

# Importance Sampling for Physically-Based Hair Fiber Models


Eugene d'Eon<sup>1</sup>

Steve Marschner<sup>2</sup>

Johannes Hanika<sup>1</sup>

<sup>1</sup>Weta Digital

<sup>2</sup>Cornell University



**Figure 1:** Our new importance sampling strategy allows easy inclusion of Marschner and related hair reflectance functions in physically-based Monte Carlo renderers. Here we show hair volumes illuminated by environment maps and area lights with unbiased global illumination (computed using a forward path-tracer with multiple importance sampling). Our sampling strategy requires no precomputation, so it is easy to vary the absorption along the fiber (second image), and to add noise to the index of refraction, roughness, and scale tilt to create subtle heterogeneity along each fiber. Each image is 1024 samples/pixel.

## Abstract


We present a new strategy for importance sampling hair reflectance models. To combine hair reflectance models with increasingly popular physically-based rendering algorithms, an efficient sampling scheme is required to select scattered rays that lead to lower variance and noise. Our new strategy, which is tied closely to the derivation of physically-based fiber functions, works well for both smooth and rough fibers based on the Marschner et al. model and also for Lambertian fibers. It should be directly usable with future hair reflectance models that allow for more general cross-sections and more complex surface properties, provided the lobes are derived in a similar, separable fashion. Our strategy includes lobe selection and can efficiently sample complex lobe shapes like the Marschner TRT function. The scheme is easy to implement and requires no precomputation, allowing fully heterogeneous variation of all fiber parameters.

**CR Categories:** I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing;

**Keywords:** Hair scattering, BCSDf, importance sampling

## 1 Introduction

Believable computer generated characters require believable hair. For rendered hair to appear realistic the *Bidirectional Curve Scattering Distribution Function* (BCSDF) used to simulate the light



**Figure 2:** Refracted pathways through a hair fiber. Single Reflection ( $R$ ), double transmission ( $TT$ ), transmission-reflection-transmission ( $TRT$ ) and so on.

interaction with the individual hair fibers must be accurate—all of the visibly relevant behaviours seen in plausible illuminations of typical hair volumes must be well approximated.

Previously in graphics, the fiber model of Kajiya and Kay [1989] remained predominant until Marschner et al. [2003] introduced the *factored-lobe* analytic BCSDf that remains the basis of most high-accuracy, parametric hair rendering today. Factored-lobe BCSDfs decompose the reflectance into separate modes of propagation—direct reflection ( $R$ ), double transmission, ( $TT$ ) and paths with one or more internal reflections ( $TRT$ ,  $TRRT$ , ...) (Figure 2). The total reflectance function  $S$  is the sum of all such *lobes*  $S_p$ , indexed by  $p$ , the *number of internal path segments* traversed by light rays contributing to that mode (Figure 2),

$$S(\theta_i, \theta_o, \phi) = \sum_{p=0}^{\infty} S_p(\theta_i, \theta_o, \phi). \quad (1)$$

Many Monte Carlo rendering algorithms, such as path tracing, require efficient schemes for randomly selecting samples (directions and weights) with a statistical distribution that is close to the reflectance function of the material being simulated. We derive a practical, analytic, easy-to-implement importance-sampling strategy for several recent physically-based factored-lobe BCSDfs, allowing their efficient combination with these rendering algorithms.

## 2 Related work

Our new sampling strategy works for fiber models including those of Zinke and Weber [2007], who introduced the formalism of BCSDFs. They also derived a near-field model, important for close rendering of hair fibers. We use the energy-conserving longitudinal term of d'Eon et al. [2011] and Zinke and Weber's [2007] method of treating azimuthal roughness—these remove the requirement for root-solvers and careful treatment of caustics required in the original Marschner model.

The Marschner, Zinke and d'Eon models are the most physically accurate parametric models, but no analytic sampling scheme has been previously presented for these full models. Hery and Ramamoorthi [2012] presented importance sampling for the R lobe of Marschner et al. [2003] and Ou et al. [2012] presented importance sampling for the full model of Sadeghi et al. [2010]. The Sadeghi model allows more intuitive user control of the hair appearance, but isn't as physically accurate as other models. Our sampling of the  $N_R$  function closely corresponds to both of these related works, but our approach also handles the complex  $TT$  and  $TRT$  lobes that closely match the transport of light within dielectric cylinders. In addition, we handle lobe selection automatically in a fashion that works well regardless of the angle of incidence or absorption level and we present the first longitudinal sampling scheme that accounts for all angular variables.

Our sampling scheme can be used for direct illumination sampling (where indirect illumination may be approximated by Dual Scattering [Zinke et al. 2008]) and also for indirect illumination, where sampling at each vertex of a random path is used to create photon maps within hair volumes [Moon et al. 2008] or within a variety of unbiased path-tracing algorithms [Veach 1997]. For an extensive review of hair shading, animation and modeling techniques we refer the reader to the survey by Ward et al. [2007].

## 3 Hair BCSDF Importance Sampling

We begin with some notation. Following the convention of Marschner et al. [2003], directions are measured in a spherical coordinate system centered on the fiber axis, with  $\theta_i$  and  $\theta_o$  measuring *inclinations* to the normal plane of the fiber and  $\phi = \phi_o - \phi_i$  measuring the *azimuthal* difference between the incident and reflected directions. Each lobe is factored into a *longitudinal scattering function*  $M_p$  and an *azimuthal scattering function*  $N_p$ :

$$S_p(\theta_i, \theta_o, \phi) = M_p(\theta_i, \theta_o)N_p(\theta_i, \theta_o, \phi). \quad (2)$$

Note that our notation differs from previous work by including all inclination-dependent factors in  $M_p$ , such as the  $\cos^2 \theta_d$  term of the original Marschner model. Notation is simplified by referring to the longitudinal *difference angle*  $\theta_d = (\theta_o - \theta_i)/2$ . The relative *index of refraction*  $\eta$  of the hair to the surrounding medium is typically fixed at 1.55. Given  $\theta_d$ ,  $\eta$ , and the offset  $h \in [-1, 1]$  of an incoming ray from the fiber axis (Figure 2), a Bravais analysis gives the azimuthal distributions compactly using  $\gamma_i = \arcsin(h)$ ,  $\gamma_r = \arcsin(\frac{h}{\eta})$ ,  $\eta' = \frac{\sqrt{\eta^2 - \sin^2(\theta_d)}}{\cos(\theta_d)}$  to predict the relative change in azimuth

$$\Phi(p, h) = 2p\gamma_i - 2\gamma_r + p\pi \quad (3)$$

for each mode  $p$  being considered. For rough fibers we use the azimuthal functions  $N_p$  of d'Eon et al. [2011] evaluated using a 70-point Gaussian quadrature

$$N_p(\phi) = \frac{1}{2} \int_{-1}^1 dh A(p, h) D_p(v_{p,N}, \phi - \Phi(p, h)). \quad (4)$$

Here, the attenuation terms  $A(p, h)$  and the wrapped Gaussian  $D_p$  of variance  $v_{p,N}$  are those given by d'Eon et al. [2011]. The ef-

fects due to tilted scales are simulated approximately by perturbing the *specular cone angles* slightly,  $\theta_{\text{cone},R} = -\theta_i + 2\alpha$ ,  $\theta_{\text{cone},TT} = -\theta_i - \alpha$ ,  $\theta_{\text{cone},TRT} = -\theta_i - 4\alpha$ , where  $\alpha$  is the *scale tilt angle*.

We now describe our importance sampling strategy, using the hair reflectance model of d'Eon et al. [2011] as an example. Alternative longitudinal functions (and sampling for those functions) can be substituted and the remaining elements of the approach remain unchanged.

### 3.1 Strategy Outline

A good importance sampling method for a BCSDF starts with a fixed incoming direction ( $\theta_i$  and  $\phi_i$ ) and produces random outgoing directions ( $\theta_o$  and  $\phi_o$ ) with a distribution that is close to  $S(\theta_i, \theta_o, \phi) \cos^2 \theta_o$  (one cosine for the spherical coordinate system and another for the geometric projection to the fiber). We begin with lobe selection, by using uniform random sampling of a position  $h$  along the cross section of the hair fiber. Given this offset  $h$ , we first use the attenuations through a smooth fiber at this offset  $h$  for selecting a lobe, then use the same  $h$  to choose  $\theta_o$  and  $\phi$ . Because the longitudinal functions  $M_p$  in Equation 2 are independent of azimuth, we sample  $\theta_o$  first, and then sample the azimuthal direction. For the former procedure, we derive perfect importance sampling for the energy-conserving longitudinal term of d'Eon et al. [2011]. Then we show how to use Equation 3 to importance sample the azimuthal functions  $N_p$  for smooth ( $v_{p,N} = 0$ ) hairs. To sample rough hairs, we use an additional Gaussian random variable applied to the exitant azimuth predicted by the smooth hair sampling. The result is an analytic, easy to implement, and efficient sampling scheme requiring no precomputation.

### 3.2 Importance Sampling the Longitudinal M function

We use the longitudinal scattering function  $M(v, \theta_{\text{cone}}, \theta_o)$  of d'Eon et al. [2011] for all lobes

$$M(v, \theta_c, \theta_o) = \frac{\text{csch}(1/v)}{2v} e^{\frac{\sin \theta_c \sin \theta_o}{v}} I_0 \left[ \frac{\cos \theta_c \cos \theta_o}{v} \right] \frac{1}{\cos \theta_o}. \quad (5)$$

For additional numerical recipes for evaluating this function, see [d'Eon 2013]. The derivation of this  $M$  function was to convolve a Dirac circle  $\delta(\theta - \theta_{\text{cone}})$  on the surface of a sphere by a spherical Gaussian and then take the longitudinal distribution of that result to be the spreading in  $\theta$  due to roughness. To sample this with a single random number seems intractable. However, with two uniform random variables  $\xi_1$  and  $\xi_2$  each within  $[0, 1]$  we can importance sample  $M \cos^2 \theta_o$  perfectly. We start with the importance sampling of a spherical Gaussian of variance  $v$  (being careful of numerical issues for low variance [Jakob 2013])

$$u(\xi_1) = v \log \left( e^{1/v} - 2\xi_1 \sinh \frac{1}{v} \right) \quad (6)$$

and then take the final inclination  $\theta_o$  of a point chosen uniformly on the Dirac specular-cone circle parameterized by  $\phi = 2\pi\xi_2$  and tilted away from the north pole by the angle  $\theta$  sampled from the spherical Gaussian,  $\cos \theta = u(\xi_1)$ . Putting this all together, sampling  $\theta_o$  given variance  $v$  and the specular cone angle  $\theta_{\text{cone}}(\theta_i, \phi, \alpha)$  is

$$\theta_o(\xi_1, \xi_2, v, \theta_{\text{cone}}) = \arcsin(u(\xi_1) \cos \theta' + \sqrt{1 - u(\xi_1)^2} \cos(2\pi\xi_2) \sin \theta') \quad (7)$$

where  $\theta' = \frac{\pi}{2} - \theta_{\text{cone}}$ . This is perfect importance sampling, so the sample weight is 1 and the PDF is  $M \cos^2 \theta_o$ .

We note that this  $M$  function is not reciprocal due to the  $1/\cos \theta_i$  term. Lack of reciprocity is common in other forms of energy-conserving  $M$  functions that renormalize Gaussians over the finite

domain of inclinations based on the incoming angle  $\theta_i$ . Also, analogous to shading normals for surfaces, the tilted scales treatment generally yields non-reciprocal scattering functions in any case. We haven't found any issues with ignoring these details in our implementation, however.

### 3.3 Importance Sampling Azimuthal N functions

The azimuthal functions  $N_p$  of the Marschner model describe a large family of lobe shapes that model the distributions of light exiting a dielectric fiber when illuminated from a single direction. However, despite their complexity and variety, they can be sampled analytically and efficiently. The key insight is to exploit the original derivation of Marschner et al. [2003]—the fiber function is a uniform integral over the fiber cross-section  $h \in [-1, 1]$  (Equation 4) and the laws relating an incident azimuth to an exitant azimuth are compact expressions. In fact, the derivation of the original Marschner model was precisely the calculation of the outgoing densities that arise from this sampling scheme.

We first select a random cross-section offset  $h = 2\xi_h - 1$  (where  $\xi_h \in [0, 1]$  is a uniformly distributed random number). For smooth fibers, the exitant normal-plane inclination is  $\theta_o = \theta_{\text{cone}_p}(\theta_i)$  from which  $\theta_i$  is then known. In the case of no azimuthal roughness, the set of directions  $\Phi(p, 2\xi_h - 1)$ , using Equation 3, each with weight  $A(p, 2\xi_h - 1) \leq 1$ , is an efficient scheme for importance sampling the functions  $N_p$  derived by Marschner et al. [2003]. In the case  $p = 0$ , this is identical to prior sampling schemes for  $N_R$  [Hery and Ramamoorthi 2012; Ou et al. 2012].

In the case of azimuthal roughness the functions we wish to sample are the convolution of  $\Phi(p, h)$  by a wrapped Gaussian  $D_p(v_{p,N})$  with variance  $v_{p,N}$  (Equation 4). These are easily sampled by selecting relative azimuths

$$\phi = \Phi(p, 2\xi_h - 1) + g \sqrt{v_{p,N}} \quad (8)$$

where  $g$  is a Gaussian distributed random variable with standard deviation 1 (sampled with a lookup table or Box-Muller). Again, the sample weight is  $A(p, h) \leq 1$  and the PDF is Equation 4 with  $A(p, h)$  replaced by 1. We mention that sampling the  $N_p$  functions is quite cheap—much cheaper than evaluating them or computing their PDF, due to the quadrature.

### 3.4 Lobe Selection

We now return to lobe selection. The key to lobe selection is to delay the choice until after randomly selecting  $h$ . Given  $h$ , the attenuations through a smooth, ideally specular fiber,  $A_{\text{spec}}(p, h)$ , are known—they are the product of Fresnel terms and absorption terms for the single set of paths predicted for the smooth fiber (assuming  $M$  is a delta function at the ideal specular cone angle when computing  $\theta_i$ ). We use these attenuations for the smooth fiber to select a lobe  $p$  to sample, even for rough fibers. When the inclination is high and Fresnel dominates, or absorption is high, we tend to select the R lobe, for example. A random number  $\xi_p \in [0, 1]$  can be used to choose a lobe  $p$  in proportion to  $w_p = A_{\text{spec}}(p, h) / \sum_p A_{\text{spec}}(p, h)$ . Selecting  $p$  proportional to  $w_p$  incurs a selection weight of  $1/w_p$ . The attenuation through the actual path (which deviates from the specular cone path after sampling  $M$ ) has a throughput of  $A(p, h)$ , so the final sample weight is  $w = A(p, h) / w_p$ , which is  $\sum_p A_{\text{spec}}(p, h)$  times the ratio  $A(p, h) / A_{\text{spec}}(p, h)$ . Provided the roughness is reasonably low, the Fresnel and absorption terms for slightly deflected paths will not differ significantly from the ideal specular case, keeping the final sample weights low. However, for extreme inclinations and roughnesses, we clamp the upper bound of the final sample weights to

2.0 (and have noted no measurable energy loss in a variety of hair renders).

## 3.5 Sampling Summary and PDF Computation

To summarize, our importance sampling strategy for the full fiber function with rough surfaces is:

- An incoming direction  $\theta_i$  is given
- We uniformly choose a random offset  $h \in [-1, 1]$  along the fiber cross-section
- We compute attenuations  $A_{\text{spec}}(h, p)$  for each lobe  $p$  assuming no deflection away from the specular cone angles
- We select a lobe in proportion to the specular attenuations
- We now know which  $M_p$  function (of variance  $v_p$ ) to sample using Equation 7, giving  $\theta_o$  and  $\theta_d$ .
- We sample a random Gaussian variable  $g$  and use Equation 8 to compute the relative outgoing azimuth.
- We return a sample weight of  $w = A(p, h) / w_p$

The PDFs of this sampling scheme are exactly analogous to evaluation of the model  $S$ , but with the attenuations  $A$  replaced by the selection weights  $w_p$ , which are easily evaluated for  $R$  and  $TT$  by solving for the single root  $h_p$ . For the more complicated TRT lobe, the quadrature evaluation of  $N_{TRT}$  with the selection weights  $w_p$  in place of  $A$  is required, and this is of comparable cost to evaluation of TRT itself. The total PDF is the sum of these individual lobe pdf terms.

## 4 Lambertian Fibers

We briefly mention two new (to our knowledge) results regarding diffuse fibers. Zinke and Weber [2007] derived the BCSDf for an opaque, cylindrical fiber with Lambertian reflectance properties. They expressed the exact solution as an integral and gave an approximate expansion of it. We note that the exact solution has a simple form (taking care that this holds only for  $\phi \in [0, 2\pi]$ ):

$$S_{\text{Lambert}}(\theta_i, \theta_o, \phi) = k_d \left| \frac{(\pi - \phi) \cos(\phi) + \sin(\phi)}{4\pi} \right| \quad (9)$$

where  $k_d$  is the diffuse albedo. We also note that importance sampling this BCSDf is done easily with three uniform random variables by, again, selecting  $h \in [-1, 1]$  uniformly and then sampling a Lambertian BRDF about normal  $(h, 0, \sqrt{1-h^2})$  (in the coordinate space where we have chosen the fiber tangent to be aligned with the  $y$  axis and the incoming direction to be  $(0, \sin \theta_i, \cos \theta_i)$ ). After sampling an outgoing direction in this coordinate space, a rotation back to world space gives the desired direction, and a sample weight of  $k_d$  is returned. We mention this approach because the CDF for the azimuthal angle seems to have no analytic inversion.

## 5 Results

We have implemented our new importance sampling for the d'Eon et al. [2011] hair model in a path tracer. Russian roulette is used together with next event estimation and multiple importance sampling at each vertex. The geometric term includes a sine between the hair fiber axis and the ray direction instead of the cosine term applied to surfaces. Timings for a single call to our sampling function are comparable to the time of a single BCSDf evaluation (the most expensive calculations are the quadrature for evaluating the TRT function and its PDF).

Figures 1 and 3 show several results under IBL or area light illumination. We easily support a variety of heterogeneous fiber



Figure 3: Various IBL illuminations of hair.




Figure 4: Limiting the number of indirect bounces (from left to right: 0, 1, 2, ..., 29) clearly shows how much the appearance of hair depends on multiple scattering. The concentration of melanin in the hair is varied from top to bottom to simulate a variety of hair types in one dataset.

properties—a challenge for sampling strategies that use precomputation, especially for smooth fibers, which have strongly peaked reflectances and would require very high resolution datasets to store their sampling tables. With our new importance sampling, the images converge at about 1024 samples or less and are fast enough to allow reasonable iteration time for achieving a desired look, even with very low absorption hairs where most of the reflectance is due to high-order scattering (Figure 4). The requirement for such a high sample count is reduced with lower variance illumination, but sunlight illumination of light, shiny hair remains a challenge to render both efficiently and accurately.

## 6 Conclusions

We have presented new methods for importance sampling physically-based fiber reflectance functions. The analytic form of our sampling scheme makes it easy to adopt in a variety of Monte Carlo renderers. We demonstrated the utility of our new methods for rendering physically-based images in a path tracer.

## 7 Acknowledgements

Joe Letteri, Wenzel Jakob, Guillaume Francois, and Luca Fascione provided helpful feedback, Nicholas Wilson modeled the hair curls and Jean-Marie Aubry kindly provided the hair dynamics simulations.

## References

- D'EON, E., FRANCOIS, G., HILL, M., LETTERI, J., AND AUBRY, J. 2011. An energy-conserving hair reflectance model. In *Computer Graphics Forum*, vol. 30, 1181–1187.
- D'EON, E. 2013. [publons.com/p/2867/](https://publons.com/p/2867/). *Publons* 2867.
- HERY, C., AND RAMAMOORTHY, R. 2012. Importance sampling of reflection from hair fibers. *Journal of Computer Graphics Techniques (JCGT)* 1, 1 (June), 1–17.
- JAKOB, W. 2013. *Light transport on path-space manifolds*. PhD thesis, Cornell University.
- KAJIYA, J., AND KAY, T. 1989. Rendering fur with three dimensional textures. In *ACM Siggraph Computer Graphics*, vol. 23, ACM, 271–280.
- MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, 780–791.
- MOON, J. T., WALTER, B., AND MARSCHNER, S. 2008. Efficient multiple scattering in hair using spherical harmonics. In *ACM Transactions on Graphics (TOG)*, vol. 27, ACM, 31.
- OU, J., XIE, F., KRISHNAMACHARI, P., AND PELLACINI, F. 2012. ISHair: importance sampling for hair scattering. In *Computer Graphics Forum*, vol. 31, 1537–1545.
- SADEGHI, I., PRITCHETT, H., JENSEN, H., AND TAMSTORF, R. 2010. An artist friendly hair shading system. *ACM Transactions on Graphics (TOG)* 29, 4, 1–10.
- VEACH, E. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University.
- WARD, K., BERTAILS, F., KIM, T., MARSCHNER, S., CANI, M., AND LIN, M. 2007. A survey on hair modeling: Styling, simulation, and rendering. *IEEE Transactions on Visualization and Computer Graphics*, 213–234.
- ZINKE, A., AND WEBER, A. 2007. Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics*, 342–356.
- ZINKE, A., YUKSEL, C., WEBER, A., AND KEYSER, J. 2008. Dual scattering approximation for fast multiple scattering in hair. *ACM Transactions on Graphics (TOG)* 27, 3, 1–10.