Physically-Accurate Fur Reflectance: Modeling, Measurement and Rendering

Ling-Qi Yan¹

Chi-Wei Tseng²

Henrik Wann Jensen²

Ravi Ramamoorthi²

¹University of California, Berkeley

²University of California, San Diego

Figure 1: A rendering of the Wolf scene under environment lighting using (left) our physically-based double cylinder fur reflectance model with parameters from our database of animal fur samples, and (right) energy conserving Marschner model [Marschner et al. 2003; d'Eon et al. 2011] with best-fit parameters. Insets showing detailed comparisons from top to bottom using our model, Marschner model and Kajiya-Kay model. Since the Marschner model consists of only specular lobes, it often produces dark regions (limbs and tail). Furthermore, since the TT lobe is extremely strong in the Marschner model, especially for light colored fur fibers, it completely fails in heterogeneous regions (head) where dark colored fur is covered by light colored fur. The Kajiya-Kay model produces empirically plausible but hard-and-solid appearance, and it doesn't fit the measured reflectance data in Sec. 6.

Abstract

Rendering photo-realistic animal fur is a long-standing problem in computer graphics. Considerable effort has been made on modeling the geometric complexity of fur, but the reflectance of fur fibers is not well understood. Fur has a distinct diffusive and saturated appearance, that is not captured by either the Marschner hair model or the Kajiya-Kay model. In this paper, we develop a physically-accurate reflectance model for fur fibers. Based on anatomical literature and measurements, we develop a double cylinder model for the reflectance of a single fur fiber, where an outer cylinder represents the biological observation of a cortex covered by multiple cuticle layers, and an inner cylinder represents the scattering interior structure known as the medulla. Our key contribution is to model medulla scattering accurately—in contrast, for human hair, the medulla has minimal width and thus negligible contributions to the reflectance. Medulla scattering introduces additional reflection and transmission paths, as well as diffusive reflectance lobes. We validate our physical model with measurements on real fur fibers, and introduce the first database in computer graphics of reflectance profiles for nine fur samples. We show that our model achieves significantly better fits to the measured data than the Marschner hair reflectance model. For efficient rendering, we develop a method to precompute 2D medulla scattering profiles and analytically approximate our reflectance model with factored lobes. The accuracy of the approach is validated by comparing our rendering model to full 3D light transport simulations. Our model provides an enriched set of controls, where the parameters we fit can be directly used to render realistic fur, or serve as a starting point from which artists can manually tune parameters for desired appearances.

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

Keywords: fur, double cylinder, single scattering, near field

1 Introduction

In computer graphics, accurate physically-based fur rendering is often required for creating realistic furry appearance for animals. This is a long-standing problem [Kajiya and Kay 1989] with many approaches proposed to address the geometric complexity of fur. However, current fur reflectance models are mostly derived empirically or from those for human hair, such as the Kajiya-Kay [1989] or Marschner model [2003]. Fur has a distinct diffusive and saturated appearance, as shown in Fig. 1, which is not fully captured by these models since their focus is largely on the specular reflection and refraction. We also conduct measurements for a number of different types of fur, to confirm that previous hair reflectance models cannot fit reflectance profiles from fur fibers (see Figs. 8, 11, 12).

These observations motivated us to look into the literature on differences between hair and fur [Carlee and Horelick 2011; Deedrick and Koch 2004a; Deedrick and Koch 2004b] (Sec. 4). Briefly, a single fur fiber cannot be modeled as a simple dielectric cylinder similar to hair models. It often contains a non-negligible scattering structure inside called the medulla (Fig. 3), which significantly affects the appearance of a fur fiber. In contrast, for human hair, the medulla usually takes up less than one-third of the fiber diameter and can be neglected. A key insight of this paper is to take

lingqi@berkeley.edu, c5tseng@eng.ucsd.edu, {henrik,ravir}@cs.ucsd.edu

ACM Reference Format

Yan, L., Tseng, C., Jensen, H., Ramamoorthi, R. 2015. Physically-Accurate Fur Reflectance: Modeling, Measurement and Rendering. ACM Trans. Graph. 34, 6, Article 185 (November 2015), 13 pages. DOI = 10.1145/2816795.2818080 <http://doi.acm.org/10.1145/2816795.2818080>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH Asia '15 Technical Paper, November 02 – 05, 2015, Kobe, Japan.
Copyright 2015 ACM 978-1-4503-3931-5/15/11 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2816795.2818080>

medulla scattering into account for a novel physically-accurate fur reflectance model. Our major contributions are:

Double Cylinder Fur Fiber Model: In Sec. 5, we develop the physical double cylinder model for a fur fiber. As shown in Fig. 5, our model consists of three anatomically-based components—cuticle, cortex and medulla. The surface of the outer cylinder represents the cuticle, the inner cylinder represents the scattering medulla, and the cortex lies between them. We also introduce a multi-layer cuticle model to better capture reflection effects (Fig. 4).

Measurement and Database of Fur Reflectance: To validate our physical model, we use a gantry setup to measure 2D reflectance profiles of single fur fibers from nine different animals, as well as a tenth measurement on human hair for reference (Sec. 6). We fit our physical model to the measured data, showing quantitative agreement in the space of 2D reflectance profiles (Fig. 8). We also show that existing hair reflectance models like Kajiya-Kay and Marschner do not fit key features of the appearance (Figs. 11, 12). These results clearly indicate that our accurate modeling of cuticle reflection and medulla scattering are critical for fur rendering.

We provide the first database in computer graphics, of reflectance measurements and fit parameters for nine types of animal fur. The parameters can directly be plugged into our rendering model, or provide a baseline for an artist to slightly vary parameters to obtain different types of appearance. Fit parameters are listed in Table 2.

Rendering Model for Animal Fur: In Sec. 7, we develop an efficient and accurate reflectance model for single fur fibers, considering near-field scattering [Zinke and Weber 2007]. Conceptually, we trace chief specular rays within the double cylinder while we accumulate roughness (from cuticle scales) and keep track of scattering effects (from the medulla). We precompute medulla scattering to avoid explicit volumetric simulation. For consistency with previous hair models, we make a further approximation of factored lobes [Marschner et al. 2003; d'Eon et al. 2011] and represent the longitudinal and azimuthal 2D reflectance profiles separately. These approximations enable an analytic solution for the tracing step, and simplify evaluation of medulla scattering to a table lookup. In Fig. 8, we validate that our rendering model is comparable in accuracy to full volumetric simulation. We can use this fur fiber reflectance model in any global illumination renderer and for multiple types of animal fur (Fig. 1 and Sec. 8).

2 Related Work

Physically-based hair reflectance models: Marschner et al. [2003] proposed a physically-based hair reflectance model. The hair fibers are considered as rough dielectric cylinders, where three scattering paths contributing to the primary and secondary highlights are modeled: R , TT and TRT . The reflectance lobe from each path is separated into a product of longitudinal and azimuthal scattering profiles. Zinke et al. [2007] formalized hair reflectance models by introducing the notion of the Bidirectional Curve Scattering Distribution Function (BCSDF). Sadeghi et al. [2010] reformulated the model of [Marschner et al. 2003] into an artist friendly representation. d'Eon et al. [2011] extended Marschner's model from an energy conserving perspective by modeling higher-order scattering lobes such as $TRRT$, and by fixing energy-conserving issues at grazing angles. Recently, d'Eon et al. [2014] proposed a “non-separable” reflectance lobe representation by relating longitudinal contributions with relative azimuths, while still keeping a factored representation longitudinally and azimuthally. These methods produce excellent results for hair, but are not suitable for fur (see Figs. 1, 11), since they exclude scattering from the medulla, which is prominent in fur fibers. Concurrent work [Khurgum and Marschner 2015] focuses on elliptical hair fibers explicitly, revealing different optical properties compared to circular sections.

Non physically-based hair/fur reflectance models: Kajiya and

Kay [1989] introduced a methodology for rendering fur using 3D textures, together with an empirical fur shading model. The model approximates fur fibers as opaque cylinders. By extending the Phong model, it produces a diffuse lobe and a specular lobe centered around the fiber's tangent. Goldman et al. [1997] empirically improved the Kajiya-Kay model by giving it different opacity values for different viewing angles. Zinke et al. [2009] noticed the inability to fit the measured scattering from human hair fibers using Marschner's model, so they proposed an ad-hoc method by blending a diffuse lobe with Marschner's model to capture the diffusive hair reflectance observed in their measurement data. Though these methods generate plausible rendering results, they are not physically based, nor energy conserving. Moreover, they do not fit the observed reflectance profiles for fur fibers as accurately as our physically-based model.

Multiple scattering inside hair volume: Since hair contains many fibers, multiple scattering is difficult to compute. The dual scattering approximation [Zinke et al. 2008] assumes local similarity of hair strands, and derives an analytical multiple scattering model, which is later extended by [Xu et al. 2011] to enable real-time rendering and editing under environment lighting. Shadow map related methods [Lokovic and Veach 2000; Yuksel and Keyser 2008; Sintorn and Assarsson 2009] use transparency to compensate the missing transmittance of light in [Kajiya and Kay 1989]. Moon et al. [2006] applies photon mapping into the hair volume, trading noise with overblur or bias. In this paper, we focus on the optical properties from a single fur fiber, and use a standard multiple scattering renderer [Jakob 2010] to obtain global effects.

Importance sampling for hair: To be efficient, a reflectance model requires an importance sampling method in a global illumination renderer. Hery and Ramamoorthi [2012] proposed an importance sampling scheme for reflectance lobes of the Marschner model. Ou et al. [2012] extended the sampling scheme to separately sample different lobes. d'Eon et al. [2013] proposed an efficient technique based on the extended model from [d'Eon et al. 2011] by first performing a lobe selection based on the energy of each scatter type, then importance sampling the longitudinal and azimuthal scattering profiles respectively. We develop an effective importance sampling method for our fur reflectance model, based upon [d'Eon et al. 2013].

Near field scattering and far field approximation: Far-field approximation based methods [Kajiya and Kay 1989; Marschner et al. 2003; d'Eon et al. 2011] regard single hair fibers as thin curves, and assume collimated incident light rays over the width of the fiber. This approximation gives hair fibers a flat appearance at a close viewing distance. Zinke et al. [2007] proposed an analytical near-field solution to render fibers viewed up close. Our method is based on near-field scattering, and we refer to the same integration technique in [d'Eon et al. 2011] when far-field approximation is needed.

Precomputation and empirical models: Precomputation-based rendering methods such as [Sloan et al. 2002; Wang et al. 2005; Nguyen and Donnelly 2005; Ogaki et al. 2010; Ren et al. 2013] work by solving a subset of the problem in advance. These methods are usually very efficient. However, since precomputed data can take up significant storage, they are practical only under a confined set of inputs. Empirical models such as [Stam 1995; Peers et al. 2006; Donner et al. 2009] use approximations and numerical methods to get the reflectance or scattering profiles. These models often require measurements or simulations covering the entire parameter space. Our model includes a precomputation step in a low-dimensional 3D parameter space to account for medulla scattering, storing a 1D profile for each combination of these parameters. We compress the precomputed profiles to 20MB.

Figure 2: (Left) Longitudinal-azimuthal parameterization for hair/fur fibers. θ is angle to plane orthogonal to cylinder axis; ϕ is angle within plane. (Middle & Right) Illustration of Marschner model with factored representation longitudinally and azimuthally.

3 Background

Hair reflectance models treat hair fibers as cylinders, and use the Bidirectional Curve Scattering Distribution Function (BCSDF) to represent reflectance properties of a fiber,

$$L_r(\omega_r) = \int L_i(\omega_i) S(\omega_i, \omega_r) \cos \theta_i d\omega_i, \quad (1)$$

where L_i and L_r are the incoming radiance from direction ω_i , and outgoing radiance in direction ω_r respectively, and S is the BCSDF.

As shown in Fig. 2, we follow the longitudinal-azimuthal (θ, ϕ) parameterization in [Marschner et al. 2003],

$$L_r(\theta_r, \phi_r) = \int_{-\pi}^{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} L_i(\theta_i, \phi_i) S(\theta_i, \theta_r, \phi_i, \phi_r) \cos^2 \theta_i d\theta_i d\phi_i, \quad (2)$$

where the single cosine term becomes squared because the solid angle $d\omega_i = \cos \theta_i d\theta_i d\phi_i$ in this parameterization.

Kajiya-Kay model considers hair fibers as opaque solid cylinders. The reflectance is separated into a diffuse component and a specular component. Following [Zinke and Weber 2007], the BCSDF is

$$S(\theta_i, \theta_r, \phi_i, \phi_r) = k_d + k_s \frac{\cos^n(\theta_r + \theta_i)}{\cos \theta_i} \quad (3)$$

where k_d and k_s are diffuse and specular coefficients respectively. Note that the Kajiya-Kay model is azimuthally independent.

Marschner model: The BCSDF proposed by Marschner et al. [2003] regards hair fibers as glass-like dielectric cylinders. As shown in Fig. 2, it takes different specular paths $p \in R, TT, TRT$ into consideration, where R stands for reflection and T for transmission. The contribution of p is factored into a product of M and N profiles, representing longitudinal and azimuthal events:

$$\begin{aligned} S(\theta_i, \theta_r, \phi_i, \phi_r) &= \sum_p S_p(\theta_i, \theta_r, \phi_i, \phi_r) / \cos^2 \theta_d \\ &= \sum_p M_p(\theta_h) \cdot N_p(\phi; \eta') / \cos^2 \theta_d. \end{aligned} \quad (4)$$

where $\theta_h = (\theta_r + \theta_i)/2$ and $\theta_d = (\theta_r - \theta_i)/2$ are the longitudinal half angle and difference angle respectively, $\phi = \phi_r - \phi_i$ is the relative outgoing azimuth and $\eta' = \sqrt{\eta^2 - \sin^2 \theta_d} / \cos \theta_d$ is the cortex's virtual refractive index, accounting for inclined longitudinal incident directions. The types of specular paths were later extended by d'Eon et al. [2011] to handle multiple internal reflection events such as $TRRT$, etc. Also, since the original Marschner model [2003] suffers energy conservation issues addressed by d'Eon et al. [2011], we regard [d'Eon et al. 2011] as a correct implementation of the Marschner model in all our renderings and comparisons, while we keep the name Marschner model throughout the paper.

Figure 3: Structure of human hair and animal fur fibers. From left to right: section of a human hair fiber, cuticle of a human hair fiber, section of a cougar fur fiber, cuticle of a corsac fox fur fiber. Note major differences in the size of the medulla and complexity of the cuticle. Images authorized by [Wei 2006; Galat et al. 2011].

Figure 4: (Left) A microscope photograph of a sample of polar bear fur. Note the obvious coating formed by the cuticle scales. (Right) A slice of cuticle scales on human hair shaft. Images authorized by [Carrlee and Horelick 2011; Hashimoto 1988].

Near/far field scattering models: Near-field scattering specifies an actual offset h azimuthally as the incoming position (Fig. 2). For far field approximation, parallel light is assumed, covering a fiber's width. Thus, far field approximation yields the azimuthal scattering function N_p by integrating over all possible offsets $h \in [-1, 1]$,

$$N_p(\phi) = \frac{1}{2} \int_{-1}^1 N_p(h, \phi; \eta') dh \quad (5)$$

The integral for far field approximation could be solved either analytically [Marschner et al. 2003] but only for simple types of paths, or numerically [d'Eon et al. 2011] but with more computation.

4 Differences between hair and fur fibers

In this section, we describe key differences between hair and fur fibers. While the subject has received little attention in computer graphics, a number of references in other fields [Carrlee and Horelick 2011; Deedrick and Koch 2004a; Deedrick and Koch 2004b; Stamm et al. 1977] discuss microscopic variations, that we summarize here.

Hair and fur fibers share some common structures. They are often cylindrical shaped with some extent of eccentricity. As shown in Fig. 3, from outer to inner, a single fiber is divided into three layers: the *cuticle* which covers the fiber's surface with inclined scales, the *cortex* which contains nearly all colored pigments within the fiber, and the *medulla* which lies in the center of the fiber with complex internal structure that scatters light that goes through.

In spite of these common structures, hair and fur fibers do have several structural differences. Here we only introduce the most important features for our model that result in clearly different optical properties. For a comprehensive and detailed study, we refer the reader to the literature mentioned at the beginning of this section.

Medulla: The most obvious difference is that animal fur fibers usually have significant medullas inside. For human hair, the medulla is very small, and it can often be neglected. However, animal fur fibers can have medullas that hold up to the total size of the cylinder (Fig. 3).

The structure within the medulla volume is often complicated, while some animals, such as polar bears, have hollow medullas in their fur. In any case, the medulla acts as an internal scattering structure, giving the fur a generally diffusive appearance. The

Figure 5: Schematic of our double cylinder model in longitudinal section (left) and azimuthal section (right). Our model considers the medulla and cuticle effects as introduced in Sec. 4. We mark new types of paths TrT , $TrRrT$, $TttT$, $TtrtT$, $TttRttT$ that our model introduces. For clarity we hide TT and TRT paths that were previously considered by [Marschner et al. 2003] in Fig. 2. These TT , TRT paths enter the cortex, but miss the medulla entirely.

medulla could be filled with solid transparent materials or simply air, which indicates that the medulla could have a different refractive index compared to the cortex. There are usually no pigments inside the medulla.

Cuticle: While a human hair shaft has cuticle scales that resemble roof shingles, cuticles on fur can have complex shapes (Fig. 3). The outer surface of animal fur fibers is usually rougher than that of human hair. Following [d’Eon et al. 2011], we account for roughness of both longitudinal and azimuthal sections of a fur fiber in our model, assuming they have the same value for simplicity.

Additionally, as Fig. 4 shows, the cuticle layer forms a clear coat over the cortex, within which, multiple cuticle scales stack up and form a layered structure. These properties increase cuticle reflectance compared to Fresnel reflectance from a dielectric interface. In our model, we consider the cuticle as multiple layers of dielectric slab with air outside both sides of each layer. The unpolarized reflectance for each layer is given by [Stokes 1860] as

$$F(\theta_i, 1) = \frac{1}{2} \left\{ F_s(\theta_i) + \frac{[1 - F_s(\theta_i)]^2 F_s(\theta_i)}{1 - F_s^2(\theta_i)} \right\} + \frac{1}{2} \left\{ F_p(\theta_i) + \frac{[1 - F_p(\theta_i)]^2 F_p(\theta_i)}{1 - F_p^2(\theta_i)} \right\} \quad (6)$$

where F_s and F_p are s -polarized and p -polarized Fresnel reflectance respectively. Considering l layers together, the reflectance is then given by [Stokes 1860] as

$$F(\theta_i, l) = \frac{l \cdot F(\theta_i, 1)}{1 + (l - 1) \cdot F(\theta_i, 1)} \quad (7)$$

As pointed out by [Stamm et al. 1977] and [Hashimoto 1988], the internal composition of the cuticle layer may give rise to a different refractive index than the cortex. Additionally, the air interface between layers could be absent. We simplify all these properties by extending l into a real number, while using the same refractive indices for the cuticle and the cortex. The value of l is decided in our fitting step in Sec. 6. l is usually within a typical range of (0.5, 1.4), while a single dielectric interface produces a value of $l \approx 0.5$, since $F(\theta_i, 1)$ is for a double-sided slab.

5 Double cylinder fur fiber model

We now develop our physical double cylinder model for single fur fibers, based on the above observations. Our model consists of three structural components—cuticle, cortex and medulla—each with

Parameter	Definition
κ	medullary index (rel. radius length)
η_c	refractive index of cortex
η_m	refractive index of medulla
α	scale tilt for cuticle
β	roughness of cuticle (stddev.)
$\sigma_{c,a}$	absorption coefficient in cortex
$\sigma_{m,s}$	scattering coefficient in medulla
g	anisotropy factor of scattering in medulla
l	layers of cuticle

Table 1: Parameters used in our double cylinder model.

Figure 6: (Left) A photograph of a real fur fiber under bright field microscopy with medulla filled half with air and half with a mounting medium. (Right) Our Monte Carlo simulated back-lit microscopic appearance of fur fiber samples with unmounted and mounted medulla (two images stitched). Photograph publicly licensed by [Deedrick and Koch 2004a].

their respective physically-based optical properties. As shown in Fig. 5, the surface of the outer cylinder represents the (multi-layer) cuticle, the inner cylinder represents the medulla, and the cortex lies between them. The two cylinders are coaxial with relative radius 1 for the outer cylinder and κ for the inner cylinder, which is known as the *medullary index*, i.e., the ratio between the radius of the medulla and the radius of the fur fiber.

Table 1 lists all the parameters used in our model. The parameters derive largely from those in [Marschner et al. 2003]. These include the refractive index η and roughness β , the scale tilt α and absorption coefficient $\sigma_{c,a}$ for the outer cylinder (cortex). In addition, we need to consider the (probably different) refractive indices in the cortex and medulla. Moreover, since the medulla is a scattering medium, we must include its scattering coefficient $\sigma_{m,s}$ and its phase function with anisotropy g . Note that the cortex doesn’t have scattering structures inside it, and the pigments are seldom found in the medulla, as stated in Sec. 4. Hence, the scattering of the cortex and the absorption of the medulla are not required in our model. We don’t explicitly use the eccentricity parameter, but it is taken into account in the refractive indices, as in [Marschner et al. 2003].

With the double cylinder model, the types of paths through a single fiber are significantly enriched. Apart from those that were previously considered in hair models, our model introduces new types of paths that are not captured by previous methods such as TrT , $TrRrT$, $TttT$, $TtrtT$ and $TttRttT$, as shown in Fig. 5. The upper-case letters indicate interactions with the outer cylinder interface, while the lower-case ones are with the inner cylinder interface. When light transmits through the inner cylinder interface, volumetric scattering events happen.

The medulla’s significant contribution to the optical properties of a single fiber could be observed when the medulla is filled with a medium of similar refractive index as the cortex. This makes the medulla much more homogeneous and significantly reduces scattering. As shown in Fig. 6, when it is back-lit, the medulla filled with mounting medium appears nearly invisible. However, if filled with air, it is dark due to internal reflections and scattering. Our model produces similar results by tuning $\sigma_{m,s}$ and η_m accordingly. To our knowledge, this phenomenon cannot be simulated with any other hair/fur reflectance models.

Figure 7: (Left) The spherical gantry we use to measure individual fibers' reflectance profiles. (Right) Illustration of setup for two-dimensional far-field reflectance measurements.

6 Database: measurements and validation

To justify the double cylinder model, and to experimentally observe the influence of the cuticle and the medulla, we take full 2D far-field reflectance measurements of a fur fiber. We create a database of reflectance profiles using fur fibers of 9 animal species (plus a human hair). From each dataset, we also fit a set of parameters for the double cylinder model. The raw reflectance data and the fit parameters are available on <http://viscomp.ucsd.edu/projects/fur>.

Measurement Setup: Measurements are made using the UCSD spherical gantry (Fig. 7). Our gantry has two robotic arms, on which the light source and the sensor are attached at approximately 2 feet and 3 feet away from the sample respectively. We use a 150-watt DC-regulated quartz halogen bulb, a digital camera with 35 mm lens and a 12-bit 1/1.8" CCD sensor. We straighten a fur fiber on the sample plate, fix the incident direction of light at a point on the fiber, and record radiance towards discretized directions over the entire outgoing sphere. More details are found in [Tseng 2015].

We fix the light source at $(\theta_i, \phi_i) = (-40^\circ, 0^\circ)$ and capture images with the camera sweeping over $\phi_r \in [-20^\circ, 200^\circ]$ at 5° steps and $\theta_r \in [10^\circ, 50^\circ]$ at 2° steps. For simplicity, we omit the r subscripts and just use (θ, ϕ) to describe the 2D reflectance profiles. The swept range covers all the specular and diffusive scattering lobes introduced in prior human hair reflectance models, but only spans across a quarter-sphere for efficiency reasons. In other words, we implicitly assume that the fur fiber is symmetric over the incident plane. To ensure the validity of this assumption, before taking 2D measurements, we spin each fur fiber along its tangent until a qualitatively symmetric 1D normal plane reflectance profile is captured. If such symmetry in reflectance pattern is never observed, we simply dispose of the sample. The proportion of samples disposed due to asymmetry is about 10%.

For each direction, we capture 5 images for 5 stops in shutter speed (25 images in total). Each image is first cropped to a 25-by-25-pixel patch which contains the fiber. Next, the 5 patches for each shutter speed are averaged to eliminate temporal noise. Finally, we leverage the 5 averaged patches under different shutter speeds to reconstruct an HDR radiance signal. During this process, the aperture of the camera is fixed at $f/8$, and the 5 stops in shutter speed to construct HDR radiance values are 12.5, 25, 50, 100 and 200 ms. In each of the 25-by-25-pixel image patches, the sample fiber takes up approximately 100 to 200 pixels, depending on its actual width. Note that, for measurements and subsequent fits, we consider the data as gray-scale images.

Database: The 9 animal species in the database are bobcat, cat, deer, dog, mouse, rabbit, raccoon, red fox and springbok. Fur fibers are donated by a taxidermy store. We also measure the reflectance profile of a human hair for comparison and verification against prior work. Our goal is to investigate the range of reflectance patterns from fur fibers. We do not focus on taxonomic details; the names of the 9 species are only for reference. The recorded profiles show that the 10 samples (including a human hair fiber) do span a qualitatively

Figure 8: 2D reflectance profiles measured from different animals' fur fibers (left), synthesized from full 3D volumetric path tracing of a double cylinder (middle), and from our factored rendering model from Sec. 7 (right). The signals are in arbitrary units and displayed in logarithmic space to visualize perceptual brightness.

Figure 9: Illustration of general positions and shapes of all the major lobes we observed in our database.

large space in the BCSDf domain, and we consider them to include the most important reflectance phenomena.

The entire database is displayed in the left column of Fig. 8. Note that a small region of the 2D profile around $\phi = 180^\circ$ could not be measured accurately, where the camera points directly at the light source¹. Also note that many profiles resolve subtle patterns of bright and dark stripes. We believe that these stripes arise from thin-film interference between light reflected off the front and the back side of the cuticle layer.² Modeling these interference patterns in our double cylinder model is beyond the scope of the paper.

A schematic of the key lobes observed in these measurements is shown in Fig. 9, and some examples are given in the left column of Fig. 11. The name for each lobe corresponds to its contributor path, where their assignments are verified through virtual experiments using volumetric light transport simulation on 3D double cylinder models. The Marschner model fails to produce satisfying fits to the reflectance profiles of fur fibers, as shown in the right column of Fig. 11, since it supports only the R , TT and TRT lobes. Furthermore, even for these 3 lobes, there exist significant inconsistencies between what the Marschner model predicts and what we observe in the reflectance profiles. Some key features are:

- The reflectance from a fur fiber may have a blurry and partly occluded TT component around $\phi = 180^\circ$ (mouse and raccoon) or even an absent TT component (cat). Note especially the sharp edge at $\phi \approx 150^\circ$ where TT vanishes for the raccoon. Additionally, the TRT component may become dimmed (mouse) or absent (cat). These phenomena are due to light being scattered away from the original TT and TRT paths by the medulla.
- The R component in the reflectance profile of a fur fiber is usually extremely bright and blurry, and cannot be explained with Fresnel reflection alone (dog, mouse, raccoon and springbok). The phenomena are caused by the cuticle, where the reflectance is boosted by reflection off the front and the back surfaces of the cuticle.
- Forward and/or backward diffusive lobes appear in many profiles. The centers of the lobes mostly lie around the normal plane where $\theta \approx 0^\circ$ (cat, red fox, springbok). These lobes are most likely due to the medulla, which provides a volume filled with randomly distributed dielectric interfaces and is prone to multiple scattering. Additionally, less-scattered light through the medulla can still be observed in some cases as a glow around $(\theta, \phi) = (40^\circ, 180^\circ)$ (mouse, red fox).

All of these phenomena can be well modeled by the double cylinder model proposed in the previous section (middle column of Fig. 8).

Parameters and Fitting: We now fit parameter values (Table 1) for our double cylinder model from the measured 2D reflectance profiles. We can then use these parameters to define an explicit double cylinder geometry, and run a 3D volumetric light transport simulation. We compare the simulated reflectance profiles to the measurements (second column of Fig. 8). We also compare (third column) to the reflectance generated by our factored rendering model introduced in Sec. 7 with the same set of parameters.

As our model has more parameters than previous hair models, fitting parameters over 1D slices of the measured reflectance profiles as in [Marschner et al. 2003; Sadeghi et al. 2013] would lead to over-fitting. We therefore fit directly to the measured 2D profiles.

¹Since the light cannot be a perfect point light, a single fur fiber could not fully block it, thus resulting in extremely bright values. In practice, we narrow the range of the light using four paper-made walls around it, reducing the unmeasurable range to at most 10 degrees azimuthally.

²We check the difference in incident angle between rays that contribute to two consecutive bright fringes on our profiles. An analysis in optical path difference shows that to resolve the pattern requires the thickness of the cuticle to be approximately 10 micrometers, which conforms in magnitude with values reported in the literature.

Figure 10: Renderings of a fur ball with 9 different sets of fit parameters. All the images are rendered using 1024 samples per pixel with top-front area lighting.

For each fur sample, we initialize a first estimate of the parameters manually. Then, we run full 3D volumetric light transport simulation over the double cylinder to generate a synthetic 2D reflectance profile. We use an expectation-maximization algorithm to iteratively find the optimized parameters with lowest root-mean-squared error between the measured and simulated 2D reflectance signals. This optimization is quite expensive (but need only be done once offline), since simulation must be used in the inner loop of the optimization to generate simulated reflectance profiles and must be noise-free so as to provide a smooth energy landscape. However, at the resolutions of our measured reflectance profile (45×21) and using 512 samples per pixel, the optimization procedure can converge to an acceptable minimum within several hours. A similar optimization is used to fit parameters of the Marschner and Kajiya-Kay models, where needed for comparison, with the only difference being that we directly use the analytical form of these models to generate 2D profiles, instead of using Monte Carlo simulation.

Table 2 lists the optimized parameters in our double cylinder model for each fur sample. We also render a fur ball for each material in Fig. 10 using our practical rendering model that will be introduced in Sec. 7. Note that, since our measurements are gray-scale, we need to set the RGB values for the cortex absorption manually to introduce color in the renderings. For all the fur balls, we set $\sigma_{c,a} = 0.2, 0.4, 0.6$, replacing the original fit parameter $\sigma_{c,a}$.

Several observations can be made from Fig. 10. First, none of these renderings display large black regions. This is difficult for Marschner model to replicate, since its lobes are all purely specular. Figure 18 further validates this observation, showing that the Marschner model produces large black areas under area lighting. Second, forward/backward scattered lobes affect the appearance significantly. Strong forward scattering (mouse, red fox) blurs regions in the middle of these fur balls, where there are usually more fur fibers behind, so that the forward scattering energy penetrates and further scatters. Backward scattering (springbok) blurs the top and bottom regions, where the fur layer is usually thin and not viewed perpendicularly, and it visually enhances the reflection lobe. These effects indicate that the Kajiya-Kay model is not enough to represent the scattered lobe, since it is strictly uniform

Parameter	Unit	Bobcat	Cat	Deer	Dog	Mouse	Rabbit	Raccoon	Red fox	Springbok	Human
κ	unitless	0.78	0.85	0.87	0.69	0.60	0.66	0.59	0.69	0.85	0.34
η_c	unitless	1.40	1.43	1.54	1.55	1.38	1.36	1.23	1.43	1.55	1.21
η_m	unitless	1.27	1.35	1.42	1.37	1.38	1.34	1.23	1.38	1.32	1.21
α	degree	4.44	3.97	2.93	2.47	1.05	4.41	1.20	2.25	0.03	0.87
β	degree	4.86	4.94	5.35	4.21	4.70	6.97	5.27	4.86	8.43	2.03
$\sigma_{c,a}$	diameter ⁻¹	0.75	0.48	1.81	0.37	0.50	0.83	0.38	0.73	0.96	0.83
$\sigma_{m,s}$	diameter ⁻¹	3.18	2.58	2.75	3.17	2.93	2.53	3.45	2.99	3.06	4.30
g	unitless	0.54	0.62	0.39	0.18	0.65	0.31	0.35	0.63	0.03	0.38
l	unitless	0.50	0.59	0.69	0.53	0.89	0.65	1.51	0.53	0.54	1.49
Simulation NRMSE		8.1%	6.7%	8.4%	9.1%	7.8%	9.4%	12.2%	5.9%	8.4%	15.4%
Rendering model NRMSE		7.2%	5.3%	7.9%	9.1%	8.5%	8.4%	10.1%	6.3%	7.0%	19.3%

Table 2: (Top) Optimized parameters fit from our measured data using our simulated double cylinder model. Our rendering model shares exactly the same set of parameters. All length-related parameters are calculated assuming the azimuthal section of every fiber is a unit circle. All angle-related parameters are in degrees. (Bottom) Normalized RMS error of our simulated model and our rendering model.

Figure 11: Comparisons of synthesized profiles using our rendering model (left column) and Marschner model (right column, first three rows). We also compare with empirically synthesized profiles by blending a diffusive (D) lobe into the Marschner model as in [Zinke et al. 2009] (right column, last row). From top to bottom: dog fur profiles, human hair profiles, mouse fur profiles. Lobes are marked in the synthetic 2D profiles.

azimuthally, as shown in the comparison in Fig. 12. Third, for reflection lobes with similar roughness and cortex refractive indices, the one with more cuticle layers l always produces stronger reflection lobes (cat slightly brighter than bobcat, deer much brighter than dog). Thus, both our medulla and cuticle models make a qualitative difference in the appearance of fur, which cannot be captured by previous methods.

Validation: As shown in Fig. 8, we demonstrate quantitative similarities between the measured 2D profiles, simulated ones by path tracing an actual double cylinder, and profiles generated using our rendering model,³ which will be introduced in Sec. 7. The RMS errors are shown in Table 2, and are comparable for the simulation and rendering model, being below 10% in almost all cases.

We compare visually with the Marschner model in Fig. 11. Since

³Since our rendering model uses near field scattering, we integrate over azimuthal offset h to generate these far field 2D profiles, using the same technique as that introduced by [d'Eon et al. 2011].

Figure 12: Comparisons of synthesized profiles using our rendering model (left column) and the Kajiya-Kay model with specular (S) and diffuse (D) lobes (right column).

the Marschner model is purely specular, it fails in every measured profile to produce similar results, except in the case of human hair. Even when a diffuse lobe is blended into it, as suggested by [Zinke et al. 2009], it still cannot match the measured reflectance. This is because the blended-in diffuse lobe is always symmetric longitudinally at $\theta = 0^\circ$ and doesn't distinguish forward and backward scattering effects. We also show comparisons for 2D profiles generated using the Kajiya-Kay model. Since the Kajiya-Kay model is azimuthally independent and its shape changes slowly, we only compare two representative fur fibers in Fig. 12. Kajiya-Kay produces only a diffuse lobe and a specular cone, eliminating all other interesting lobes.

7 A Practical Rendering Model

The physical double cylinder model enables us to run light transport simulation on an optically faithful representation of a fur fiber. However, as evaluating integrals over multiple volumetric scattering events is costly, the model remains impractical for rendering. In this section, we develop a practical rendering model, which can be plugged into global illumination renderers.

7.1 Overview

Under the far-field approximation, most prior reflectance models for hair obtain a compact BCSDf for the fiber. However, as introduced in Sec. 2, to model more sophisticated appearances of fur at a close distance, it is essential to include near-field effects; if necessary, these could be integrated to obtain far-field approximation. Thus, we analyze the potential scattering paths for an incoming ray hitting the double cylinder, based on its incident position and direction.

We make this approach practical through a few key approximations. First, our near-field model incorporates factored approximation for

rendering, similar to far-field hair models [Marschner et al. 2003; d'Eon et al. 2011]. However, unlike those models, which consider a longitudinal and an azimuthal reflectance profile being generated collectively by all collimated incident rays over the width of the fiber, we consider an azimuthal and a longitudinal distribution being generated by each ray we trace through the fiber. This approach is equivalent to having varying BCSDFs over h in Fig. 5, which makes our model compatible with the far-field BCSDF framework, yet enables near-field effects. In [Zinke and Weber 2007], the approach we use is classified as a near-field scattering model with constant incident illumination. Hence, we can leverage previous work, keeping the R , TT and TRT terms identical, but adding lobes for TrT , $TttT$, $TtrtT$, $TrRrT$ and $TttRttT$ paths, indicated in Figs. 5 and 13. Note that our framework can also handle general higher-order scatterings, but in our experiments, the contributions from those paths were insignificant.

It remains to find the forms of azimuthal and longitudinal scattering functions for each lobe. Conceptually, given the incident position and direction of a ray that enters the fiber, we ray-trace its chief specular ray [d'Eon et al. 2011; Zinke et al. 2008] on the 2D azimuthal and longitudinal cross sections of the fiber. We consider the direction at which the chief specular ray leaves the fiber as the center of a reflectance lobe, and accumulate the attenuation factors, and the azimuthal and the longitudinal roughnesses along the specular path, in the form of a Gaussian outgoing lobe. This is shown for the $TttT$ path $PQ'Q'P$ in the azimuthal section in Fig. 13. In practice, rays are traced using closed-form analytical formulae. Longitudinal distributions are further simplified to conform to previous work, reducing to Gaussians with offsets/width for effective cuticle tilt/roughness. Finally, for those paths that enter the medulla, and are *scattered* by it, we must also include a scattered lobe (see broad yellow lobe at Q' in Fig. 13). We precompute medulla scattering separately for 2D azimuthal/longitudinal profiles (Fig. 14), based on [Donner et al. 2009], reducing scattering to a table lookup.

Formally, the near-field scattering distribution for each ray is,

$$S(\theta_i, \theta_r, \phi_i, \phi_r, h) = \frac{\sum_p M_p^u(\theta_i, \theta_r) N_p^u(h, \phi)}{\cos^2 \theta_i} + M^s(\theta_i, \theta_r, \phi) \frac{\sum_p N_p^s(h, \phi)}{\cos^2 \theta_i} \quad (8)$$

$$p \in \{R, TT, TRT, TrT, TttT, TtrtT, TrRrT, TttRttT\},$$

where $\phi = \phi_r - \phi_i$, while M_p and N_p are respectively longitudinal and azimuthal scattering profiles that depend on the fur parameters in Table 1. We generalize equation 4 in a few respects. Note the h parameter (Figs. 5, 13) for near-field scattering. The superscripts s and u stand respectively for paths *scattered* by the medulla, and *unscattered*. (We will see that a single longitudinal M^s is adequate for all scattered paths p). The paths p include R , TT and TRT from previous hair models, as well as new terms specific to fur. Note that only unscattered lobes will be present for R , TT , TRT , TrT , $TrRrT$ paths that never transmit into the medulla, while $TttT$, $TtrtT$ and $TttRttT$ will have both scattered and unscattered lobes. Finally, we choose to replace the $1/\cos^2 \theta_a$ term with $1/\cos^2 \theta_i$ per [Hery and Ramamoorthi 2012] to mitigate the energy conservation issues at grazing angles at the cost of losing reciprocity, as analyzed by d'Eon et al. [2011]. This substitution also simplifies our rendering model, since the cosine terms are completely cancelled out in equation 2.

Simple Lobes: Consider rays that do not enter the medulla at all ($p \in \{R, TT, TRT, TrT, TrRrT\}$). In this case, there is no scattering from the medulla, and we can drop the superscripts. The azimuthal profile is simply

$$N_p(h, \phi) = A_p(h) \cdot D_p(h, \phi), \quad (9)$$

where A_p is the attenuation along path p , considering Fresnel terms on the cuticle and absorption along interior paths. D_p is the az-

Figure 13: Illustration of evaluating the azimuthal scattering function for the type of path $TttT$.

imuthal distribution of scattered energy, a Gaussian with width determined by considering the roughness from all the surfaces the ray has interacted with. Both A_p and D_p share the same representation with unscattered lobes in equations 12 and 13. For the longitudinal profile of a ray, we follow [Marschner et al. 2003; d'Eon et al. 2011] and approximate it with a Gaussian distribution⁴,

$$M_p(\theta_i, \theta_r) = G(\theta_r; -\theta_i + \alpha_p, \beta_p), \quad (10)$$

where α_p is the accumulative angular tilt of the chief specular ray on path p due to interaction with the cuticle scales, and β_p is the roughness for path p , which is empirically given by accumulating cuticle roughness. These expressions reduce to previous work for the lobes R , TT , TRT from the hair model. We now proceed to develop general formulae for paths that interact with the medulla.

7.2 Unscattered Lobes

We first discuss lobes that are not scattered by the medulla, giving expressions for M_p^u and N_p^u . Consider the azimuthal profile in Fig. 13. Conceptually, we simply ray-trace the chief specular ray incident at P in the 2D azimuthal cross-section, accumulating attenuation and surface roughness along the path to give the intensity and width of the outgoing Gaussian lobe centered at P' . In practice, ray-tracing can be replaced with simple analytic formulae.

An interesting observation is that certain types of paths only happen in specific zones over the offset h (see Fig. 13). For example, depending on the size of the medulla, the traditional hair model paths T , TT and TRT are only possible for large h , and must enter the medulla otherwise. In practice, we solve geometrically for the boundaries of these zones, corresponding to $\gamma_{i2} = 90^\circ$ and $\gamma_{t2} = 90^\circ$, and only consider the relevant paths within each zone.

Azimuthal Scattering Profile: The azimuthal profile is given by equation 9, but additional handling is required for chief specular rays that hit and enter the medulla ($p \in \{TttT, TtrtT, TttRttT\}$). We split the contribution into two terms: the multiple scattered light for which we compute M^s and N^s in the next sub-section based on precomputed medulla scattering; and the unscattered light considered here, which is simply attenuated as it passes through the medulla. For instance, in a $TttT$ path as shown in Fig. 13,

$$N_p^u(h, \phi) = A_{Tt}(h) \cdot A_{tt}(h) \cdot A_{tT}(h) \cdot D_p^u(h, \phi). \quad (11)$$

⁴There are known energy leak issues at grazing angles since Gaussian lobes are unbounded, as analyzed and accurately solved in [d'Eon et al. 2011]. However, the accurate solution requires heavy computation and has numerical precision issues. Instead, we simply fold energy cutoffs outside $[-90, 90]$ degrees back (e.g. 93 degrees back to 87 degrees), which works well in practice.

Here, $A_{tt}(h)$ is the attenuation factor for a chief specular ray transmitting through the medulla (QQ' in Fig. 13).

With this background, a simple expression can be written for general higher-order paths, which can also be used for our $TtrtT$ and $TttRttT$ lobes. If we regard p as a string of length n , while p_i ($i = 1 \dots n$) represents each vertex in p , we can write,

$$A_p(h) = \prod_{i=1}^n F(p_i) \cdot \prod_{i=2}^n \exp(-\sigma_t(p_{i-1}p_i) \cdot |p_{i-1}p_i|) \quad (12)$$

where F is the (extended) Fresnel term, $|p_{i-1}p_i|$ is the length of segment $p_{i-1}p_i$, and $\sigma_t(p_{i-1}p_i)$ equals either $\sigma_{c,a}$ of the cortex or $\sigma_{m,s}$ of the medulla, as the segment's extinction coefficient.

Similarly, at every intersection, the direction of p alters by angle $\Gamma(p_i)$. Thus, the outgoing azimuth could be computed by accumulating these deviation angles as $\Phi_p(h) = \pi + \sum_{i=1}^n \Gamma(p_i)$, where π accounts for the inversion of the incoming direction. For the distribution term D , similar to [d'Eon et al. 2011; Zinke et al. 2008], we accumulate the roughness i.e., β^2 at each intersection p_i along the path p . Since the outgoing distribution is a Gaussian lobe G centered at $\Phi_p(h)$, we can derive an analytic form,

$$D_p^u(h) = G\left(\Phi_p(h), \sqrt{\sum_{i=1}^n \beta^2(p_i)}\right) \quad (13)$$

where $\beta(p_i) = 0$ if the intersection p_i is not on the cuticle. Explicitly, the width is given by the number of upper-case (cortex) letters, and is β for $p \in \{R\}$, $\sqrt{2}\beta$ for $p \in \{TT, TrT, TttT, TtrtT\}$, and $\sqrt{3}\beta$ for $p \in \{TRT, TrRrT, TttRttT\}$.

Longitudinal Scattering Profile: The unscattered longitudinal profile is given by equation 10. All that remains is to determine the center α_p and width β_p of the unscattered lobes. We follow previous work for R and TT lobes, setting $(\alpha_R, \beta_R) = (\alpha, \beta)$ for R , corresponding to the cuticle tilt and roughness, while $(\alpha_{TT}, \beta_{TT}) = (-\alpha/2, -\beta/2)$. For other lobes, we approximate

$$\begin{aligned} \alpha_p &= \alpha_{TT} - n_R \alpha \\ \beta_p &= \beta_{TT} + n_R \beta + (n \bmod 2)(\beta/2), \end{aligned} \quad (14)$$

where n is the length of p , and n_R is the number of R s appearing in p . The general idea is that, every reflection on the cuticle decreases its tilt angle and increases roughness. The final $\beta/2$ compensates for lobes exiting backwards, since they are usually wider than forward lobes [d'Eon et al. 2014]. Thus, $\alpha_p = \alpha$ for $p \in R$, $\alpha_p = -\alpha/2$ for $p \in \{TT, TrT, TttT, TtrtT\}$ and $\alpha_p = -3\alpha/2$ for $p \in \{TRT, TrRrT, TttRttT\}$. Similarly, $\beta_p = \beta/2$ for $p \in \{TT, TttT\}$, $\beta_p = \beta$ for $p \in \{R, TrT, TtrtT\}$ and $\beta_p = 2\beta$ for $p \in \{TRT, TrRrT, TttRttT\}$. Note that, to our knowledge, longitudinal lobes' parameters were empirically given in most previous work; equation 14 provides a reasonable approximation to extend prior work to general paths that may enter the medulla.

7.3 Scattered Lobes

When the path $p \in \{TttT, TtrtT, TttRttT\}$ goes through the inner cylinder, scattering events happen, and this medulla scattering needs to be taken into account (the broad yellow lobe from Q' in Fig. 13). We rely on precomputation, and make a number of significant approximations to the full volumetric multiple scattering computation, to enable a practical rendering model.

Precomputation of Medulla Scattering: Our precomputation approach is similar in spirit to the empirical BSSRDF model by

⁵In the precomputation, we assume all sub-paths are entering the medulla horizontally, since the azimuthal section is rotationally-invariant.

Figure 14: Precomputing medulla scattering. We enumerate $\sigma_{m,s}$ and g , and vary azimuthal offset h' and longitudinal incident angle θ'_i respectively.⁵ We store the yellow-marked scattered lobe in every outgoing azimuth ϕ' and longitudinal outgoing angle θ'_r .

Donner et al. [2009]. However, since we use factored lobes, we precompute scattering profiles by 2D volumetric path-tracing separately for azimuthal and longitudinal components, for all possible combinations of scattering parameters, as illustrated in Fig. 14. Here, we don't consider surface effects (reflection, refraction, Fresnel, etc) for the medulla; instead we compute these effects in the evaluation steps. In effect, we are precomputing 4D tables $C^N(\phi'; h', \sigma_{m,s}, g)$ azimuthally and $C^M(\theta'_r; \theta'_i, \sigma_{m,s}, g)$ longitudinally (1D profiles for 3D sets of parameters). We use primes to distinguish notation from the main parameters. Our precomputation is entirely scene-independent, and only needs to be done once. We also make C^N and C^M available online, so other researchers can use them directly. After compression, these tables takes up only about 20MB. The appendix discusses details.

Azimuthal Scattering Profile: The scattering lobe is usually large and diffusive, and therefore not significantly affected by the smaller effects of surface roughness. We also ignore refraction by the cortex-air interface, and assume that the light leaving the medulla is attenuated by a constant factor corresponding to the thickness of the cortex. Moreover, we assume that the precomputed scattered lobe from the medulla doesn't change its shape after transmitting outside. Finally, we assume that there is only one scattered lobe that is not further reflected by the cortex; the appendix discusses a first step towards relaxing this assumption.

In analogy to equation 11, the azimuthal scattering profile from a $TttT$ path is,

$$N_p^s(h, \phi) = A_{Tt}(h) \cdot A'_{Tt}(h) \cdot D_p^s(h, \phi). \quad (15)$$

Note that the distribution D^s is not normalized, and accounts for the reduction in energy $1 - A_{tt}(h)$ due to the unscattered lobe already considered. The final attenuation A'_{Tt} is now approximated simply as $\exp[-\sigma_{c,a} \cdot (1 - \kappa)]$, corresponding to the thickness of the cortex. We simply need to add one more attenuation term for $TtrtT$,

$$N_p^s(h, \phi) = A_{Tt}(h) \cdot A_{tr}(h) \cdot A'_{Tt}(h) \cdot D_p^s(h, \phi). \quad (16)$$

Note that the scattered lobe arises only on the final rt segment; any scattering in the earlier segment will be considered as part of $TttT$. A similar expression can be used for the $TttRttT$ lobe. More generally, the attenuation of a scattered lobe consists of two parts. First, the energy reaches into the medulla. Second, the scattered energy is further absorbed, transmitting through the cortex.

The distribution term D for a scattered lobe is simply a query into the precomputed azimuthal scattering profile $C^N(\phi'; h', \sigma_{m,s}, g)$. Since we precompute the medulla as a unit circle, h' and $\sigma_{m,s}$ need to be normalized by radius κ , leading to

$$D_p^s(h, \phi) = C^N(\phi - \Phi_p(h); h'/\kappa, \sigma_{m,s}/\kappa, g), \quad (17)$$

where Φ_p is the angle at which the ray enters the medulla.

Figure 15: Illustration for computing longitudinal scattered lobe M^s . Refractions are considered here at P , Q , P' and Q' .

Longitudinal Scattering Profile: The longitudinal scattered lobe M^s is p -independent, because every type of path p has the same θ_i and θ_r . However, it is ϕ -dependent. As Fig. 14 illustrates, the longitudinal profiles we precompute are for $\phi = 0$ (the upper lobe) and $\phi = \pm 180^\circ$ (the lower lobe). We compute M^s at both azimuths, and linearly interpolate the results for any ϕ . Thus, for simplicity we omit ϕ in the following, and take $\phi = 0$ for illustration.

We query the precomputed longitudinal distribution for the medulla C^M , considering refractions through the cuticle and the cortex. As Fig. 15 shows, θ'_i and θ'_r could be solved geometrically. Similar to the azimuthal case, we query the precomputed distribution at θ'_r ,

$$M^s(\theta_i, \theta_r) = \mu \cdot F_t \cdot C^M(\theta'_r; \theta'_i, \sigma_{m,s}/\kappa, g), \quad (18)$$

where $F_t = (1 - F(\theta_r + \alpha, l)) \cdot (1 - F(\theta'_r))$ is the product of (extended) Fresnel transmittance, and μ is the normalization factor which is described in detail in the appendix.

7.4 Implementation Details and Validation

Importance Sampling: Our importance sampling scheme is similar to [d'Eon et al. 2013], where we first perform a lobe selection, then sample this lobe azimuthally and longitudinally. Here, we treat all scattered lobes together. If an unscattered lobe p is chosen, we sample a Gaussian around its azimuthal outgoing center $\Phi_p(h)$ and its longitudinal outgoing center $-\theta_i + \alpha_p$. This leads to a near-perfect importance sampling scheme. If the (summed up) scattered lobe is chosen, we perform a cosine-weighted sampling longitudinally and a uniform sampling azimuthally, taking advantage of the fact that the scattered lobe is smooth. For multiple importance sampling, which queries the PDF at a given outgoing direction, we first perform a lobe selection similarly, then compute the corresponding PDF value at the outgoing direction, depending on whether the scattered or unscattered lobe is selected.

Non-separable lobes: In [d'Eon et al. 2014], the non-separable lobes representation was introduced to accurately capture light scattering through a hair fiber. In this representation, the center and width of longitudinal lobes further depends on the relative azimuth ϕ . Indeed, there are quality improvements in rendering results from the original paper. However, we find that the main difference is the shape of the R lobe, which is the only one that spans a large range over azimuthal angles. Thus, when applied to our double cylinder model, we simply represent our R lobe as non-separable, and leave other lobes using the traditional representations.

Validation: We compare 2D profiles generated using our rendering model with measured data and simulations in Fig. 8. Our rendering model closely matches the measured reflectance profiles, and has comparable error (Table 2) as a full simulation, in some cases even being closer to the measurements. This is not surprising, since the physical double cylinder model is exactly the same. Minor discrepancies are due to approximations, such as factoring longitudinal and azimuthal scattering profiles, empirical longitudinal lobe centers and widths, and medulla scattering approximations.

	Wolf	Chipmunk	Cat	Fur pelt
# Strands	1.6 M	503 K	729 K	12.5 K
# Segs	8	8	10	5
Area light		✓		✓
Env. light	✓	✓	✓	
# Samples	1600	1600	2500	1024
Time	60.8 min	23.6 min	56.2 min	12.4 min

Table 3: Statistics for all of our scenes. We represent fur fibers using line segments for each fur fiber. # Strands is the number of fur fibers, and # Segs is the number of segments along each fur fiber. # Samples is the number of samples per pixel.

8 Results

In this section, we show results generated using our rendering model, and visual comparisons to previous methods. All results use our rendering model as a shader within Mitsuba [Jakob 2010], run on an Intel 6-core 3.6 GHz i7 4960X CPU, hyperthreaded to 12 threads. Statistics of each scene such as number of strands, number of samples and timings are listed in Table 3. The times are total wall clock running time, including global illumination; the cost of evaluating our model is comparable to that of the Marschner model, which produces almost identical timings. For all the scenes, all parameters are directly derived from our database in Table 2, except for the absorption term $\sigma_{c,a}$ which accounts for different colors.⁶ We also account for the fact that the medulla does not absorb light, by multiplying $\sigma_{c,a}$ by the medullary index κ . For wolf and chipmunk, since we don't have corresponding fur samples, we refer to parameters for dog and mouse from our database instead.

In comparisons, we consider [d'Eon et al. 2011] as a correct implementation of the Marschner model, so that the dark appearance generated is due to specularly of the model itself, rather than energy conservation issues. For the Marschner model, we use its optimized fitting parameters from Fig. 11, rather than ad-hoc settings from industry such as enlarging azimuthal roughness. For the Kajiya-Kay model, we enable global illumination by normalizing equation 3, since the model itself is not energy conserving.

Wolf: Figure 1 shows our rendering result for a Wolf model with a side-by-side comparison with the Marschner model. The wolf model is placed on a turn-table to demonstrate consistency of our rendering model, by rotating it in our accompanying video. The environment lighting is manually blurred prior to rendering. Insets are provided to compare with the Marschner and Kajiya-Kay models. All the renderings are path traced (for global illumination between fibers and the environment) using 1600 samples per pixel (spp) at a resolution of 1920×1080 . Our double cylinder model produces a diffusive and saturated appearance, while the Marschner model is highly specular and dark. Note that, our method actually produces brighter highlights with the layered cuticle model than Marschner, but these are less visually obvious due to low local contrast. The Kajiya-Kay model produces a hard and solid appearance even with global illumination. Intuitively, compared to our model, it is visually similar to a BRDF vs BSSRDF (BCSDF) comparison.

Cat: In Fig. 16, we render a close-up view of a cat head with depth of field effects under environment lighting, showing details of each fur fiber. We use our optimized fit parameters for cat fur in our database. All renderings use a sampling rate of 2500 spp with resolution 1024×1024 . We can see that our rendering produces both a blurred area (top left of the eye) and distinct appearance (around the whisker) due to different arrangements and orientations.

We also show how an artist can manually vary key parameters of the

⁶For comparisons, we use colored textures to assign $\sigma_{c,a}$ in the Marschner model for each fiber rooted at texture coordinate (u, v) as $\sigma_{c,a} = -\log(T(u, v))/4$, where $T(u, v) \in [0, 1]$ is the texture color at (u, v) , considering each color channel.

Figure 16: Renderings of the Cat scene under environment lighting using our rendering model, with (left) increasing scattering coefficients $\sigma_{m,s}$ and (right) increasing medulla size. Note the differences between the first and fifth images, where the fifth image produces a strong TT term on the eyebrows and a clear secondary highlight on the forehead, which are characteristics from the Marschner model, because it is using a small medullary index or radius $\kappa = 0.25$.

Figure 17: Renderings of the Chipmunk scene using (top) our double cylinder rendering model and (bottom) Marschner model illuminated with a strong area light and a dim environment light.

model to get a range of appearances. Figure 16 shows renderings with varying scattering coefficients $\sigma_{m,s}$, as well as medulla size κ . We observe that, artistically, $\sigma_{m,s}$ closely controls color saturation, and κ determines the “specularity”, or the extent of similarity between hair and fur. In the video, we rotate environment lighting to show color transitions of the cat fur from gold to dark orange under different lighting conditions.

Chipmunk: Figure 17 is rendered with a sharp and strong area light, and a relatively dim environment light. The skin of the chipmunk is dark colored. The Marschner model again produces unrealistically specular and dark appearance, since the light easily penetrates the fibers and hits the skin. However, primary (uncolored) and secondary (colored) highlights are still visible in the Marschner model. Both models are rendered using 1600 spp at resolution

Figure 18: Renderings of the Fur pelt scene under area lighting. (Top left) Our rendering model. (Top right) Marschner model. (Bottom left) Kajiya-Kay model. (Bottom right) Marschner model blended with diffuse lobe.

1920 × 1080. In the video, we rotate the area light to see moving highlights, while other parts never get fully dark in our model.

Fur pelt: Figure 18 contains a pelt of fur placed on a checker board rendered using 1024 spp with resolution 1024 × 1024. A large area light illuminates the pelt from the top-left. Our model gives a realistic diffusive and saturated appearance, while the Kajiya-Kay model looks hard and solid. The Marschner model produces classic primary and secondary highlights, but leaves other regions black. By blending a diffuse lobe into the Marschner model as proposed by [Zinke et al. 2009] (a solution that is widely adopted by the industry), one can generate a diffusive appearance. However, the blending technique lowers the intensity of the original lobes in the Marschner model, especially for the reflected lobe R , which leads to a flat appearance. Furthermore, this approach is empirical and the existence of the diffuse lobe cannot be explained physically.

9 Discussion and Limitations

More complex animal fur: In the real world, animal fur could be even more complicated than what our model can represent. Our observations are focused only on some of the structural features that affect the optical properties — we capture the most important optical properties of fur fibers that have been omitted from previous work but we ignore other aspects. For example, currently we only focus on symmetric fibers, and we discard measured samples without symmetry in the reflectance pattern. However, we believe that our model is a good starting point to further account for other interesting features such as irregular-shaped azimuthal sections. For more complete observations on animal fur, we refer readers to [Deedrick and Koch 2004a; Deedrick and Koch 2004b].

Medulla scattering approximations: A key aspect of the rendering model is an efficient precomputation of medulla scattering, which was not required in previous work. Since we seek a factored lobe approximation, we have considered separate simulations for azimuthal and longitudinal profiles. While this factored precomputation produces accurate results, it does involve an approximation of 2D scattering in orthogonal azimuthal/longitudinal planes, that needs further evaluation. Future work could explore the limits of this approximation, and more general scattering lobes.

Application to human hair: Previous hair models could also benefit from our work. The most important reason is that there could be visible medullas inside human hair fibers [Deedrick and Koch 2004a]. Moreover, artists may want hair to exhibit a more diffusive and saturated appearance, which is not easily handled using previous models, especially when the number of hair strands is limited.

10 Conclusions and Future Work

We present a physically-accurate fur reflectance model to accurately capture the appearance of animal fur fibers. Our model is near-field based, treating fur fibers as double cylinders, taking into account the existence of scattering medullas inside. We also derive a fast evaluation algorithm for rendering, built upon precomputed empirical medulla scattering profiles. We demonstrate that our rendering model fits the measured data well, with errors comparable to a full volumetric simulation. We also introduce the first database in computer graphics of reflectance measurements on a number of animal fur fibers, including both raw 2D scattering profiles and fit parameters. We show that our model is capable of generating a variety of realistic animal fur appearances, with significantly more realistic results than previous methods.

In the future, we would like to explore possibilities to make our model completely analytical without precomputation. An extension to handle more features such as irregular fiber sections, complicated cuticle scale arrangements and discontinuous medullas would also be interesting. Industry may benefit if a re-interpretation from an artistic viewpoint is derived. Finally, we could accelerate our rendering using methods like cone tracing [Qin et al. 2014]. A fast and accurate way to do approximate multiple scattering using our model would contribute greatly in real-time fur rendering applications.

11 Acknowledgements

The authors are indebted to Alexander Alvarado and Jared Reisweber, who created the cat and chipmunk models respectively. We also thank Matt Chiang and Brent Burley for helpful discussions. The fur samples are donated by Lyons and O'Haver Taxidermy. This work was supported in part by NSF grant 1451828, Intel, and gifts from Pixar, Sony, Adobe and Qualcomm to the UC San Diego Center for Visual Computing.

References

- CARRLEE, E., AND HORELICK, L. 2011. The alaska fur id project: A virtual resource for material identification. In *Objects Specialty Group Postprints*, American Institute for Conservation of Historic and Artistic Works, vol. 18, 149–171.
- DAVIS, A. 2006. Effective propagation kernels in structured media with broad spatial correlations, illustration with large-scale transport of solar photons through cloudy atmospheres. In *Computational Methods in Transport*, F. Graziani, Ed., vol. 48 of *Lecture Notes in Computational Science and Engineering*. Springer Berlin Heidelberg, 85–140.
- DEEDRICK, D. W., AND KOCH, S. L. 2004. Microscopy of hair part i: A practical guide and manual for human hairs. *Forensics Science Communication*.
- DEEDRICK, D. W., AND KOCH, S. L. 2004. Microscopy of hair part ii: a practical guide and manual for animal hairs. *Forensics Science Communication*.
- D'EON, E., FRANCOIS, G., HILL, M., LETTERI, J., AND AUBRY, J.-M. 2011. An energy-conserving hair reflectance model. In *EGSR 11*, 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*, 25:1–25:4.
- D'EON, E., MARSCHNER, S., AND HANIKA, J. 2014. A fiber scattering model with non-separable lobes. In *ACM SIGGRAPH 2014 Talks*, 46:1–46:1.
- DONNER, C., LAWRENCE, J., RAMAMOORTHY, R., HACHISUKA, T., JENSEN, H. W., AND NAYAR, S. 2009. An empirical BSSRDF model. *ACM Trans. Graph.* 28, 3, 30:1–30:10.
- GALATK, A., GALATK, J., KRUL, Z., AND GALATK JR, A., 2011. Furskin identification. <http://www.furskin.cz>.
- GOLDMAN, D. B. 1997. Fake fur rendering. In *SIGGRAPH 97*, 127–134.
- HASHIMOTO, K. 1988. The structure of human hair. *Clinics in Dermatology* 6, 4, 7–21.
- HERY, C., AND RAMAMOORTHY, R. 2012. Importance sampling of reflection from hair fibers. *Journal of Computer Graphics Techniques (JCGT)* 1, 1, 1–17.
- JAKOB, W., 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with three dimensional textures. In *SIGGRAPH 89*, 271–280.
- KHUNGURN, P., AND MARSCHNER, S. 2015. Azimuthal scattering from elliptical hair fibers. *Accepted to ACM Transactions on Graphics with minor revisions*.
- LOKOVIC, T., AND VEACH, E. 2000. Deep shadow maps. In *SIGGRAPH 00*, 385–392.
- MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. *ACM Trans. Graph.* 22, 3, 780–791.
- MOON, J. T., AND MARSCHNER, S. R. 2006. Simulating multiple scattering in hair using a photon mapping approach. *ACM Trans. Graph.* 25, 3, 1067–1074.
- NGUYEN, H., AND DONNELLY, W. 2005. Hair animation and rendering in the nalu demo. *GPU Gems 2*, 361–380.

- OGAKI, S., TOKUYOSHI, Y., AND SCHOELLHAMMER, S. 2010. An empirical fur shader. In *ACM SIGGRAPH ASIA 2010 Sketches*, ACM, 16.
- OU, J., XIE, F., KRISHNAMACHARI, P., AND PELLACINI, F. 2012. Ishair: Importance sampling for hair scattering. In *ACM SIGGRAPH 2012 Talks*, 28:1–28:1.
- PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHY, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. *ACM Trans. Graph.* 25, 3, 746–753.
- QIN, H., CHAI, M., HOU, Q., REN, Z., AND ZHOU, K. 2014. Cone tracing for furry object rendering. *Visualization and Computer Graphics, IEEE Transactions on* 20, 8, 1178–1188.
- REN, P., WANG, J., GONG, M., LIN, S., TONG, X., AND GUO, B. 2013. Global illumination with radiance regression functions. *ACM Trans. Graph.* 32, 4, 130:1–130:12.
- SADEGHI, I., PRITCHETT, H., JENSEN, H. W., AND TAMSTORE, R. 2010. An artist friendly hair shading system. *ACM Trans. Graph.* 29, 4, 56:1–56:10.
- SADEGHI, I., BISKER, O., DE DEKEN, J., AND JENSEN, H. W. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.* 32, 2, 14:1–14:12.
- SINTORN, E., AND ASSARSSON, U. 2009. Hair self shadowing and transparency depth ordering using occupancy maps. In *Symposium on Interactive 3D Graphics and Games*, 67–74.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph.* 21, 3, 527–536.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Rendering Techniques 95*, 41–50.
- STAMM, R. F., GARCIA, M. L., AND FUCHS, J. J. 1977. The optical properties of human hair i. fundamental considerations and goniphotometer curves. *Journal of the Society of Cosmetic Chemists* 28, 9, 571–599.
- STOKES, G. G. 1860. On the intensity of the light reflected from or transmitted through a pile of plates. *Proceedings of the Royal Society of London* 11, 545–556.
- TSENG, C.-W. 2015. *A Physically-Based Reflectance Model For Mammalian Fur Fibers Based On Anatomy And Gonioreflectometry Measurements*. Master's thesis, University of California San Diego.
- WANG, R., TRAN, J., AND LUEBKE, D. 2005. All-frequency interactive relighting of translucent objects with single and multiple scattering. *ACM Trans. Graph.* 24, 3, 1202–1207.
- WEI, X., 2006. What is human hair? a light and scanning electron microscopy study.
- XU, K., MA, L.-Q., REN, B., WANG, R., AND HU, S.-M. 2011. Interactive hair rendering and appearance editing under environment lighting. *ACM Transactions on Graphics* 30, 6, 173:1–173:10.
- YUKSEL, C., AND KEYSER, J. 2008. Deep opacity maps. *Computer Graphics Forum* 27, 2, 675–680.
- ZINKE, A., AND WEBER, A. 2007. Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics* 13, 2, 342–356.
- ZINKE, A., YUKSEL, C., WEBER, A., AND KEYSER, J. 2008. Dual scattering approximation for fast multiple scattering in hair. *ACM Trans. Graph.* 27, 3, 32:1–32:10.
- ZINKE, A., RUMP, M., LAY, T., WEBER, A., ANDRIYENKO, A., AND KLEIN, R. 2009. A practical approach for photometric acquisition of hair color. *ACM Trans. Graph.* 28, 5, 165:1–165:9.

Appendix

Precomputation and Compression: We enumerate scattering coefficient $\sigma_{m,s} \in [0, 20]$ and anisotropy $g \in [0.0, 0.8]$, and we vary an additional parameter specifying different incoming directions. For azimuthal profiles, it is the offset $h' \in [-1, 1]$ assuming all sub-paths are entering the medulla horizontally (Fig. 14). For longitudinal profiles, it is the incident angle $\theta'_i \in [-\pi/2, \pi/2]$. We discretize the range of g into 16 steps, while for all other parameters, we use 64 steps.

We separately simulate the medulla's azimuthal and longitudinal scattering profiles using volumetric path tracing in 2D, assuming that the incident path carries unit energy. We trace a smooth unit circle azimuthally and a double slab with distance 2 in between longitudinally. For all scattering events, we use the planar Henyey-Greenstein phase function [Davis 2006]

$$\rho(\theta, g) = \left(\frac{1}{2\pi} \right) \frac{1 - g^2}{1 + g^2 - 2g \cos \theta} \quad (19)$$

to consider anisotropy, where θ is the angle each scattering event deviates from its earlier path.

Each precomputed outgoing profile is stored using 720 bins covering every direction, recording the exiting energy. For longitudinal scattering, we further normalize the upper and lower lobes respectively, making both of them PDFs. Then we accumulate their CDFs for convenience in the next steps. Since these profiles are generally smooth, we further compress each profile by dividing it into 4 segments, and fit each segment with a quadratic function.

Refining azimuthal scattered lobes: Due to Fresnel effects, the scattered lobe could be reflected back by the surface of the medulla and the cuticle, thus undergoing further scattering, making it even smoother. We approximate this proportion as $\lambda = 1 - (1 - F(\pi/2))(1 - F(\pi/2, l))$, which is simply the leftover energy after perpendicularly transmitting through two interfaces, ignoring multiple internal reflections. So the azimuthal scattered lobes can be extended to

$$N_p^s(h, \phi) = A_p^s(h) \cdot [D_p^s(h, \phi) \cdot (1 - \lambda) + I_j^s \cdot \lambda], \quad (20)$$

where $I_j^s = [1 - \exp(-\sigma_{m,s} \cdot |p_j p_{j+1}|)]/2\pi$ is uniform distribution of the scattered energy, under the assumption that the reflected back scattered lobe will become isotropic after more scattering events. j is the location of the vertex where the path enters the medulla and scatters, and A^s is the effective attenuation for the scattered lobe.

Normalizing longitudinal scattered lobe: Since our final longitudinal scattered lobe M^s is normalized, we have $\int M^s d\theta_r = 1$. According to equation 18, we have $\int F_t \cdot C^M d\theta_r = 1/\mu$. Then, approximately we have $\int F_t d\theta_r \cdot \int C^M d\theta_r = 1/\mu$. Denoting these two separated integrals as U and V , we found that U can be numerically calculated when a certain species of fur is loaded in, which introduces no overhead when rendering on the fly. For V , we first convert the integration domain into θ'_r , using $\frac{d\theta'_r}{d\theta_r} \approx \frac{\cos \theta_r}{\eta_c \cos \theta'_r}$ by differentiating the formula for Snell's law at P' and Q' . Then we know the range of θ'_r that C^M should be integrated in. Since the precomputed C^M is a normalized PDF, the integration is a query on its CDF, which is simply accumulated in the precomputation step. Thus, we have $\mu = (UV)^{-1}$ as our normalization factor.