

Importance Sampling of Reflection from Hair Fibers

Christophe Hery
Pixar Animation Studios

Ravi Ramamoorthi
University of California, Berkeley

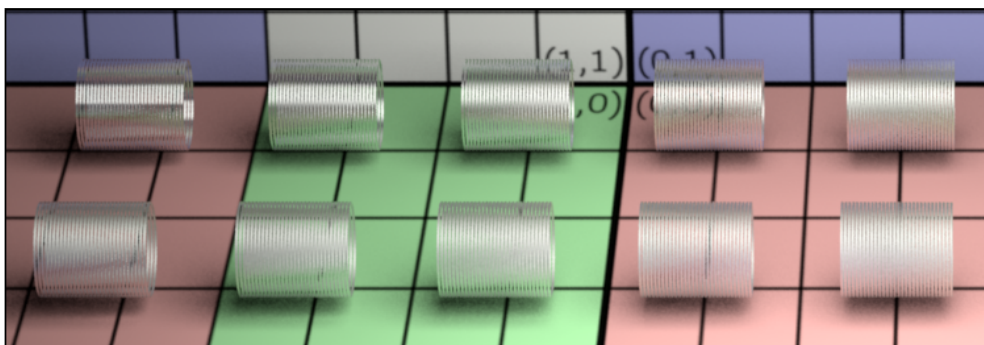


Figure 1. Our importance sampling method for the Marschner specular lobe BRDF. From left to right, top to bottom, cone angle $\beta = 1, 3, 5, 7, 9, 11, 13, 15, 17$ and 19° .

Abstract

Hair and fur are increasingly important visual features in production rendering, and physically-based light scattering models are now commonly used. In this paper, we enable efficient Monte Carlo rendering of specular reflections from hair fibers. We describe a simple and practical importance sampling strategy for the reflection term in the Marschner hair model. Our implementation enforces approximate energy conservation, including at grazing angles by modifying the samples appropriately, and includes a Box-Muller transform to effectively sample a Gaussian lobe. These ideas are simple to implement, but have not been commonly reported in standard references. Moreover, we have found them to have broader applicability in sampling surface specular BRDFs. Our method has been widely used in production for more than a year, and complete pseudocode is provided.

1. Introduction

Hair and fur are important visual features, that are increasingly common in production environments. They are also the building blocks for accurate rendering of seemingly unrelated effects such as clothing, where we model individual fibers of yarn. Standard surface reflection algorithms no longer apply directly, since a hair fiber does not have a surface normal in the conventional sense (a single pixel corresponds to the entire cylinder of micro-surface normals), but only an overall orientation or tangent direction.

For many years, the standard hair reflection model was the extension of the Phong model proposed by Kajiya and Kay [1989]. This model was adapted for production by Goldman [1997]. In 2003, [Marschner et al. 2003] proposed a comprehensive physically-based light scattering model from human hair fibers, that has become the basis for most subsequent work, including this paper. While the Marschner model defines an effective hair “BRDF”, efficient Monte Carlo rendering also requires practical techniques for importance sampling (we use BRDF importance sampling [Lawrence et al. 2004] within a multiple importance sampling framework [Veach and Guibas 1995]). To date, no importance sampling method has been published, and personal communications indicate the lack of widespread existence of such a method.

In this paper, we describe a simple and practical importance sampling scheme for the single scattering or reflection term R in the Marschner hair model. While we do not address the other (TT and TRT) terms, they are often considered separately for easier artistic design [Sadeghi et al. 2010], and have a similar form. In fact, TRT is commonly split between a glint component (which can be achieved via some noise calls), and a regular reflection lobe, derived from R , only with a longitudinal shift in the opposite direction. And TT , the double transmittance, is usually blocked because of shadowing and thus requires a more global scatter approach for blonde hair [Zinke et al. 2008].

Our method was originally developed simultaneously with the recent comprehensive work by d’Eon et al. [2011] and addresses some of the same issues. In particular, they correctly deal with energy conservation, modifying the BRDF accordingly. However importance sampling is still an open problem.

In summary, we describe a simple practical approach to importance sampling the reflected lobe in the Marschner hair model. We enforce approximate energy conservation at grazing angles by clamping and simplifying the weight

of the Monte Carlo estimator, which is a “trick” that is also useful in other contexts like specular BRDF sampling. There are a number of interesting practical issues, that we describe in detail with pseudocode faithful to our actual production-ready implementation.

The panel in Figure 1 shows some results of our importance sampling method for the Marschner specular lobe BRDF inside a global illumination renderer for both direct lighting from an area light source and for tracing reflections. These images use 48 samples, and each cylinder is composed of 30 hair fibers shaped like circles and assembled next to one another. With a fairly low number of samples, each hair fiber is reflection tracing to its surrounding walls and ground. Observe that we can resolve both glossy and semi-glossy specular defined by the Marschner reflection model. Our algorithm has been used in production at Pixar for the past year in rendering hair for Monsters University and other shows (see Fig. 5 for an example). It has produced satisfactory results with no tweaking required, beyond what is reported in this paper and described in our complete pseudocode.

2. Background

The reflected radiance is given in the standard way by

$$L_r(\omega_r) = \int L_i(\omega)S(\omega, \omega_r)V(\omega) \cos \theta d\omega, \quad (1)$$

where the integral is over all incident directions, S is the scattering function (equivalent to the BRDF) for hair, L_i is the lighting, and ω and ω_r are incident and reflected directions. V is the visibility function that is ray traced, or computed using an approach like deep and multilayer shadow maps [Lokovic and Veach 2000; Xie et al. 2007]. For notational simplicity, incident directions are not subscripted. The important difference from surface reflection is that the angles are measured with respect to the normal *plane* (perpendicular to the hair tangent direction), rather than a single surface normal. Thus, θ is the incident angle to the plane, which ranges from $[-\pi/2, \pi/2]$.

Our goal is to importance sample $S(\omega, \omega_r)$ to determine incident directions, given that we know the reflected direction ω_r . In doing so, we will pick a number J of samples ω_j , and compute

$$L_r \approx \frac{1}{J} \sum_{j=1}^J \frac{L_i(\omega_j)S(\omega_j, \omega_r)V(\omega_j)}{p(\omega_j)} \frac{\cos \theta_j}{\cos^2 \theta_{d,j}}, \quad (2)$$

where in the standard way, the Monte Carlo estimator takes the value of the integrand divided by the probability $p(\omega_j)$ of generating the sample. For reasons of notational simplicity in later derivations, we include the $\cos^2 \theta_d$ term in [Marschner et al. 2003] in the denominator above, outside of S . We will see that this term approximately cancels, and in any event we do not attempt to importance sample it.

To do the importance sampling, we need to know the form of the scattering function, which is given in [Marschner et al. 2003] by

$$S(\theta_i, \phi_i, \theta_r, \phi_r) = M(\theta_i, \theta_r)N(\phi_i, \phi_r). \quad (3)$$

This form is already factored, allowing us to use many of the techniques for BRDF importance sampling as in Lawrence et al. [2004]. Note that we have not explicitly considered the Fresnel term, nor the division by $\cos^2 \theta_d$ (included directly in equation 2). If the Fresnel term is desired, it can simply multiply the value of the estimator (numerator in equation 2), but we do not consider it in the importance sampling itself.

Finally, for the reflection lobe, we use the formulae,

$$\begin{aligned} M &= g(\beta, \theta_h - \alpha) \\ N &= \cos \frac{\phi}{2}, \end{aligned} \quad (4)$$

where g is a (normalized) Gaussian lobe, θ_h is the half-angle between incident and reflected directions, α is an offset to capture the shift in reflected angle because of the tilt of the surface scales, and $\phi = \phi_r - \phi_i$ is the azimuthal angle in the range $[-\pi, \pi]$. The form for M is directly from [Marschner et al. 2003], while the form for N is a common simplification that can be derived [Sadeghi et al. 2010]. Note also that unlike surface reflection, we are considering angles to the normal plane, so the formula for θ_h is just $\theta_h = (\theta_i + \theta_r)/2$.

This paper now describes how to generate sample directions in a probability distribution corresponding to the scattering function in equation 3, and how to compute equation 2.

Algorithm 1. (Algorithm Pseudocode for Generating BRDF samples)

```
1 void sample (vector  $\omega_r$ ; vector Rnd [ ]; in Geom; out BRDFsamp)
    // Basic Setup, compute  $\theta_r$  and  $\phi_r$ 
2 float  $\phi_r = \text{atan}(\omega_r[2], \omega_r[1])$ ;
3 float  $\theta_r = \frac{\pi}{2} - \text{acos}(\omega_r[0])$ ;
4 float  $\theta_{\max} = \pi/2 - \text{abs}(\theta_r/2 - \alpha)$ ;

    // Now, loop over the required number of samples
5 uniform float k;
6 for ( $k = 0$ ;  $k < \text{numDirections}$ ;  $k += 1$ ) do
7     vector Rand = Rnd[k];
    // Box-Muller Transform for sampling  $M$ 
8     float  $\theta_s = \beta * \text{sqrt}(-2.0 * \log(\text{Rand}[0])) * \cos(2\pi * \text{Rand}[1])$ ;
    // Account for edge conditions
9     if ( $\text{abs}(\theta_s) > \theta_{\max}$ ) then
10         $\theta_s = \text{sign}(\theta_s) * \theta_{\max}$ ;
11    end
12     $\theta_h = \theta_s + \alpha$ ; // Account for tilt from cuticle scales
13     $\theta_i = 2.0 * \theta_h - \theta_r$ ; // Convert to  $\theta_i$ 
14    if ( $\text{abs}(\theta_i) > \pi/2$ ) then
15         $\theta_i = \text{sign}(\theta_i) * (\pi - \text{abs}(\theta_i))$ ; // Set  $\theta_i$  to  $[-\pi/2, \pi/2]$ 
16    end
17    float  $\text{cos}_i = \cos(\theta_i)$ ; // Frequently used trig function
    // Inverse-CDF for  $N$  and generate sample direction
18    float  $\Delta\phi = 2.0 * \text{asin}(2.0 * \text{Rand}[2] - 1.0)$ ;
19    float  $\phi_i = \phi_r + \Delta\phi$ ;
20    vector  $\omega_i = \text{vector}(\sin(\theta_i), \text{cos}_i * \cos(\phi_i), \text{cos}_i * \sin(\phi_i))$ ;
21    BRDFsamp  $\rightarrow \text{dir}[k] = \text{Geom} \rightarrow \text{transformFromLocal}(\omega_i)$ ;

    // Sample weights and pdf
22    uniform float  $\text{denom} = -0.5 / \beta / \beta$ ;
23    float  $M = \frac{1}{\beta\sqrt{2\pi}} * \exp(\theta_s * \theta_s * \text{denom})$ ;
24    float  $N = 2.0 * \text{sqrt}(\text{Rand}[2] * (1.0 - \text{Rand}[2]))$ ; //  $\cos \text{asin}(u)$ 
25    BRDFsamp  $\rightarrow \text{pdf}[k] = M * N / (8.0 * \text{cos}_i)$ ;
    // If desire  $\cos\theta_d$ :  $\text{cos}_d = \max(\cos((\theta_i - \theta_r)/2), 1.0e-5)$ 
    // BRDFsamp  $\rightarrow \text{wt}[k] = K_s * (\text{cos}_i * \text{cos}_i) / (\text{cos}_d * \text{cos}_d)$ ;
    // Simpler practical form, that conserves energy below
26    BRDFsamp  $\rightarrow \text{wt}[k] = K_s$ ;
27 end
```

3. Sampling

This section is the main body of the paper and describes how to generate the ω_j samples, and assign their directions, probabilities and weights of the Monte Carlo estimator in equation 2. Section 4 discusses some refinements needed for multiple importance sampling. We include pseudocode for the entire process, in Algorithm 1. Our system is implemented as a RenderMan shader, and the pseudocode is taken directly from our source code, with only minor editing for readability and to conform to the notational conventions in the text.

3.1. Basic Setup

The basic sampling function definition takes as inputs the reflected direction ω_r , an array of random numbers *Rnd* (each element is a vector since as we shall see, we will need 3 independent random numbers), and a structure for the geometry (that will be used to transform into local coordinates later). The output will be the BRDF samples (their directions, weights and probabilities for computing the estimator). The random numbers can be generated in the standard way, with stratified or quasi-Monte Carlo methods. ω_r is assumed to be available in a local coordinate frame aligned with $u - v - w$ directions as in Marschner et al.'s work [2003], where u is the tangent along the hair, and v and w represent the normal plane. We first compute θ_r and ϕ_r .

3.2. Box-Muller for Sampling the Gaussian for M

We begin by generating samples according to the Gaussian for the M term. The standard approach is based on an inverse-cumulative distribution function. However, the Gaussian cannot be analytically integrated and inverted, which means we would need to resort to computing the *inverse erf* function or numerical inversion. While *erf* is a standard numerical function in most packages including RenderMan, the *inverse erf* is found in Mathematica and Matlab, but is not standard in most shading languages, including RenderMan. There are a number of routines to compute the *inverse erf* [Acklam 2003], but they can be expensive and difficult to port. Instead, there is a simple trick using two random variables known as the Box-Muller transform [Box and Muller 1958].

While the usual derivation is for the 2D normal distribution, each random variable X or Y is also a normal distribution, and we can use either

for the 1D Gaussian for M . Note that unlike standard inverse-CDF methods, we are using two random numbers to generate a single sample. There are also rejection sampling-based methods to avoid the trigonometric calculations, but we did not use them in our implementation. The Wikipedia page (http://en.wikipedia.org/wiki/Box_muller) on the Box-Muller transform [Box and Muller 1958] happens to have an excellent discussion of the alternatives.

In line 8 of the algorithm, we first sample the 1D Gaussian to generate θ_s . Note that the basic Box-Muller value is multiplied by β to account for the variance. To obtain θ_h , we will now account for the offset α and edge effects.

3.3. Accounting for Edge Cases

The normal distribution function or Gaussian has no limits on its domain, but angles must generally be within $[-\pi, \pi]$. This leaves the question of how to handle samples that lead to angles outside these limits (not so much a problem for θ_s itself, but for the result in θ_i). Note this only occurs in the tail of the Gaussian and so any suitable method will generally lead to only minimal bias. However, these edge cases must be addressed explicitly in some way to avoid generating numerical garbage.

One physically-based approach [Lawrence et al. 2004] is to simply set the weight for these samples to 0. While this is physically accurate for the BRDF as written, it leads to some potentially undesirable properties with losing some incident energy; a constant white dome with a specular albedo of 1 should ideally reflect an energy of 1, but setting samples to 0 loses energy. An alternative would be to reject those samples and renormalize the weight of the remaining samples. This approach is reasonable, but wastes samples. (Note that in neither case do we explicitly fire rays for the samples in question, but performing the computations to generate the sample is in itself wasteful; we would ideally like to fully use all samples that we generate).

Therefore, in practice, we impose a maximum value on the domain of the Gaussian and clamp to that. This only affects samples deep in the tail, and this clamping introduces minimal bias. More sophisticated changes to the Gaussian function itself, to handle this in a more principled fashion, as explored by d'Eon [2011], are a subject of future work. In particular, in line 4, we compute the maximum value for θ_s to ensure that $|\theta_i| < \pi$ (simple algebra will verify the result), and in line 10 we clamp the sampled value to this maximum. Only then do we compute the values for θ_h and finally θ_i .

3.4. Inverse CDF for sampling N and final sample direction

We now apply a fairly standard inverse-CDF method for sampling the N term. Recall from equation 4 that $N(\phi) = \cos(\phi/2)$. Note that ϕ lies in the interval from $[-\pi, +\pi]$. However, to convert this to a pdf, we need to normalize by a factor of 4. The PDF and CDF are simply

$$\begin{aligned} \text{pdf}(\phi) &= \frac{1}{4}N(\phi) = \frac{\cos(\phi/2)}{4} & \frac{1}{4} \int_{-\pi}^{\pi} \cos \frac{\phi}{2} d\phi = 1 \\ \text{cdf}(\phi) &= \frac{1}{2} \left(1 + \sin \frac{\phi}{2} \right), \end{aligned} \quad (5)$$

where the offset is to ensure the CDF is 0 at $\phi = -\pi$. Inverting this directly gives line 18 in the pseudocode.

We can now go ahead and construct the incident vector or geometric sampling direction (note that the construction is in the hair coordinate system, and therefore somewhat different from the standard spherical coordinates). Finally, we transform this into the appropriate reference frame.

3.5. Sample pdf

Finally, we need to compute the estimator in equation 2, which requires both the value for a sample, as well as the probability distribution function. *Note that if we only need to do BRDF sampling, we need only the final weight (value/pdf), in line 26. The explicit pdf calculations in lines 22-25 of the pseudocode are only needed for multiple importance sampling.*

First, consider the probability of choosing a given sample direction, separately considering the angles θ_i and ϕ_i . We need to compute the probability distribution function $\text{pdf}(\theta_i, \phi_i)$ with proper normalization,

$$\int_{\theta_i=-\pi/2}^{\pi/2} \int_{\phi_i=-\pi}^{\pi} \text{pdf}(\theta_i, \phi_i) \cos \theta_i d\theta_i d\phi_i = 1, \quad (6)$$

where the cosine is needed for the solid angle measure in our hair coordinates (compare to the sine for standard spherical coordinates). To compute this probability, we observe that we have sampled so far not in terms of (θ_i, ϕ_i) but in terms of the half angle (θ_h, ϕ) . In other words, we actually have,

$$\int_{\theta_h} \int_{\phi} M(\theta_h) \frac{N(\phi)}{4} d\theta_h d\phi = 1, \quad (7)$$

where the normalizing factor of 4 is because $\text{pdf}(\phi) = (1/4) \cos \phi/2$. To convert this to the form of equation 6, we must change variables, or use the Jacobian $J(\theta_h, \phi; \theta_i, \phi_i) = \partial(\theta_h, \phi)/\partial(\theta_i, \phi_i)$, with a term for the area measure $|\det(J)|$ as is standard for change of variables for integration.

In our case, the Jacobian is even simpler than in surface reflection, since we simply have $\phi = \phi_r - \phi_i$ and $\theta_h = (\theta_i + \theta_r)/2$. From this, it is clear that $|d\phi| = |d\phi_i|$ and $d\theta_h = d\theta_i/2$. Since the Jacobian is diagonal, the factor $|\det(J)|$ is simply 1/2. Therefore,

$$\begin{aligned} & \int_{\theta_i} \int_{\phi_i} \frac{M(\theta_h)}{2} \frac{N(\phi)}{4} d\theta_i d\phi_i = 1 \\ \Rightarrow & \int_{\theta_i} \int_{\phi_i} \frac{M(\theta_h)}{2 \cos \theta_i} \frac{N(\phi)}{4} \cos \theta_i d\theta_i d\phi_i = 1, \end{aligned} \quad (8)$$

where in the last line we have accounted for the $\cos \theta_i$ in the solid angle measure. By inspection (compare to equation 6) from the equation above,

$$\text{pdf}(\theta_i, \phi_i) = \frac{M(\theta_h)N(\phi)}{8 \cos \theta_i}, \quad (9)$$

which is directly expressed in line 25 of the pseudocode.¹ Line 24 introduces a neat trick to avoid explicitly applying a trigonometric function. We know that $\Delta\phi$ is obtained by an inverse sine. Noting that $\cos(\sin^{-1}(u)) = \sqrt{1-u^2}$ and simplifying the algebra, we obtain the result.

3.6. Computing the Estimator and Energy Conservation

Finally, we must compute the sample's contribution to the estimator in equation 2. One condition we would like to ensure is *energy conservation*, that the hair appears uniform when placed in a lighting dome of uniform radiance. This requires the hair BRDF to be properly normalized. In our case, it will be a probability function, essentially requiring the scattering function to have the same normalization as the pdf. Therefore, we use

$$S(\omega_i, \omega_r) = \frac{M(\theta_h)N(\phi)}{8}, \quad (10)$$

¹The factors of 2 and 8 in lines 24 and 25 could easily be pre-cancelled and other trivial optimizations applied. We retain the original form for readability.




Figure 2. While our BRDF sampling offers full convergence at 1 sample under a uniform white dome with no shadowing (fig a), light sampling needs all the way to 256 samples (fig d) to resolve the correct normalization. For reference, we show images with 1 and 16 light samples (fig b and c).

from which it follows that the reflectance-dependent part in equation 2 is given by

$$\frac{S(\omega_i, \omega_r)}{p(\omega_i)} \cdot \frac{\cos \theta_i}{\cos^2 \theta_d} = \frac{\cos^2 \theta_i}{\cos^2 \theta_d}, \quad (11)$$

since all other factors involving M and N cancel. Indeed, this is the beauty of good importance sampling, that most factors cancel, leaving an estimator with very low variance. We robustly compute (avoiding small values) the cosine denominator $\cos \theta_d = \max(\cos((\theta_i - \theta_r)/2), 1.0e - 5)$. In the pseudocode, we also include the overall specular color K_s in the weight.

Our final form in line 26 is even simpler. For sharp specular lobes, incident and reflected cosines will be very similar, as will that of the difference angle. Thus, the right-hand side in equation 11 can simply be replaced with 1. We also note that Marschner et al.’s original derivation [2003] uses a mirror where $\cos \theta_i = \cos \theta_r = \cos \theta_d$, and the rationale for using a denominator with θ_d for rough surfaces, as opposed to θ_i is not clear, except from conditions of reciprocity. Therefore, we directly use the very simple form in line 26, and we have not found this to change the results significantly. Besides simplicity, this formula enforces a form of exact energy conservation; the scattering function is now exactly a probability distribution function (with edge cases handled not with an analytic formula, but implicitly through our earlier discussion; implicitly both the probabilities and value of the scattering function are modified in the same way to give a net Monte Carlo weight of 1).²

² Note that the first part of equation 6, and by extension equation 8 requires the Gaussian normal distribution function to integrate to 1, which it does over an infinite domain. The integral is approximately 1 over the restricted angular domain, but our computations do not




Figure 3. As can be seen from the figure, BRDF sampling converges rapidly to the true result (ground truth was obtained using light sampling with 256 samples). This example includes full shadow tracing.

Finally, the overall rendering system will take the weights produced from the BRDF sampler, and multiply them with the lighting for the sample directions, modulated by visibility, and average over all Monte Carlo samples. Note that the overall rendering system cares only about the weight and the BRDF direction. However, we do compute the pdf explicitly, both for instructive purposes, and since it is useful for multiple importance sampling, as discussed next.

Refer to Figure 2, where a common validating "white furnace" test is done. The idea is to make sure that the BRDF correctly integrates to white under a non shadowing uniform white dome. Similarly Figure 3 demonstrates the fast convergence of our approach under arbitrary lighting with full shadowing computations.

strictly account for the way we handled edge cases to clamp the range of values. This does not create practical problems, especially since the estimator is also set up to compensate and ensure energy conservation, as discussed above.

Algorithm 2. (BRDF Value and PDF for Multiple Importance Sampling)

```
1 void ValuePDF (vector  $\omega_r$ ; vector Dir [ ]; out BRDFsamp)
    // Basic Setup, compute  $\theta_r$  and Rperp to calculate  $\phi$ 
2 float  $\theta_r = \frac{\pi}{2} - \text{acos}(\omega_r[0])$ ;
3 vector Rperp = normalize (vector (0.0,  $\omega_r[1]$ ,  $\omega_r[2]$ ) );
4 uniform float denom =  $-0.5 / \beta / \beta$ ;

    // Now, loop over the required number of samples
5 uniform float k ;
6 for ( $k = 0$ ;  $k < \text{numDirections}$ ;  $k += 1$ ) do
    // Compute  $M$  term
7 vector  $\omega_i = \text{Dir}[k]$ ;
8 float  $\theta_i = \frac{\pi}{2} - \text{acos}(\omega_i[0])$ ;
9 float  $\text{cos}_i = \text{sqrt}(1.0 - \omega_i[0] * \omega_i[0])$ ;
10 float  $\theta_s = (\theta_i + \theta_r) / 2 - \alpha$ ;
11 float  $M = \frac{1}{\beta \sqrt{2\pi}} * \exp(\theta_s * \theta_s * \text{denom})$ ;

    // Compute  $N$  term
12 vector Lperp = normalize (vector (0.0,  $\omega_i[1]$ ,  $\omega_i[2]$ ) );
    // Trig identity  $\cos \phi = 2 \cos^2(\phi/2) - 1$ 
13 float  $N = \text{sqrt}((1.0 + \text{Rperp} \cdot \text{Lperp}) * 0.5)$ ;

    // Compute Value and PDF
14 BRDFsamp  $\rightarrow \text{pdf}[k] = M * N / (8.0 * \text{cos}_i)$ ;
15 BRDFsamp  $\rightarrow \text{value}[k] = K_s * \text{BRDFsamp} \rightarrow \text{pdf}[k]$ ;
    // If we desire to keep the  $\cos \theta_d$  term, we can
    multiply this by  $(\cos \theta_i / \cos \theta_d)^2$ .
16 end
```

4. Multiple Importance Sampling

In practice, the BRDF sampling routine above will often be combined with light sampling in a multiple importance sampling (MIS) framework [Veach and Guibas 1995]. One requirement of MIS is that we are able to compute the BRDF value and pdf for an arbitrary direction generated by light sampling. Moreover, there may be cases where we want to use light sampling; we still need to be able to evaluate the normalized BRDF value in those cases for an arbitrary incident direction. Therefore, we describe the Value and PDF

function in algorithm 2, which is largely similar to BRDF sampling.

The main difference is that we now are given as input a list of incident directions (in Dir []). We start with the basic setup as for BRDF sampling, computing θ_r and the constant *denom* term. Instead of computing ϕ_r , we compute the corresponding vector *Rperp* instead (a similar vector *Lperp* will be computed later for incident directions).

We now proceed to compute *M*, starting by reading in ω_i and determining θ_i . Note that we have ω_i so the cosine can be computed directly without a trigonometric function call. θ_s is now computed directly from the formula (since $\theta_h = \theta_s + \alpha$). From this, we apply the standard formula for *M*.

For computing *N*, we compute *Lperp* for the incident direction, just as we calculated *Rperp*. $\cos\phi$ is simply the dot product between these vectors. $\cos(\phi/2)$ is obtained directly from a well-known trigonometric identity, without using any explicit trigonometric functions.

Finally, we need to compute the value and pdf. The pdf is computed just as for the BRDF sampling case, discussed earlier. The value is simply the final weight times the pdf, and we have already seen that the weight is simply K_s .

Figure 4 compares light sampling with BRDF sampling and MIS. In this environment, the advantage of MIS is rather minimal. However, as is common with MIS, situations with more isolated light sources would demonstrate benefits from light sampling.

5. Discussion and Use in Production

The method described in this paper was originally developed for production use for hair rendering on Pixar's upcoming Monsters University feature. It is one component of a significant shift involving the computer-generated animation industry, where previous ad-hoc shading models within a rasterization pipeline are increasingly being replaced by physically accurate lighting and reflectance, within a raytracing and importance sampling framework. Indeed, from mid-2011, Pixar's industry standard Renderman 16.0 software has included support for (multiple) importance sampling, inspired in large part by our initial shaders for this purpose (which in addition to the work described in this paper also handled standard diffuse and specular BRDFs in a similar fashion). Since our application is to hair rather than surface BRDFs, and the original development pre-dates this change in RenderMan, our actual implementation uses independent shaders, pseudocode for which is given here.

The technique described in this paper has been in production use at Pixar




Figure 4. Here we compare the quality of BRDF sampling versus light sampling. Under this environment, 4 BRDF samples (fig c) are equivalent to about 24 light samples (fig b). We also provide (fig d) an MIS render (a combination of light 4 samples and 4 BRDF samples): the differences are rather subtle since light sampling is much worse than BRDF sampling, but if you look closely, you can see better definition of the individual strands.

for Monsters University and other shows for more than a year now, and has been generally well received, with almost no additional tweaks required beyond what is described here. To our knowledge, it has performed satisfactorily in all settings. Given the “in-production” nature of Monsters University, we are unable to provide too many example images at this time. Figure 5 shows one example from the publicly-released trailer. We see how the parameters (reflectance K_s , color, width of highlight defined by β and shift α) can be used to create interesting appearances, allowing sufficient flexibility for artistic direction. In this case, the image actually uses two Marschner lobes (both sampled with our algorithm) for creating the right look.

While many of the basic ideas in this paper are a fairly direct application of the literature, our informal discussions with many other production houses indicated they were not aware of, nor able to independently develop, a suitable



Figure 5. An example image in production from *Monsters University*, with the hair rendered using our method. This image uses a small diffuse component, but is primarily rendered with our specular model. It uses two Marschner lobes, one with $K_s = 0.17, \beta = 5^\circ, \alpha = -2^\circ$ with a white color, and a secondary lobe with $K_s = 1, \beta = 10^\circ, \alpha = 5^\circ$ with a saturated blue color. Image copyright (2012) Pixar. All Rights Reserved.

method for importance sampling. We are therefore publishing our complete implementation in the hope it is more broadly useful to the industry. There is also considerable room for future work, such as including the other Marschner terms, like TT and TRT, as well as in applying the concepts to more recent hair BRDF models such as d'Eon's [2011].

Acknowledgements

We would like to thank Jaakko Lehtinen, Morgan McGuire, and the anonymous reviewers for their helpful comments.

References

- ACKLAM, P., 2003. An algorithm for computing the inverse normal cumulative distribution function. Available online, <http://home.online.no/~pjacklam/notes/invnorm/>. 6
- BOX, G., AND MULLER, M. 1958. A note on the generation of random normal deviates. *The Annals of Mathematical Statistics* 29, 2, 610–611. 6, 7
- D'EON, E., FRANCOIS, G., HILL, M., LETTERI, J., AND AUBRY, J. 2011. An energy-conserving hair reflectance model. *Computer Graphics Forum (EGSR 11)* 30, 4, 1181–1187. 2, 7, 15
- GOLDMAN, D. 1997. Fake fur rendering. In *SIGGRAPH 97*, 127–134. 2
- KAJIYA, J., AND KAY, T. 1989. Rendering Fur with Three Dimensional Textures. In *SIGGRAPH 89*, 271–280. 2
- LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHY, R. 2004. Efficient BRDF Importance Sampling Using a Factored Representation. *ACM Transactions on Graphics (SIGGRAPH 2004)* 23, 3. 2, 4, 7
- LOKOVIC, T., AND VEACH, E. 2000. Deep Shadow Maps. In *SIGGRAPH 00*, 385–392. 3
- MARSCHNER, S., JENSEN, H., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. *ACM Transactions on Graphics (SIGGRAPH 03)* 22, 3, 780–791. 2, 4, 6, 10
- SADEGHI, I., PRITCHETT, H., JENSEN, H., AND TAMSTORF, R. 2010. An artist friendly hair shading system. *ACM Transactions on Graphics (SIGGRAPH 10)* 29, 4. 2, 4
- VEACH, E., AND GUIBAS, L. 1995. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *SIGGRAPH 95*, 419–428. 2, 12

- XIE, F., TABELLION, E., AND PEARCE, A. 2007. Soft shadows by ray tracing multilayer transparent shadow maps. In *EuroGraphics Symposium on Rendering*, 265–276. 3
- ZINKE, A., YUKSEL, C., WEBER, A., AND KEYSER, J. 2008. Dual scattering approximation for fast multiple scattering in hair. *ACM Transactions on Graphics (SIGGRAPH 08)* 27, 3. 2

Author Contact Information

Christophe Hery
Pixar Animation Studios
1200 Park Avenue
Emeryville, CA 94608
chery@pixar.com

Ravi Ramamoorthi
525 Soda Hall
University of California
Berkeley, CA 94720
ravir@eecs.berkeley.edu

Hery and Ramamoorthi, Importance Sampling of Reflection from Hair Fibers,
Journal of Computer Graphics Techniques (JCGT) 1, 1, 1–17, 2012
<http://jcgt.org/published/0001/01/01/>

Received: June 15, 2012

Recommended: June 16, 2012

Published: June 22, 2012

Corresponding Editor: Jaakko Lehtinen

Editor-in-Chief: Morgan McGuire

© 2012 Hery and Ramamoorthi (the authors).

The authors provide this document (the Work) under the Creative Commons CC BY-ND 3.0 license available online at <http://creativecommons.org/licenses/by-nd/3.0/>. The authors further provide the first page as a separable document under the CC BY-ND 3.0 license, for use in promoting the work.

