Ward BRDF, we generate halfway vector samples, which we then convert into incident direction samples. The formulas for sampling θ_h and ϕ_h with the Ashikhmin model are:

$$\phi_h = \arctan\left(\sqrt{\frac{n_u+1}{n_v+1}}\tan\left(\frac{\pi\xi_1}{2}\right)\right)$$

$$\theta_h = \arccos\left((1-\xi_2)^{\frac{1}{n_u\cos^2(\phi_h)+n_v\sin^2(\phi_h)+1}}\right).$$