A Fiber Scattering Model with Non-Separable Lobes - Supplemental Report

Eugene d'Eon[1] Steve Marschner[2] Johannes Hanika[1]
[1]Weta Digital [2]Cornell University

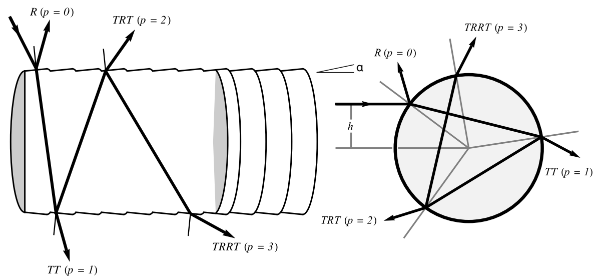

Figure 1: *Refracted pathways through a hair fiber. Single Reflectance (R), double transmittance (TT), transmission-reflection-transmission (TRT) and so on.*

1 Introduction

The rendering of hair, fur, and other fibrous structures (dieletric or conductive) is an important problem in rendering. Sometimes fibers are very close to the camera and span several pixels, requiring an explicit cylindrical geometric representation together with common BRDF and BSDF reflectance models. By far the most common case, however, is very thin fibers that are much more efficiently reprented either as ideal curve primitives or as statistical scattering events in an anisotropic medium. In this latter cases the reflectance from fibers is described using a far-field radiometric model called a *Bidirectional Curve Scattering Distribution Function* (BCSDF) [Zinke and Weber 2007].

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 realistic hair rendering today. The key to factored BCSDFs is to 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 1). (Glossy metallic fibers use only the R lobe). The total reflectance function S is the sum of all such *component scattering functions* (or *lobes*) S_p

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

The modes are indexed by p, the number of internal path segments traversed by light rays contributing to that mode. Following the convention of Marschner et al. [2003], incident and reflected directions are measured in a spherical coordinate system centered on the fiber axis, with θ_i and θ_o measuring *longitudinal* 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:

$$S_p(\theta_i, \theta_o, \phi) = M_p(\theta_i, \theta_o, \phi) N_p(\theta_i, \theta_o, \phi). \tag{2}$$

In previous work, M_p has been assumed independent of ϕ, and N_p depends only weakly on the θs. In this sense the functions S_p are partly separable.

Marschner et al. [2003] began by deriving an exact far-field solution for circular fibers with no tilted scales and a smooth dielectric surface. They then showed how to approximately treat surface roughness, tilted scales, and elliptical cross-sections by modifying the exact solution in various ways. Comparisons of this general approximate fiber model to measured reflectance data from human hair have been performed previously [Marschner et al. 2003; Zinke et al. 2009]. However, to the best of our knowledge, no previous work has studied the radiometric accuracy of these analytic models relative to the ideal representation at their foundation—cylindrical primitives with glossy surfaces. Irrespective of the applicability of these models to simulating human hair reflectance, we find that new important effects are required for analytic BCSDFs to accurately simulate reflectance from cylidrincal primitives with rough surfaces.

We performed a comprehensive Monte Carlo simulation using a fiber model consistent with the original derivation of Marschner et al. [2003]—a rough dielectric cylinder with tilted scales and surface roughness following a Beckmann microfacet distribution (Figure 1). We found the simulation to be in good general agreement with previous analytic models, specifically the treatment of roughness by Zinke and Weber [2007] and d'Eon et al. [2011]. However, we note an important behaviour that was not previously predicted. The shift and blur of the specular cone that is caused by deviations from the smooth cylinder depends significantly on the inclination angle of the incident light and also varies around the cone. That is, the M_p functions depend on all three variables θ_i, θ_o, and ϕ, requiring an entirely *non-separable* model for S_p.

To incorporate these behaviours, we derive higher accuracy perturbations from the exact smooth model, resulting in new M_p functions that bring the model into much better agreement with simulation. Our model is analytic and is easy to deploy in existing rendering systems that use factored BCSDFs. For efficient use in Monte Carlo renderers we also provide practical, analytic, easy to implement importance sampling scheme extending the approach of d'Eon et al. [2013] to handle the non-separable lobes using a short iteration loop.

2 A New Fiber Reflectance Model

In this section we describe details of our ground-truth simulation and our evaluation of the accuracy of previous fiber models. We then describe each lobe of our new model individually and give derivations of our new non-separable functions that exhibit higher accuracy behaviours.

We seek to measure the accuracy of a modified version of the model of Marschner et al. [2003]. Specifically, the parametric hair reflectance model described of the form Equation 1 as a function of *incoming* and *reflected directions*, denoted $\vec{\omega}_i$ and $\vec{\omega}_o$, respectively. Notation is simplified by referring to the longitudinal *difference angle* $\theta_d = (\theta_o - \theta_i)/2$. The relative *index of refraction* η of the hair to the surrounding medium is typically fixed at 1.55. From θ_d, η, and the offset $h \in [-1, 1]$ to the fiber cross section (Figure 1) a Bravais analysis gives the azimuthal distributions compactly using $\gamma_i = \arcsin(h)$, $\gamma_t = \arcsin(\frac{h}{\eta'})$, $\eta' = \frac{\sqrt{\eta^2 - \sin(\theta_d)}}{\cos(\theta_d)}$ to predict the relative change in azimuth

$$\Phi(p, h) = 2\, p\, \gamma_t - 2\, \gamma_i + p\pi \tag{3}$$

for each mode p being considered. To treat rough fibers we use the azimuthal functions N_p of d'Eon et al. [2011] evaluated using a

70-point Gaussian quadrature

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

We use the attenutation terms $A(p,h)$ and the wrapped Gaussian D_p of variance $v_{p,N}$ given by d'Eon et al. [d'Eon et al. 2011]. We use the longitudinal scattering function $M(v, \theta_{\text{cone}}, \theta_o)$ of d'Eon et al. [2011] for all lobes

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

For additional numerical recipes for evaluating this function, see [d'Eon 2013]. We remark that this M function leads to a non-reciprocal reflectance. However, for low fiber roughnesses this does not cause significant problems, and the derivation of non-separable lobe widths and inclinations can apply readily to alternative M functions, such as renormalized Gaussians.

The M function provides rough spreading in the longitudinal directions acting like a Gaussian of variance v_p centered at the non-rough *specular cone exitance* θ_{cone_p}. Note that with previous models, h (and therefore ϕ) does not appear in the evaluation of M and so the lobes are separable. Our new non-separable model arises by adding h variation to both the longitudinal variances v_p and to θ_{cone_p}.

2.1 Ground Truth Simulation

We ran a geometrical optics Monte Carlo simulation of scattering within rough dielectric cylinders with homgoeneous interior absorption. We performed a full suite of simulation experiments with parameters

$\theta_i \in [0, 0.25, 0.5, 0.75, 1, 1.1, 1.2, 1.3, 1.4, 1.5, 1.55]$

$\alpha \in [0, \pm 2\,\text{degrees}, \pm 4\text{degrees}]$

$\mu_a \in [0.0, 0.003, 0.006, 0.015, 0.03, 0.06, 0.12, 0.25, 0.5, 1.0, 2.0, 4.0,\ 8.0, 16.0]$

$\beta \in [0.0025, 0.005, 0.01, 0.02, 0.03, 0.04, 0.06, 0.08, 0.15, 0.3]$

$\eta \in [1.55, 1.55/1.33]$

where α is the tilt of the surface scales, μ_a is the internal absorption coefficient, β is the Beckmann roughness, and η is the relative index of refraction of the hair to the surrounding medium (we considered human hair in air and in water). For each simulation the trace geometry was a unit radius analytic infinite cylinder with a rough dielectric BSDF ([Stam 2001; Walter et al. 2007]) with Beckmann roughness β and Smith shadowing. For each simulation with fixed θ_i a random offset $h \in [-1, 1]$ was chosen uniformly (see Figure 1). Importance sampling of the BSDF determined the traced pathways (ie. R vs T was chosen proportional to each Fresnel term—grazing the edge of the fiber made the likelihood of scoring an "R" sample higher, for example). 10 million rays were traced per simulation. The outgoing distribution was accumulated into a 400 x 100 latlong image. The sample weight was the product of weights returned by the surface BSDF importance sampling and also multiplied by the Beer-Lambert absorptions along the internal pathways. The first four lobes were stored in separate buffers. All paths past TRRT were stored in a single buffer. Tilted scales were simulated with a shading normal (so rare occurences of hitting the scales from the back side result in energy loss). Some energy loss is also to be expected from the shadowing terms in the BSDF especially for high roughness.

2.2 Approximating cone angles and widths

For fibers with rough surfaces and tilted scales, the ground truth simulation shows significant variation in both the width and the inclination of the R, TT, and TRT highlights as a function of ϕ, α, and θ_i. Previous fiber models have assumed constant lobe widths $\beta_R, \beta_{TT}, \beta_{TRT}$ that are user-specified. In this section we show how to connect these explicitly to ϕ, α, θ_i, and roughness β and show that these new expressions create lobes in much closer agreement with simulation. While the exact values can always be computed using ray tracing through the cylinder, simple approximations to these inclinations and widths are of great help in formulating a practical scattering model.

We call the surface normal of the average cylinder the *macro surface normal* or macro-normal and the surface normal that is perturbed by scale tilt and roughness the *micro surface normal* or micro-normal. For this section of the paper we are concerned only about *longitudinal* perturbations to the path, so we need only consider longitudinal perturbations to the normals, and we will assume the micro-normal, macro-normal, and the fiber axis are coplanar. Under this assumption the micro-normal at the k^{th} intersection point can be described by its inclination α_k relative to the macro-normal. Each α_k is an independent random variable, with mean α and variance β^2.

The path taken by a ray depends on the micro-normals encountered at each of the refraction or reflection points (see Figure 1); a single ray entering at inclination θ_i and offset h gives rise to an exiting ray for each p that departs the cylinder at an inclination $\theta_p(\theta_i, h, \alpha_1, \ldots, \alpha_{p+1})$. As observed by Marschner et al. [2003], if the normals are all perpendicular to the cylinder axis ($\alpha_k = 0$), then the ray exiting the cylinder will have exactly the opposite inclination as the incident ray; but in the presence of tilted normals, the inclination of the exiting ray varies. Our goal is to calculate the derivatives of θ_p with respect to the α_ks, from which we will ultimately derive the longitudinal center and width of the lobe (that is, the mean and variance of θ_p, taken over the random distribution of α_ks).

To compute the angle of the exiting ray, we express the ray directions at each interface as unit vectors in a basis aligned with the macro- or micro-normal, keeping only the components perpendicular to the normal. We call the component in the direction of the fiber axis x, and the component in the perpendicular direction y. The disk $x^2 + y^2 < 1$ is the Nusselt analog of the hemisphere about the surface normal.

In this coordinate system, the coordinates of the incident direction are $(x_i, y_i) = (\sin\theta_i, h\cos\theta_i)$. The progress of the ray can then be computed using a series of transformations in (x, y) space, the most important being rotations about the y axis

$$Q(\alpha_k, x, y) = (x\cos\alpha_k - z\sin\alpha_k, y),$$

where $z = \sqrt{1 - x^2 - y^2}$, which are used to convert between the coordinates of the micro- and macro-normal. For instance, the effect of a reflection on a ray coming from the outside of the fiber is calculated by starting with the (outward-pointing) incident direction (x, y) relative to the macro-normal, rotating it to align with the micro-normal, negating it (for reflection), and rotating back to the frame of the macro-normal:

$$R[\alpha_k](x, y) = Q(\alpha_k, -Q(-\alpha_k, x, y)).$$

For a ray from the inside, the rotations are in the opposite sense:

$$R'[\alpha_k](x, y) = Q(-\alpha_k, -Q(\alpha_k, x, y)).$$

In these Nusselt coordinates, the effect of refraction going into the fiber is simply a dilation:

$$T[\alpha_k](x, y) = Q(-\alpha_k, -Q(-\alpha_k, x, y)/\eta)$$

and going out of the fiber:

$$T'[\alpha_k](x, y) = Q(\alpha_k, -Q(\alpha_k, x, y)\,\eta).$$

Again the sign of the rotation reverses for vectors on the inside of the fiber. Finally, the relationship between the transmitted direction at one point on the surface and the incident direction after propagation to the next point is simply

$$P(x,y) = -(x,y).$$

With these tools we can express the functions that map incident to exiting directions for the three modes as follows:

$$(x_R, -y) = F_R(\alpha_1, x, y) = R[\alpha_1](x_i, y)$$

$$(x_{TT}, -y) = F_{TT}(\alpha_1, \alpha_2, x, y) = (T'[\alpha_2] \circ P \circ T[\alpha_1])\,(x_i, y)$$

$$(x_{TRT}, -y) = F_{TRT}(\alpha_1, \alpha_2, \alpha_3, x, y) = (T'[\alpha_3] \circ P \circ R'[\alpha_2] \circ P \circ T[\alpha_1])(x_i, y).$$

Note that $\eta|y|$ is invariant along the whole path, since we are only considering normal tilt in the x direction.

With the computation set up this way, with all functions written as compositions of Q with scaling interspersed, the derivatives can be computed by multiplying the derivative matrices (this procedure is similar to the approach taken by automatic differentiation). The derivative of Q is

$$\begin{bmatrix} -z\cos\alpha_k - x\sin\alpha_k & \cos\alpha_k + (x/z)\sin\alpha_k & (y/z)\sin\alpha_k \\ 0 & 0 & 1 \end{bmatrix}$$

and the derivatives of the Fs can be easily computed by evaluating this derivative for the various events along the path and multiplying.

To find the lobe tilt angles for the specular cone θ_{cone}, which is to say, the deviation of a ray that hits average normals at all interfaces, we evaluate x_R, x_{TT}, and x_{TRT} for $\alpha_1 = \alpha_2 = \alpha_3 = \alpha$ and solve for $\theta = \arcsin x$. To find the lobe width, we evaluate the derivatives $dx_*/d\alpha_k$ at $\alpha_k = \alpha$ and compute β times the norm of the resulting vector of derivatives.

While the above provides an efficient way to calculate the three Fs and their derivatives $dF_*/d\alpha_k$, we use a simpler approximation, by expanding the Fs about zero tilt:

$$F_*(\alpha, x_i, y) = F_*(0, x_i, y) + \alpha(D_\alpha F_*)(0, x_i, y).$$

For all three lobes $F_*(0, x_i, y) = (-x_i, -y)$. The relevant derivatives of F take on a simple form at $\alpha = 0$:

$$\frac{dx_R}{d\alpha_1} = -2z$$

$$\frac{dx_{TT}}{d\alpha_{1,2}} = [\eta z' - z, \eta z' - z]$$

$$\frac{dx_{TRT}}{d\alpha_{1\ldots3}} = [\eta z' - z, 2\eta z', \eta z' - z]$$

where $z' = \sqrt{1 - (x/\eta)^2 - (y/\eta)^2}$. Using the approximate model, the lobe shifts are the sums of these vectors times α, and the widths are the norms of these vectors times β.

2.3 A New R Lobe

Using the formalism derived in the previous section we derive a new R reflectance model ($p = 0$) for both conductive and dielectric fibers. Our motivation is to include two behaviours seen in ground truth Monte Carlo photon simulation that previous models lack:

- the longitudinal angle θ of the specular cone varies with the relative azimuth of the light and camera, ϕ
- the longitudinal spread of the exitance also varies with ϕ—the R function tightens when the light grazes off the sides of the fiber

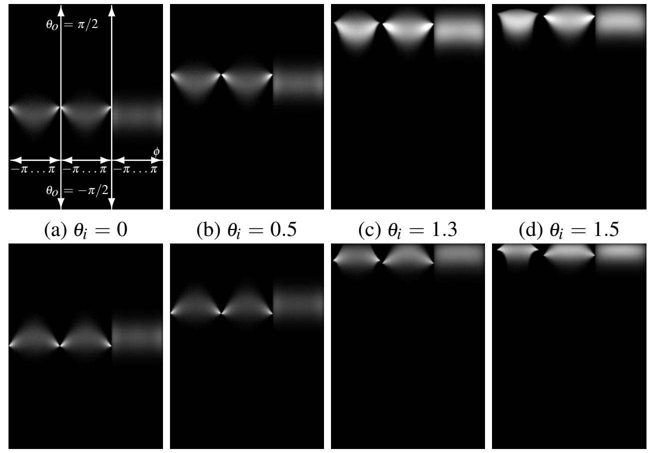

(a) $\theta_i = 0$ (b) $\theta_i = 0.5$ (c) $\theta_i = 1.3$ (d) $\theta_i = 1.5$

(e) $\theta_i = 0$ (f) $\theta_i = 0.5$ (g) $\theta_i = 1.3$ (h) $\theta_i = 1.5$

Figure 2: *Each plot shows three lat-long visualizations of the exitance from the fibre due to R reflection: MC ground-truth (left), our new proposed R lobe (middle) and [d'Eon et al. 2011] (right). Top row: Scale tilt of 4 degrees, bottom row: Scale tilt of -4 degrees. Fiber roughness of* 0.08.

From the analysis in the previous section we find that specular cone exitance for R is

$$\theta_{\mathrm{cone}_R} = \arcsin x_R =$$
$$-\arcsin\left(\sin\theta_i - 2\sin\alpha\left(\cos(\phi/2)\cos\alpha\cos\theta_i + \sin\alpha\sin\theta_i\right)\right)$$

which shows the ϕ and α dependence clearly and reduces to the expected $-\theta_i$ when $\alpha = 0$. To get a first order approximation of the width $\beta_R = \sqrt{v_R}$ of the roughened M function as a function of h and θ_i we use the derivative of the F_R function above to arrive at the simple expression:

$$\sqrt{v_R} = \beta\sqrt{2(1-h^2)} = \beta\sqrt{2}\cos(\phi/2) \qquad (6)$$

Like d'Eon et al. [2011] we special case N_R to not use azimuthal roughness. We found this new R lobe to match simulations quite closely for all but extremely grazing incidence/exitance (see Figure 2).

2.4 A Non-Separable TT Lobe

Monte Carlo simulation of double transmission (TT) through rough dielectric cylinders showed several behaviours not found in previous hair models:

- The width of the M_{TT} function depends on ϕ
- The width of the M_{TT} function widens dramatically as the longitudinal difference angle, θ_d becomes large
- The tilted scales lead to non-trivial lobe placement and total internal reflection can happen for a significant portion of the incident angles

Similar to R, our new TT lobe includes h dependence in both the cone angles and longitudinal widths. These are easily evaluated using several Nusselt Sphere rotations

$$\theta_{\mathrm{cone}_{TT}} = \arcsin\left(x_{TT}(\theta_i, \alpha, h)\right). \qquad (7)$$

To relate h to ϕ we use the solution of Equation 3 for ($p = 1$) [d'Eon et al. 2011]. However, unlike R, x_{TT} may not be real or we could also have $|x_{TT}| > 1$. This corresponds to total internal reflection as light attempts to transmit out of the fiber on the second T. This effect is seen clearly in ground-truth simulation. For tilt angles as small as 2° no TT energy escapes a smooth fiber for $\theta_i < -1.1$.

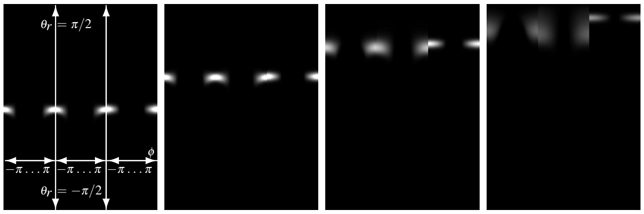

(a) $\theta_i = 0$ (b) $\theta_i = 0.5$ (c) $\theta_i = 1$ (d) $\theta_i = 1.4$

Figure 3: *Four TT lobe comparisons. Each image: MC (left), our new TT lobe (middle) and [d'Eon et al. 2011] (right). Note the variation in position and width of the lobes with ϕ (horizontal dir) and also the widening of TT as θ_i becomes large. These behaviours are not possible with a separable BCSDF.* $\eta = 1.55, \alpha = -2°, \beta = 0.04$.

To handle these cases we clamp the inputs to all square root evaluations in x_{TT} to be non-negative and for $|x_{TT}| > 1$ we return $\theta_{\text{cone}_{TT}} = -\text{sign}(\theta_i)\pi/2$ instead of returning 0 for M. This allows very rough fibers to scatter small amounts of energy into these black regions instead of creating unrealistic sharp discontinuities in the lobes.

Our new non-separable TT lobe width $\sqrt{v_{TT}}$ is estimated at $\alpha = 0$ to simplify the result. We found this to work well when applied directly to lobes shifted by $\alpha < 5°$ (Figure 3). We summarize the result, where the result uses θ_d in place of θ_i for reciprocity,

$$
\begin{aligned}
x &= \sin(\theta_d) \\
y &= h\cos(\theta_d) \\
D_{TT} &= \frac{-2\sqrt{1-x^2-y^2} - \eta\sqrt{-((-\eta^2+x^2+y^2)/\eta^2))}}{\cos(\theta_d)} \\
v_{TT} &= (\frac{1}{2}\beta D_{TT})^2.
\end{aligned}
$$

In theory a similar analysis could be used with derivatives in the azimuthal plane to estimate the azimuthal widths of the blurring in Equation 4 but we found a constant $v_{1,N} = \beta^2/2$ sufficiently accurate.

2.5 A New TRT Lobe

Similar to TT, we derive a more accurate TRT specular cone angle $\theta_{\text{cone}_{TRT}}$ to predict better longitudinal lobe positioning. This is important visually because the offset from R creates the characteristic asymmetric R + TRT highlight seen in high-variance illuminations of human hair (Figure 5). We could compute the exact lobe inclination and width as a function of h, but solving Equation 3 for h given ϕ is problematic and has up to three roots. However, the variation in the TRT longitudinal position with respect to h and therefore $\phi(h)$ was rather limited. Furthermore, when TRT is bright, its energy is mostly placed in a narrow range of ϕ, but due to a wide range of h, so we use a lobe inclination and width computed at $h = 0$ (and thus, our TRT lobe remains separable). However, it still predicts more accurate lobe placement and width than previous models (Figure 4). Our new cone angle is

$$\theta_{\text{cone}_{TRT}} = \arcsin x_{TRT}(\alpha, \sin(\theta_i), \cos(\theta_i)) \tag{8}$$

and its evaluation is most compact by computing it as the sequence of rotations derived in Section 2.2. Our new TRT lobe width $\sqrt{v_{TRT}}$ is also estimated once at $h = 0$ and $\alpha = 0$. The final result after all rotations is

$$
\begin{aligned}
D_{TRT} &= (-2\cos\theta_d + 4\eta\sqrt{1-(\sin\theta_d)^2/\eta^2})/\cos\theta_d \\
v_{TRT} &= (1/4)\beta^2 D_{TRT}^2.
\end{aligned}
$$

For the azimuthal term N_2 we use azimuthal variance $v_{2,N} = 2\beta^2$.

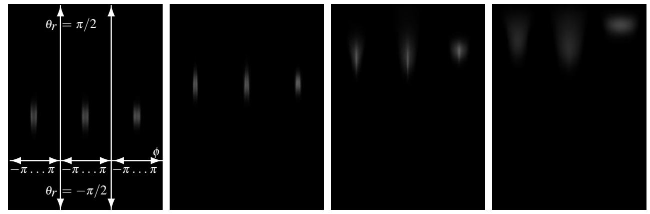

(a) $\theta_i = 0$ (b) $\theta_i = 0.5$ (c) $\theta_i = 1$ (d) $\theta_i = 1.4$

Figure 4: *Four TRT lobe comparisons. Each image: MC (left), our new TRT lobe (middle) and [d'Eon et al. 2011] (right). Note the widening of TRT as θ_i becomes large and more accurate lobe placement.* $\eta = 1.55, \alpha = -2°, \beta = 0.04$.

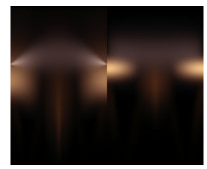

Figure 5: *The fine-tuning of lobe shapes might seem trivial but their alignment, width, and brightness contribute to the color and saturation of the total light reflected in any given direction and our new model (left) can predict quite different results from previous separable models (right).* $\beta = 0.06, \eta = 1.55, \theta_i = 1.1, \alpha = -2°, \mu_a/r = \{0.12, 0.25, 0.5\}$.

3 Importance Sampling Non-Separable Fiber Lobes

Here we describe an approximate iterative scheme to sample the non-separable lobe functions of our new model. We begin with the importance sampling approach of d'Eon et al. [2013]. Having selected h and p, the we need to combine the two individual M and N sampling schemes into an importance sampling of their non-separable product $M_p \times N_p$ (giving azimuthal deflection ϕ and longitudinal exitance θ_r and a cumulative sample weight). This is challenging because sampling ϕ changes v_p, the width of the M_p lobe. Further, the longitudinal deflection due to sampling M_p changes the Bravais index η' and, thus, the azimuthal function N_p. So, despite having sampling schemes for M and N individually, we don't know how to sample either of them because knowing how to sample M requires the result of sampling N, and vice versa. An approach we found to work well is to note that a fixed point exists: the shapes of M and N are such that, given sampling variables ξ_h and g for sampling ϕ, and $\xi_{M,1}$ and $\xi_{M,2}$ for sampling θ_r, there exists $\{\phi, \theta_r\}$ such that sampling ϕ using θ_r gives a ϕ that samples θ_r. In fact, we found that the two sampling schemes themselves are an efficient iteration procedure for finding this fixed point. So our sampling scheme is to choose all random numbers, start with a guess for ϕ, θ_r and sample one from the other back and forth until the fixed point is found and return this as our sampled direction. We describe TT here as an example, but all the lobes follow this simple pattern. We initialize the process by guessing $\theta_r = -\theta_i$—the specular cone angle for no scale tilt, resulting in the following algorithm:

```
{theta_r,phi,w} = sampleTT(h,g,xi_M1,xi_M2)
{
    // start along the specular cone
    phioffset = sqrt( v_TTN ) * g
    theta_r = -theta_i
    for( NUM_ITERATIONS times )
    {
        theta_d = 0.5 * ( theta_r - theta_i )
```

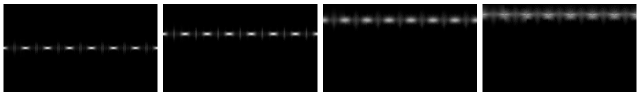

(a) $\theta_i = 0$ (b) $\theta_i = 0.5$ (c) $\theta_i = 1.0$ (d) $\theta_i = 1.2$

Figure 6: *Sampled distributions of the complete non-separable fiber model as a function of the number of iterations in our sampling scheme. Left: Eval comparison, then, moving right, 1, 2, 3, 4, 5 and 10 iterations.* $\beta = 0.04, \alpha = 0$. *10000 samples each. (b) On the first iteration the TRT caustics form vertical lines in the middle, but on the second iteration their variation with* θ *is recovered. Only for grazing angles (d) with TT are more than 2 iterations required to arrive at the desired distribution.*

```
        phi = Phi_1( h, eta_prime(h,eta,theta_d) ) + phioffset
        theta_cone = thetaConeTT( theta_i, alpha, phi, theta_d )
        v = v_TT( phi, beta, theta_d )
        theta_r = sampleM( xi_M1, xi_M2, v_TT );
    }
    return { theta_r,phi,A(p,h) / w_p }
}
```

It then remains to determine what function this samples and how well the iteration scheme converges. In short, the sampled function is the fiber function plus a very small residual error and we require, on average, only 2 iterations for most roughnesses and angles, and no more than 8 iterations worst case to sample realistic hair scattering (Figure 6). In the Appendices we discuss more formal details of the convergence and residual.

4 Application to Hair Rendering

We were initially hopeful that the newly derived non-separable fiber model would do away with the somewhat arbitrary parametric complexity of previous hair fiber models by reducing the degrees of freedom of the model to four physically-based parameters (per wavelength): α, η, β and μ_a. As Figure 7 shows, the analytic model is indeed very accurate at describing the far-field reflectance of a rough dielectric cylinder with internal absorption and tilted scales. However, when applying the model in a Monte Carlo renderer to match controlled photographs of real hair, we found that it was necessary to longitudinally widen the TT and TRT lobes significantly relative to R in a way not predicted by this physical model. The longitudinal widening would not be explained by the known elliptical cross section of hair fibers [Khungurn and Marschner 2014] and we expect some form of interaction between overlapping scales, or internal volumetric scattering by air pockets or other heterogeneity is causing this effect. The R lobe of our new model, however, shows immediate practical use for rendering hair and exhibits a fundamental property of glossy fiber reflectance that is to be expected in future hair models which consider new underlying physical representations, and consider both elliptical and non-separable effects in one combined framework.

A Importance Sampling Scheme Residual

In this appendix we derive the probability density of the samples generated by our iteration. We give the proof for a slightly simpler scenario in which a single random number is used to choose each of ϕ and θ by inverting a cumulative distribution function. That is, we have inverse-CDF sampling procedures to choose ϕ from the azimuthal lobe if we know θ, and to choose θ from the longitudinal lobe if we know ϕ. This means we have a function $P_M(\theta,\phi)$ that is the CDF of $M(\theta,\phi)$ as a function of θ for any fixed value of ϕ, and

(a) previous (b) our model (c) previous

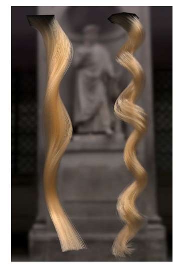

(d) our model

Figure 8: *Rendering hair fibers with our new model does not always produce dramatic appearance differences relative to previous models (here we show forward path tracing with NEE). However, in practice we find that incorporating the new physically-based behaviours lead to a more consistent appearance over varied lighting conditions.*

also $P_N(\theta,\phi)$ which is the CDF of $N(\theta,\phi)$[1] as a function of ϕ for any fixed value of θ. The two individual sampling procedures are:

$$\text{solve } P_M(\theta,\phi) = \xi_1 \text{ for } \theta \qquad (9)$$

$$\text{solve } P_N(\theta,\phi) = \xi_2 \text{ for } \phi \qquad (10)$$

where ξ_1 and ξ_2 are independent and uniformly distributed in $[0,1]$. Used separately, these procedures will generate random samples with densities $M(\cdot,\phi)$ and $N(\theta,\cdot)$ respectively. The sampling procedure we are using to sample the product of these densities is iteratively solving the joint problem:

$$\text{solve} \begin{bmatrix} P_M \\ P_N \end{bmatrix} (\theta,\phi) = \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix} \text{ for } \theta \text{ and } \phi \qquad (11)$$

The two CDFs are considered together as a function $F : \boldsymbol{R}^2 \to \boldsymbol{R}^2$, which is being solved to find θ and ϕ in one step. Since the 2D random variable (ξ_1,ξ_2) has a uniform, unit probability density, the density of (θ,ϕ) is simply the determinant of the derivative (Jacobian) of F. This derivative is

$$\nabla F = \begin{bmatrix} \frac{\partial P_M}{\partial \theta} & \frac{\partial P_M}{\partial \phi} \\ \frac{\partial P_N}{\partial \theta} & \frac{\partial P_N}{\partial \phi} \end{bmatrix} = \begin{bmatrix} M(\theta,\phi) & \frac{\partial P_M}{\partial \phi} \\ \frac{\partial P_N}{\partial \theta} & N(\theta,\phi) \end{bmatrix}$$

If at least one of the off-diagonal terms is zero, we have the desired pdf $M(\theta,\phi)N(\theta,\phi)$. However, if both off-diagonals are non-zero (as they generally are in the non-separable case), we have an extra term in our pdf, which we call the residual:

$$p(\theta,\phi) = M(\theta,\phi)N(\theta,\phi) - \frac{\partial P_M}{\partial \phi}\frac{\partial P_N}{\partial \theta}$$

The only case in which the residual term would be large is when an M lobe changes width or position quickly as a function of ϕ while N simultaneously changes rapidly with θ_d. The only lobes for which N changes rapidly with θ_d are the $\mathrm{TR}^{+}\mathrm{T}$ terms. Since we use separable lobes for these, their residual is zero.

[1] We are ignoring attenuation for simplicity, so that N integrates to 1.

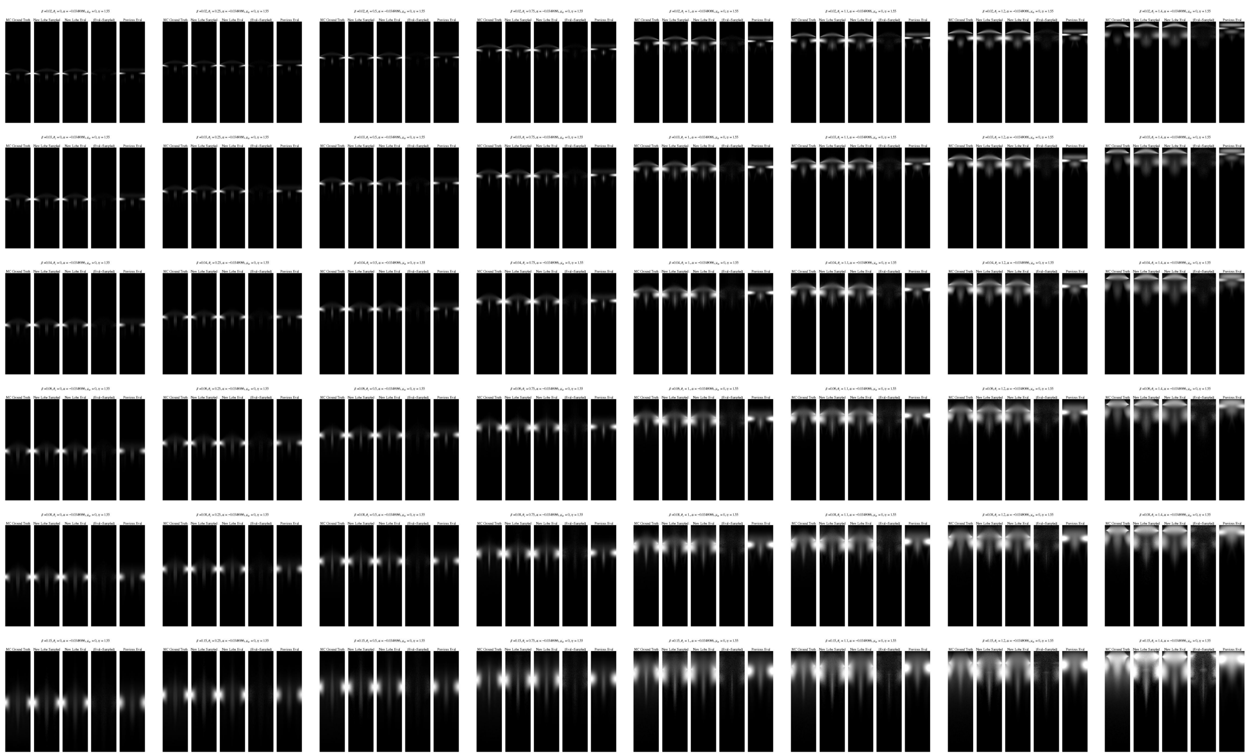

Figure 7: *Extensive Monte Carlo validation of our model and the importance sampling procedure for it. We find good matching between the sampled function (using a fixed iteration count of 6) to our new lobe functions for all but extreme roughness and inclination θ_i. Each result contains, from left to right, (a) a lobe image from the ground truth MC simulation with 10 million photons, (b) the result of sampling our new model (including lobe selection) 1 million times, splatting with the lobe weights, (c) the result of evaluating our new model over the lat/long image, (d) an absolute difference image of the same intensity of the other images (chosen by finding the brightest pixel in the left image to then map to 1), (e) evaluation of a semi-separable model ([d'Eon et al. 2011] specifically).*

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., MARSCHNER, S., AND HANIKA, J. 2013. Importance sampling for physically-based hair fiber models. In *SIGGRAPH Asia 2013 Technical Briefs*, ACM, 25.

D'EON, E. 2013. `publons.com/p/2867/`. *Publons 2867*.

KAJIYA, J., AND KAY, T. 1989. Rendering fur with three dimensional textures. In *ACM Siggraph Computer Graphics*, vol. 23, ACM, 271–280.

KHUNGURN, P., AND MARSCHNER, S. 2014. On azimuthal scattering behavior in hair fibers. *Accepted to TOG*.

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.

STAM, J. 2001. An illumination model for a skin layer bounded by rough surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, Springer-Verlag, 39–52.

WALTER, B., MARSCHNER, S., LI, H., AND TORRANCE, K. 2007. Microfacet models for refraction through rough surfaces. In *Rendering Techniques (Proc. EG Symposium on Rendering)*, 195–206.

ZINKE, A., AND WEBER, A. 2007. Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics*, 342–356.

ZINKE, A., RUMP, M., LAY, T., WEBER, A., ANDRIYENKO, A., AND KLEIN, R. 2009. A practical approach for photometric acquisition of hair color. In *ACM SIGGRAPH Asia 2009 papers*, ACM, 1–9.