

A Practical and Controllable Hair and Fur Model for Production Path Tracing

Matt Jen-Yuan Chiang¹ Benedikt Bitterli^{1,2} Chuck Tappan¹ Brent Burley¹

¹Walt Disney Animation Studios

²Disney Research Zürich

Abstract

We present an energy-conserving fiber shading model for hair and fur that is efficient enough for path tracing. Our model adopts a near-field formulation to avoid the expensive integral across the fiber, accounts for all high order internal reflection events with a single lobe, and proposes a novel, closed-form distribution for azimuthal roughness based on the logistic distribution. Additionally, we derive, through simulation, a parameterization that relates intuitive user controls such as multiple-scattering albedo and isotropic cylinder roughness to the underlying physical parameters.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism—Shading

1. Introduction

Hair and fur are ubiquitous in virtual worlds created for film. However, even though the first hair shading models appeared more than two decades ago, efficient and accurate rendering of hair and fur remains a challenging problem. The geometric complexity of fiber assemblies and the intricacy of multiply scattered light make hair rendering computationally demanding and difficult to control.

The ever-increasing availability of computational resources has made production rendering practical to move away from *ad hoc* methods towards physically based ones. Path tracing in particular has found wide-spread use in production rendering in recent years, both due to its simplicity and fidelity. However, using Monte Carlo methods to simulate actual light scattering within fiber assemblies creates new challenges for hair and fur rendering, which we address in this paper.

Hair shading models have undergone a development similar to general rendering methods, moving from *ad hoc* fiber reflectance models [KK89] towards physically motivated models [MJC*03, dFH*11] based on dielectric cylinders. However, most existing fiber reflectance models were designed within the context of scan-line or simple ray tracing renderers, which shade on the order of once per pixel. In contrast, Monte Carlo rendering methods and path tracing in particular produce many (on the order of thousands) samples for each pixel in order to realize global illumination by tracing paths against complicated three-dimensional production scenes. Compounded by the need of scattering through often an enormous number of strands (on the order of millions per character), each path might end up with hundreds of hair shader evalua-

tions. Hence individual shader evaluation need to be computationally cheap in order for such an approach to be practical.

In addition, most previous hair shading models are parameterized using the physical properties of the material, rather than parameters that directly relate to its visual appearance. Such models tend to be unintuitive to control for artists, making it difficult to quickly obtain a desired look. Additionally, multiply scattered light within the hair creates an unintuitive, non-linear relationship between parameters of a single fiber shading model and the appearance of a fiber assembly.

In this paper, we address these problems and present the following contributions:

- A physically based single fiber scattering model that allows for efficient Monte Carlo rendering of path-traced multiple fiber scattering. This entails moving from the commonly used scattering integral for far-field shading models to a near-field one. Additionally, we use a single lobe to approximate the scattering behavior of an infinite number of high-order lobes, resulting in affordable perfect energy conservation even for non-absorbing fibers. Our model is easy to implement and achieves renderings near indistinguishable from previous work at a fraction of the render time in production examples.
- A reparameterization of the absorption coefficient and roughness parameters that is more intuitive for artists and enables efficient artist workflow, while remaining physically consistent. We also provide additional controls to achieve wider range of appearances for animal fur.

The rest of the paper proceeds as follows. In Section 3 we present




Figure 1: Path-traced multiple scattering (top) vs. its Dual Scattering approximation [ZYWK08] (bottom). All the fur strands have the same absorption coefficient and are illuminated by a single one-degree distance light source so that the observed proper light wrapping (top left), soft shadow (top middle) and the milky quality of the fur (top right) are predominantly produced by faithfully simulating light bouncing around in a volume of fur. On the other hand, the Dual Scattering approximation (and other shadow-based approaches) often produces a coarse and stiff look (bottom). In this example the issue is amplified by the low fiber absorption and the unorganized orientations.

our efficient single fiber scattering model. In Section 4 we show how we parameterize our model for intuitive artist control. Finally, we analyze the efficiency of our model on several production examples in Section 5.

2. Related Work

Single Fiber Scattering Kajiya and Kay [KK89] presented the first physically based fiber reflectance model, based on a combination of an anisotropic specular and a diffuse component. Marschner et al. [MJC*03] significantly improved upon this work and modeled hair fibers as absorbing, rough dielectric cylinders. They split the scattering distribution of the fiber into different lobes based on the number of internal reflections within the cylinder and consider the first three lobes (R , TT , TRT) for their shading model. As an additional simplification, each lobe is modeled as a product of two 1D profiles, the azimuthal and the longitudinal scattering function. d'Eon et al. [dFH*11] further improved upon this work by introducing an energy conserving longitudinal scattering function as well as an azimuthal scattering function that incorporates azimuthal surface roughness. Their model also extends beyond the TRT lobe to arbitrary internal path lengths.

All of these models are considered *Bidirectional Curve Scattering Distribution Functions* (BCSDF) [ZW07] employing a *far-field* perspective, where the width of the fiber is assumed to be insignificant compared to its length and the scattering response is averaged over the width of the fiber. At low sample counts, such a far-field model is essential to avoid aliasing and noise. However, it also adds significant complexity to the shading model: Averaging the shading response over the width of the fiber involves an integral, which is non-trivial due to the rough refraction on cylindrical geometry. Marschner et al. [MJC*03] perform this integration analytically, but they only present solutions for the first three scattering events and cannot incorporate azimuthal roughness. d'Eon et al. [dFH*11] sidestep this issue by performing integration numerically using Gaussian quadrature; this allows them to incorporate azimuthal roughness and an arbitrary number of lobes, but the

quadrature comes at a significant computational cost and suffers from oscillation at low roughnesses.

The concurrent work by Yan et al. [YTJR15] modeled the unique scattering behavior of animal fur by introducing an additional internal cylinder to account for *medulla*, which is the inner core more pronounced in animal fur than in human hair. They validated such double cylinder model with measurements on real fur fibers. However, the additional complexity by adding the medulla scattering dictates a tabulated implementation that makes spatially varying parameters infeasible. Also its lack of intuitive parameterization makes producing a predictable result challenging. It is worth noting that even though their model incorporates *near-field* effects like ours, the efficiency gain was not the motivation hence never evaluated.

Multiple Fiber Scattering Unbiased path-traced multiple fiber scattering using previous single-scattering models is prohibitively expensive for production rendering. Hence its approximation has been an active research topic and remains necessary for efficiency reasons. Lokovic and Veach [LV00] precomputed the transmittance along shadow rays through hair, but can only account for unscattered forward transmission. Similarly Zinke et al. [ZYWK08] presented Dual Scattering, approximating the scattering of such forward transmission by roughening the angular distribution. They also introduced an ad hoc local scattering term assuming high local orientation similarity. Moon and Marschner [MM06] presented a photon mapping based approach to compute multiple scattering within hair. Their method was later improved by Moon et al. [MWM08] by performing light tracing through a voxelized hair volume, which stores filtered scattering functions encoded in spherical harmonics. In general, our work is orthogonal to these approaches: Our reflectance model concerns single fiber scattering only and the main goal of our efficiency improvements is to actually make path-traced multiple scattering feasible in production, therefore avoiding such biased or inconsistent approximations.

Art Controllability Sadeghi et al. [SPJT10] presented an artist friendly hair shading model explicitly based on the single-

scattering model of Marschner et al. and the multiple scattering approximation of Dual Scattering [ZYWK08]. They consider the same problem of artist controllability as us, however, their method concerns little about the physical constraint of the underlying models and explicitly relies on Dual Scattering to obtain intuitive controls in the presence of multiple scattering in hair. In contrast, our approach of intuitive parameterization is independent of the multiple-scattering method used and retains the physically based properties of the single-scattering shading model.

3. Our Model

In this section, we present our single fiber reflectance model, which is based on the model of d'Eon et al. [dFH*11]. Our model is designed to provide affordable azimuthal roughness practical enough for path-traced multiple fiber scattering in production and also achieve perfect energy conservation for previously difficult cases such as high-roughness or non-absorbing fibers.

3.1. Background

Notation Our model is defined as a reflectance function over a pair of incident and exitant directions. Following the convention [MJC*03], the directions are measured in a spherical coordinate system aligned with the fiber axis, with θ measuring their *longitudinal inclination* and ϕ measuring their azimuth. We only consider fibers with radial symmetry, and hence our model is fully parameterized by the incident inclination θ_i , the exitant inclination θ_o and the *relative azimuth* $\phi = \phi_o - \phi_i$.

Factored Model We follow the *factored lobe* representation [MJC*03], which decomposes the fiber reflectance $S(\theta_i, \theta_o, \phi)$ into a sum of separate modes of propagation as *lobes*,

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

Here, $S_p(\theta_i, \theta_o, \phi)$ represents the p -th lobe, which is the contribution of all light paths with p internal reflections within the fiber cylinder before exiting and $p \in \{R=0, TT=1, TRT=2, TRRT=3, \dots\}$. Each lobe is additionally factored into a product of two 1D functions, the *longitudinal scattering function* M_p and the *azimuthal scattering function* N_p ,

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

Longitudinal Scattering d'Eon et al. [dFH*11] proposed an energy-conserving longitudinal scattering function, which we found it to be in good agreement with the ground truth Monte Carlo simulation of light scattering in a rough dielectric cylinder. We also adopt the perfect importance sampling scheme proposed by d'Eon et al. [dMH13]. See Section 6 for more discussion.

Azimuthal Scattering The azimuthal scattering behavior is different from the longitudinal behavior in that it is highly varying across the width of the fiber. Assuming distant light sources and observers, the azimuthal scattering function proposed by d'Eon et al. [dFH*11] contains a uniform integration over the fiber width to capture such complex scattering behavior as an aggregated *far-field*

response,

$$N_p(\phi) = \frac{1}{2} \int_{-1}^1 A(p, h) D(\phi - \Phi(p, h)) dh. \quad (3)$$

Here, $h \in [-1, 1]$ describes the offset of the incident ray from the fiber axis, A is the attenuation term due to volume absorption, Φ is the exiting azimuth assuming perfectly smooth fiber cylinder surface and D is the *Gaussian detector* accounting for surface roughness.

An important property of the model of d'Eon et al. that separates it from the previous work is its ability to account for fiber surface roughness in azimuthal direction, which we found essential to achieve a rich variety of hair and fur appearance as seen in Figure 2. We will cover more in detail in Section 4.



Figure 2: Species differentiation through azimuthal roughness. Photo reference of three different species (top) and fur renderings with varying azimuthal roughness (bottom) artistically matching the overall softness.

3.2. A Path-Tracing Friendly Scattering Integral

d'Eon et al. [dMH13] recommend a 70-point Gaussian quadrature to capture the complex azimuthal scattering behavior across the fiber width in Equation 3. It is justified for the cases where the illumination is easy to sample, leaving the azimuthal scattering integral over fiber width the main source of variance. However, this creates several problems for our applications:

- In the context of path-traced global illumination in production, the illumination comes from complex scenes as well as all nearby fiber geometry, requiring a large number of paths (and hence shader evaluations) to be fully resolved. Evaluating the integrand 70 times at each shading point becomes the computational bottleneck.
- The integrand becomes increasingly peaked for low azimuthal roughness and long internal path lengths. We found that for this case the 70-point Gaussian quadrature is still not sufficient to obtain an accurate result free of oscillations.
- To allow for more efficient evaluation at runtime, d'Eon et al. propose to precompute the integral into 2D tables, but this is made impractical in a production context, where shading parameters are spatially varying and provided via texture maps or even dynamic expressions.




Figure 3: The reparameterized logistic distribution (red) closely resembles a wrapped Gaussian distribution (blue).

We observe that a Monte Carlo renderer is by nature well equipped to compute statistically unbiased solutions to difficult integrals given a large number of samples. We propose to simply move the scattering integral over the fiber width in Equation 3 into the light transport integral that is solved by the rendering system. This choice avoids the expensive quadrature computation and greatly reduces the complexity of the per sample evaluation of the azimuthal scattering function, which is now parameterized over h ,

$$N_p(\phi, h) = A_p(h)D_p(\phi - \Phi(p, h)). \quad (4)$$

It achieves a significant speed up in hair and fur shading for the production examples in Figure 10. By utilizing the true intersection point of an incident ray and the hair fiber to acquire h , Equation 4 provides spatially varying reflectance across the fiber width. Hence, on top of being significantly more efficient, our model is also more accurate when viewed close up, as seen in Figure 10. We would like to note that our choice of the scattering integral is inspired by the *position-dependent BSDF* approximation proposed by Zinke and Weber [ZW07] but with the motivation of improving performance instead of its original goal of obtaining close-up accuracy.

3.3. Logistic Azimuthal Angular Distribution

The integrand in Equation 3 contains a *Gaussian detector* [dFH*11] or commonly known as a *wrapped Gaussian* distribution,

$$D(\phi) = \sum_{k=-\infty}^{\infty} g(\beta; \phi - 2\pi k), \quad (5)$$

where $g(\beta; x)$ is a Gaussian distribution with standard deviation β . The purpose of this distribution is to model azimuthal deflection due to surface roughness.

The wrapped Gaussian requires multiple Gaussian evaluations. We found that when approaching isotropic (with standard deviation $\approx 225^\circ$), up to 5 Gaussians are required to achieve negligible energy loss ($< 1\%$). Additionally, the Gaussian distribution cannot be analytically integrated and inverted, so commonly one needs to resort to sampling with two-dimensional sampling using the Box-Muller transform or with a look up table [HR12, dMH13], which are either costly or impractical because parameters often vary per shading sample.

We propose to replace the wrapped Gaussian with the *logistic*




Figure 4: Total spherical-directional reflectance of a non-absorbing (white) fiber at inclination angle 85° as a function of the number of lobes considered. High-order lobes have diminishing contributions to the total energy yet still substantial as a whole.

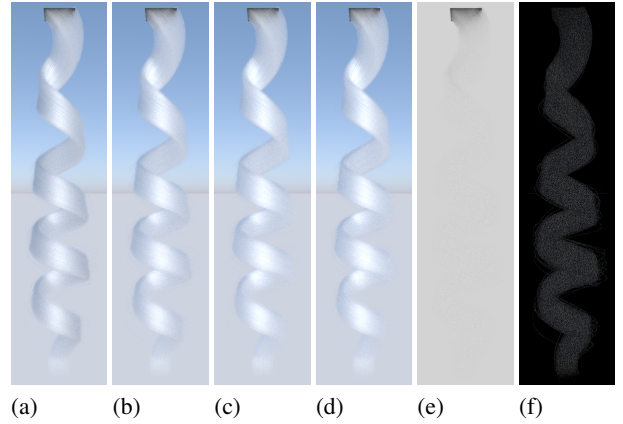


Figure 5: A lock of white straight hair rendered with sun (a half-degree distant light) and sky (an image-based light) using the first three lobes (a) and the first four lobes (b). Compared with the ground truth rendered with twenty lobes (c), the energy loss can be observed especially at near grazing angle. With a total of four lobes our method (d) conserves energy as it passes the Furnace test (e) and maintains a low difference from the ground truth (f).

distribution which has a similar shape (with reparameterization provided in Appendix A) as seen in Figure 3. What sets it apart from the wrapped Gaussian is that it can be evaluated and sampled much more efficiently. The logistic distribution is analytically integratable over arbitrary finite domain so that it can be normalized over the entire target domain, in our case, $\phi \in [-\pi, \pi]$. Hence unlike the wrapped Gaussian, *one* logistic distribution is sufficient to describe azimuthal angular distribution and achieve perfect energy conservation even at high roughness. Also, the logistic distribution has an invertible closed form CDF so it can be analytically and perfectly sampled.

3.4. A Novel Fourth Lobe

In order to achieve perfect energy conservation, Equation 1 requires us to consider all possible scattering events within the fiber, corresponding to an infinite sum of lobe evaluations. Considering only the first three lobes (R , TT and TRT) [MJC*03] results in significant energy loss especially near grazing angles for low-absorbing

fibers as seen in Figure 5. Besides, furry characters in animated film production often come with high fur density. Together with high albedo, a slight energy loss in single scattering can easily be amplified by sometimes hundreds of scattering events. Although the model of d'Eon et al. [dFH*11] can handle an arbitrary number of lobes beyond TRT , achieving energy conservation this way comes at significant cost: Figure 4 demonstrates that in order to achieve negligible energy loss ($< 1\%$) for white fur, we would need to evaluate up to 10 lobes, leading to an impractical shading cost.

Fortunately, we find that the energy carried by all infinite number of lobes can be written as a geometric series. We propose to use a fourth lobe to represent all lobes beyond TRT (starting with the very next internal scattering event $TRRT$). In this case, we get the associated energy

$$A_{\text{fourth}} = \sum_{p=TRRT}^{\infty} A_p = \frac{(1-f)^2 f^2 T^3}{1-fT}, \quad (6)$$

where A is the attenuation term as in Equation 3 and f and T are the Fresnel and the absorption factors respectively as defined in [MJC*03, dFH*11]. Also we observe that while the longitudinal angular distribution for the high-order lobes remain directional (around the specular cone), the azimuthal angular distribution becomes highly irregular. With relatively low residual energy A_{fourth} , we propose to simply approximate the azimuthal angular distribution of the fourth lobe as isotropic. With such fourth lobe, our model achieves complete energy conservation with all levels of fiber absorption at all incident angles, coming at a reasonable additional cost.

4. Controllability

Ideally, controls should directly and intuitively relate to observed appearance. Controls should also be minimal yet remain expressive enough to achieve the desired looks. This can be difficult especially in a path-tracing context where most of the appearance comes from multiple scattering. In this section we describe how we achieved these goals in our shading model's parameterization.

4.1. Roughness Reparameterization

Longitudinal The surface roughness in the longitudinal dimension is parameterized internally by the variance v of the longitudinal scattering function proposed by d'Eon et al. [dFH*11]. Although it has a physical and mathematical meaning, it is not intuitive for the artists due to unbounded range and the lack of a *perceptually uniform* relationship to its value. We define a new perceptually uniform longitudinal roughness parameter, β_M . We asked our artists to manually choose a set of images that were uniformly spaced with respect to perceived roughness, selected from the complete domain of v . Labeling these images with β_M values ranging from 0 (most directional) to 1 (isotropic), we found the following inverse mapping to correspond well to artist perception:

$$v = (0.726\beta_M + 0.812\beta_M^2 + 3.7\beta_M^{20})^2. \quad (7)$$

Azimuthal Using the same reasoning, we define a new azimuthal roughness parameter, β_N . However, instead of deriving this from perceptual experiment, we determined the physical relationship




Figure 6: Renderings of a hair block with low (left) vs. high (right) azimuthal roughness illuminated by a small frontal light source.



Figure 7: Renderings of a fur ball with increasing azimuthal roughness with (top) and without (bottom) incorporating azimuthal roughness into color parameterization process. Without it, artists would have to re-tune the color every time they adjust azimuthal roughness.

between our azimuthal and longitudinal distribution functions through the use of numerical scattering simulations on cylinders with isotropic microfacet distributions of differing roughness. Based on this discovered relationship, we reinterpreted our perceptually mapping for β_M to arrive at a comparable mapping for β_N to logistic scale factor s :

$$s = 0.265\beta_N + 1.194\beta_N^2 + 5.372\beta_N^{22}. \quad (8)$$

Figure 12 shows that both of our reparameterized longitudinal and azimuthal roughness change the multiple scattering appearance in a perceptually linear fashion, yet each in their distinct way. Longitudinal roughness primarily controls the width of the highlight, changing the perceived shininess of the fiber assembly, while azimuthal roughness controls the overall softness.

4.2. Color Reparameterization

The main factor of hair color in our model is the single-fiber absorption coefficient. However, as illustrated in Figure 8, the apparent color of the hair volume resulting from multiple scattering does not relate intuitively to the single-fiber absorption coefficient or to the single-fiber transmission color. An additional complication for color parameterization is that, as Figure 6 illustrates, azimuthal roughness also greatly affects perceived color. Azimuthal roughness behaves in effect like the phase function in a scattering medium. With lower azimuthal roughness, the hair is more forward scattering and the hair volume appears more dark and saturated with light being absorbed deeper into the volume; with higher roughness, the hair is more back scattering and the volume appears brighter and less saturated.




Figure 8: The intuitive color control of our model. A fur ball (first row) and a plastic ball (second row) are rendered with the same color chip (third row) as their color input. The color response of fur rendered with our model is very predictable and also in close agreement with surfaces. On the other hand, the last two rows are the corresponding single-fiber transmission color and absorption coefficient artists would have to pick to achieve the same fur color.

Our artists would prefer to specify the multiple-scattering color directly e.g. with a texture color map as they do for surface rendering. Conceptually this mapping can be thought as retrieving the absorption coefficient of a semi-infinite participating medium from its surface albedo. Practically, to establish this mapping, we rendered dense hair cubes, illuminated by a white dome, for a range of (scalar) absorption coefficient, σ_a , values and measured the resulting multiple-scattering albedo C averaged over the center of the cube image. Because the multiple-scattering albedo also depends on azimuthal roughness β_N , as previously shown, we repeated this process for the full range of azimuthal roughness values.

Based on the collected data points, mapping pairs of (σ_a, β_N) to the measured multiple-scattering albedo C , we found the following inverse mapping using least-squares fitting:

$$\sigma_a = (\ln C / (5.969 - 0.215\beta_N + 2.532\beta_N^2 - 10.73\beta_N^3 + 5.574\beta_N^4 + 0.245\beta_N^5))^2 \quad (9)$$

In practice, the mapping works for different hair density (Figures 8 and 9) from the one used to create the mapping. Hair density can be considered as the scattering coefficient of a semi-infinite participating medium: Changing the scattering coefficients only scales the light paths and does not change the resulting albedo. It was also pointed out by Zinke et al. [ZYWK08] that the multiple-scattering color (it is called Backscattering Attenuation by Zinke et al.) is mostly density invariant.

4.3. Additional Controls for Species Differentiation

Cuticle and Medulla Animal fiber cuticle and medullary structures can be extremely complex. While we find that the reparameterized roughness control is necessary and sufficient for perception of fur softness and variation between most species, some species with smooth cuticles and highly scattering medullary structures require an additional control. To model this observation, we expose an additional **primary reflection roughness** parameter which is a multiplier on longitudinal roughness β_M , but only for the primary



Figure 9: An example of complex color variation of furry characters in production (top). With our intuitive color parameterization, artists can specify the desired color by simply painting a texture map (bottom left) and get a predictable fur rendering result (bottom right). The texture map is visualized by applying onto the diffuse shaded surface.

reflection lobe, M_0 . This allows a layer of shiny coating on rough-looking fur as seen in Figure 13. In addition, as Yan et al. [YTJR15] pointed out, animal cuticle scales often form a layered structure, resulting in an increase of reflectance for primary reflection. For this we expose fiber **index of refraction** as an approximation.

Undercoat Human hair has just one layer of uniform fibers, but animal fur tends to include a bottom *undercoat* layer. The undercoat tends to be shorter, curly and more numerous than the top layer *guard fur*, primarily for providing thermal insulation. It is generally not feasible to model undercoat geometry explicitly due to its extreme density. We propose to approximate the effects of undercoat through shading.

As mentioned in the Section 4.2, hair density can be considered as the scattering coefficient of a participating medium. Inspired by the usage of *reduced scattering coefficient*, we propose to increase the azimuthal roughness, which has similar effects as the phase function as stated in Section 4.1, to achieve the illusion of high hair density. In addition, because under fur is more curly, we propose to also increase the longitudinal roughness to effectively simulate how the orientation deviates from the tangent direction of the explicitly modeled fur strands. With such approximation established, artists can then control the thickness of the undercoat layer by assigning high azimuthal and longitudinal roughness values along the fur length where such imaginary undercoat layer occupies. The result is shown in Figure 14.

5. Results

Performance Fur and hair rendering are computationally expensive. Brute force path tracing against hair geometry is already near the limit of affordability without considering shading cost. Figure 10 shows a production example of a furry character rendered with our choice of scattering integral and the one proposed by d'Eon et al. [dFH*11]. By avoiding the expensive 70-point Gaussian quadrature, we are able to achieve a roughly 20x speed up

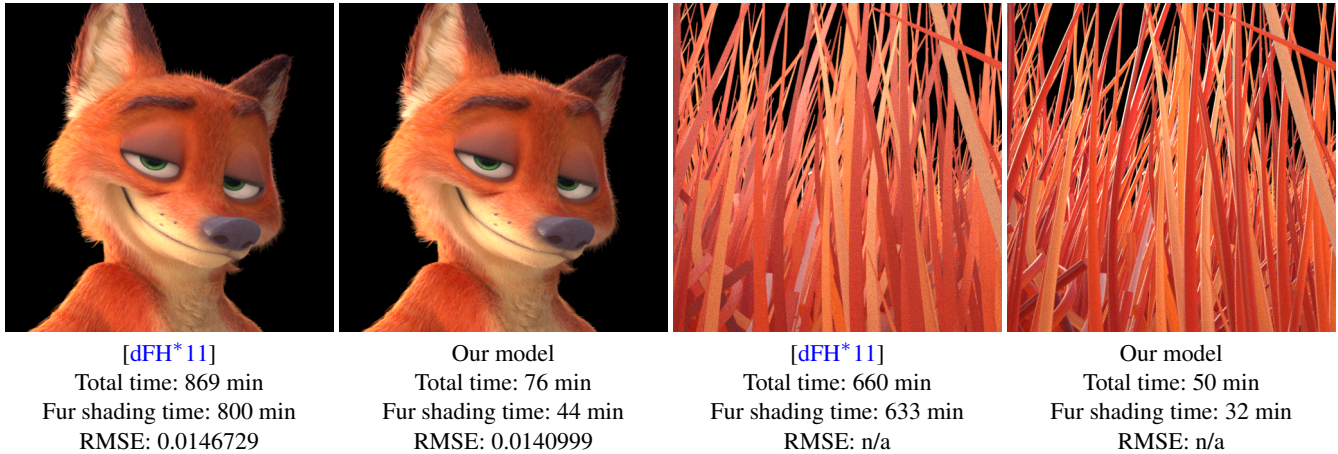


Figure 10: Comparisons of our scattering integral vs. the one proposed by d'Eon et al. [dFH*11]. A production example of a furry character illuminated by eight area lights and one image-based light (left) and its close-up view (right). Our choice of scattering integral is not only more efficient but also produces more accurate close-up results. All four images were rendered at 512 x 512 resolution with 256 samples per pixel using Hyperion, Walt Disney Animation Studio's production path tracer.

in fur shading, and a more than 10x improvement in overall render time of this particular production character. Our method takes the rendering from being extremely shading dominant (more than 80% of the total render time) to being comparable with traversal. This makes brute force path tracing of fur and hair possible in production rendering for the first time. It also implies that there is a diminishing return for further shading optimization.

Noise For the far-field comparison in Figure 10, we calculated the sampling variance of each image using a dual-buffer approach [RKZ12]. Judging by the similar sampling variance (ours: $-10.3591dB$ vs. [dFH*11]: $-10.455dB$), the two methods effectively produce images with comparable noise level. Theoretically a BCSDf [dFH*11] is designed to minimize the shading variance across fiber. We suspect that in practice there are other factors such as light sampling or geometry aliasing dominating the variance. We also evaluated the RMSE of each image against a high sample count reference image rendered with our model (ours: 0.0140999 vs. [dFH*11]: 0.0146729). It shows the BCSDf proposed by d'Eon et al. poses more error, which might be introduced by the fact that a BCSDf ignores the precise location over fiber width when evaluating illumination integral.

Shader Parameters Our unified hair and fur production shader has six parameters: **color**, **longitudinal roughness**, **azimuthal roughness**, **primary reflection roughness**, **index of refraction**, and **cuticle angle**. With these parameters we are able to reproduce a wide variety of human hair and species in our production path tracer. Figure 11 demonstrates our ability to render soft white fur with a quality unachievable by any previous method.

6. Discussion

We considered using the logistic distribution for the longitudinal angular distribution. Though the logistic is a natural choice of replacing the Gaussian distribution proposed by Marschner et al. [MJC*03], we chose to use the longitudinal scattering function

proposed by d'Eon et al. [dFH*11] for two reasons. Firstly, it produces more desirable result at high roughness, where it becomes isotropic with respect to the solid angle, whereas the longitudinal scattering function of Marschner et al. has the Gaussian defined with respect to the longitudinal angle, scattering too much energy towards the fiber's tangent direction when approaching isotropic. Secondly, the associated sampling scheme of [dFH*11] achieves perfect sampling with respect to solid angle; even though logistic distribution can be perfectly sampled there is a $1/\cos\theta_d$ in the longitudinal scattering function of Marschner et al. left unsampled.

7. Conclusions and Future Work

We have presented a new unified hair and fur shading model that is designed for path-traced multiple fiber scattering in production rendering. The implementation is similar to but more efficient than previous methods, relying on the Monte Carlo integration inherent in path tracing. Energy conservation is also made affordable for low-absorbing fibers by capturing all high-order internal scattering events using only one lobe. We have described our intuitive perceptual reparameterization and demonstrated a variety of human and animal results using just six parameters. This work has been successfully adopted by production at the Walt Disney Animation Studios. An interesting area of future work is to validate our model with the reflectance profiles of real fur provided by Yan et al. [YTJR15]. We would also like to pursue a more accurate modeling of animal undercoat distribution.

8. Acknowledgments

The authors would like to thank the Zootopia crew at the Walt Disney Animation Studios for their contribution and support of this work, especially Ernie Petti, Michelle Robinson, Brian Leach, Mitchell Snary, Alex Nijmeh, Jared Reisweber, David Kersey, Mark Siegel and Darren Robinson. We also appreciate the assistance from Dayna Meltzer. All images were rendered using Disney's Hyperion Renderer. All images © Disney Enterprises, Inc.



Figure 11: Examples of our shading model used on both human and furry characters in production.

References

- [dFH*11] D'EON E., FRANCOIS G., HILL M., LETTERI J., AUBRY J.-M.: An energy-conserving hair reflectance model. In *Proceedings of the Twenty-second Eurographics Conference on Rendering* (Aire-la-Ville, Switzerland, 2011), EGSR '11, Eurographics Association, pp. 1181–1187. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [dMH13] D'EON E., MARSCHNER S., HANIKA J.: Importance sampling for physically-based hair fiber models. In *SIGGRAPH Asia 2013 Technical Briefs* (New York, NY, USA, 2013), SA '13, ACM, pp. 25:1–25:4. [3](#), [4](#)
- [HR12] HERY C., RAMAMOORTHY R.: Importance sampling of reflection from hair fibers. *Journal of Computer Graphics Techniques (JCGT)* 1, 1 (2012), 1–17. [4](#)
- [KK89] KAJIYA J. T., KAY T. L.: Rendering fur with three dimensional textures. In *Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1989), SIGGRAPH '89, ACM, pp. 271–280. [1](#), [2](#)
- [LV00] LOKOVIC T., VEACH E.: Deep shadow maps. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2000), SIGGRAPH '00, ACM Press/Addison-Wesley Publishing Co., pp. 385–392. [2](#)
- [MJC*03] MARSCHNER S. R., JENSEN H. W., CAMMARANO M., WORLEY S., HANRAHAN P.: Light scattering from human hair fibers. *ACM Trans. Graph.* 22, 3 (July 2003), 780–791. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [MM06] MOON J. T., MARSCHNER S. R.: Simulating multiple scattering in hair using a photon mapping approach. In *ACM SIGGRAPH 2006 Papers* (New York, NY, USA, 2006), SIGGRAPH '06, ACM, pp. 1067–1074. [2](#)
- [MWM08] MOON J. T., WALTER B., MARSCHNER S.: Efficient multiple scattering in hair using spherical harmonics. In *ACM SIGGRAPH 2008 Papers* (New York, NY, USA, 2008), SIGGRAPH '08, ACM, pp. 31:1–31:7. [2](#)
- [RKZ12] ROUSSELLE F., KNAUS C., ZWICKER M.: Adaptive rendering with non-local means filtering. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 195. [7](#)
- [SPJT10] SADEGHI I., PRITCHETT H., JENSEN H. W., TAMSTORF R.: An artist friendly hair shading system. In *ACM Transactions on Graphics (TOG)* (2010), vol. 29, ACM, p. 56. [2](#)
- [YTJR15] YAN L.-Q., TSENG C.-W., JENSEN H. W., RAMAMOORTHY R.: Physically-accurate fur reflectance: modeling, measurement and rendering. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 185. [2](#), [6](#), [7](#)
- [ZW07] ZINKE A., WEBER A.: Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics* 13, 2 (Mar. 2007), 342–356. [2](#), [4](#)
- [ZYWK08] ZINKE A., YUKSEL C., WEBER A., KEYSER J.: Dual scattering approximation for fast multiple scattering in hair. In *ACM Transactions on Graphics (TOG)* (2008), vol. 27, ACM, p. 32. [2](#), [3](#), [6](#)

Appendix A: Logistic Azimuthal Angular Distribution

We start with the original logistic distribution with scale s :

$$l(x, s) = \frac{e^{-x/s}}{s(1 + e^{-x/s})^2}. \quad (10)$$

Normalizing it to a finite domain $[a, b]$ yields

$$l_f(x, s; a, b) = \frac{1}{\frac{1}{1+e^{a/w}} - \frac{1}{1+e^{b/w}}} l(x, s). \quad (11)$$

We can then reparameterize it to match the peak of a Gaussian distribution

$$l_g(x, s; a, b) = l_f(x, s\sqrt{\pi/8}; a, b). \quad (12)$$

The resulting azimuthal angular distribution over ϕ is then:

$$D(\phi) = l_g(\phi, s; -\pi, \pi). \quad (13)$$



Figure 12: Renderings of a fur ball with increasing longitudinal roughness (top) and azimuthal roughness (bottom) as described in Section 4.1.



Figure 13: Renderings of a fur ball with increasing primary reflection roughness (top) and index of refraction (bottom) as described in Section 4.3.




Figure 14: Renderings of a fur ball with increasing undercoat thickness (top), as described in Section 4.3, and a visualization of its associated distribution along the length in cross section (bottom). It demonstrates how light can be diffused from within the fur volume off the imaginary undercoat layer, while preserving the transparency of the outer guard fur at the silhouette.