

Azimuthal Scattering from Elliptical Hair Fibers

PRAMOOK KHUNGURN and STEVE MARSCHNER

Cornell University

The appearance of hair follows from the small-scale geometry of hair fibers, with the cross-sectional shape determining the azimuthal distribution of scattered light. Although previous research has described some of the effects of non-circular cross sections, no accurate scattering models for non-circular fibers exist. This article presents a scattering model for elliptical fibers, which predicts that even small deviations from circularity produce important changes in the scattering distribution and which disagrees with previous approximations for the effects of eccentricity.

To confirm the model's predictions, new scattering measurements of fibers from a wide range of hair types were made, using a new measurement device that provides a more complete and detailed picture of the light scattered by fibers than was previously possible. The measurements show features that conclusively match the model's predictions, but they also contain an ideal-specular forward-scattering behavior that is not predicted and has not been fully described before.

The results of this article indicate that an accurate and efficient method for computing scattering in elliptical cylinders—something not provided in this article—is the correct model to use for realistic hair in the future and that the new specular behavior should be included as well.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**; • **General and reference** → *Measurement*;

Additional Key Words and Phrases: Hair, fibers, optical scattering, rendering

ACM Reference Format:

Pramook Khungurn and Steve Marschner. 2017. Azimuthal Scattering from Elliptical Hair Fibers. *ACM Trans. Graph.* 36, 2, Article 13 (March 2017), 23 pages.

DOI: <http://dx.doi.org/10.1145/2998578>

This work was supported in part by the National Science Foundation (grants IIS-1513967 and IIS-1011919) and the Intel Science and Technology Center for Visual Computing. It made use of the Cornell Center for Materials Research Facilities supported by the National Science Foundation under Award Number DMR-1120296.

Authors' addresses: P. Khungurn (corresponding author), and S. Marschner, Department of Computer Science, Cornell University, Ithaca, NY; emails: {pramook, srm}@cs.cornell.edu.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2017 ACM 0730-0301/2017/03-ART13 \$15.00

DOI: <http://dx.doi.org/10.1145/2998578>

1. INTRODUCTION

The appearance of hair arises from the interaction of light with the geometry of fibers. The overall cylindrical shape of fibers causes the characteristic linear highlights on a head of hair, and the scaly surface of the cuticle covering the hair separates them into distinct colored highlights. The fibers' particular cross-sectional shape controls how scattered light depends on both the locations of the viewer and light source and the orientation of the fiber, determining whether highlights appear smooth or glittery. Scattering models that are used to render hair account for all three phenomena, but research has mainly focused on the first two. This article is about how to improve our models for the effects of cross section on the distribution of scattered light.

It is well documented that hair cross sections are generally non-circular and often roughly elliptical [Robbins 1994]. However, prior scattering models for hair are all based on an analysis of light rays interacting with circular cylinders, including the model of Marschner et al. [2003] that includes an approximate correction for eccentricity.

Creating efficient scattering models that properly account for elliptical cross sections is a difficult undertaking, and this article seeks to answer two questions:

- Is cross section *important*? Are the changes in appearance caused by non-circular cross sections significant enough to warrant the effort?
- Are ellipses a *correct* model? Does an elliptical cross section predict what happens in real hair fibers?

To answer the question of importance, we have extended previous hair models to account for elliptical cross sections. By contrast to Marschner et al.'s previous model, our model accurately describes the geometric optics of elliptical cylinders. It predicts dramatic changes in scattering behavior, particularly for transmitted and internally reflected light, as we move from circular to elliptical fibers. These changes are significant even for mild eccentricity, and the behavior differs considerably from the eccentricity correction proposed by Marschner et al.

These predictions of our elliptical theory need to be substantiated: After all, hairs are not exactly elliptical, and many details of their structure and of the physics of light are omitted by the model. The elliptical model might predict features that do not occur because of imperfections in real fibers, or its inherent assumptions might cause it to miss visually important phenomena.

The experiments of Bustard and Smith [1991] and later Marschner et al. provide promising support for an elliptical model, but their measurements were limited, being based on small numbers of data points gathered on just a few fibers. No one has examined the complete scattering function in enough detail to see whether we have modeled the whole behavior.

Therefore, to answer the question of correctness, we have made new measurements of a range of different types of hair fibers, using a new device designed to give a very complete picture of fiber scattering functions. Our goal is to capture the complete function with enough resolution to observe all the important features, but accuracy in absolute magnitudes is a secondary concern. With this



Fig. 1. Orientation-dependent glints in real hair. Here we show two photographs of Caucasian hair taken from two different angles under the same lighting: (a) with the sun behind the camera and (b) with the sun roughly 90° to the camera's left. Bright glints can be observed in panel (b) but not in panel (a).

goal in mind, the measurement system is built around image-based measurements, capturing tens of thousands of data points with each exposure.

The conclusions of our study are clear, answering both questions in the affirmative and introducing a new phenomenon:

- The structured internal-reflection glints predicted by the elliptical model dramatically contradict previous models. They also are clearly observable in real hair, as long as it is light enough in color for internal reflection to be significant.
- There is a forward-scattering feature present in all hair scattering functions, previously unremarked on in graphics articles: an ideal-specular lobe, which we call the *E mode*, that becomes particularly bright at grazing angles.

For rendering, the key implication of non-circular fibers is that scattering depends jointly on the azimuthal positions of the light and viewer, not only on the azimuthal *difference* as in the circular case. The azimuthal component of a scattering model must be a function of two variables, not one, and the brightness of fibers depends on their orientation, even when the light and viewer remain fixed. This leads to a strong glinty appearance in hair under conditions when this orientation dependence is strong (Figure 1).

It should be emphasized that this article does not propose a production-ready, practical algorithm to efficiently compute scattering from elliptical fibers in a shader. The renderings in this article are made using tabulated scattering functions, which limits the ability to model continuous variation in fiber properties. Our study has established what should be modeled to provide a truly high-fidelity hair scattering model. Models that are more efficient and include the *E mode* should be developed in the future.

In the remainder of this article, we extend previous models to accurately account for cross section (Sections 3 and 4); examine the predictions of the theory (Section 5); describe our measurement device and its results (Sections 6 and 7); and examine the agreement between theory and measurements (Sections 8 and 9).

2. PREVIOUS WORK

Kajiya and Kay [1989] presented the first analytical scattering model for hair fibers, based on the Phong surface shading model. By explicitly modeling fibers as cylinders, it captures the linear highlights characteristic of hair and fur, but since the model assumes hair fibers are opaque, it cannot model effects arising from refraction. Extensions to incorporate forward scattering into the Kajiya-Kay model have been proposed [Goldman 1997; Kim 2002].

Marschner et al. [2003] were the first to model a hair fiber as a transparent dielectric cylinder. Their model decomposes the scattering behavior of fibers into scattering modes according to the number of reflections and refractions that a light ray undergoes as it interacts with the hair fiber. They factored the scattering function for each mode into a product of a longitudinal term and an azimuthal term. The former enabled them to model the longitudinal shifts of highlights, and the latter the glints due to the fiber's cross section.

Marschner et al.'s model has been widely used and has generated much subsequent research. Zinke and Weber [2007] adapted the model to a more general mathematical framework that allowed rendering of close-up views. Zinke et al. [2009] added a diffuse component to the Marschner model and provided a procedure for acquiring model parameters from a single photograph of a hair sample. Sadeghi et al. [2010] refined the model to be more intuitive for artists. Xu et al. [2011] gave a fast approximation to the model for interactive rendering and appearance editing. Some research has also gone toward analytic importance sampling for variants of this model [Hery and Ramamoorthi 2012; Ou et al. 2012]; in our implementation we use a tabulated approach.

The Marschner model, however, has two important shortcomings. First, it is not energy conserving, and, second, it deals with the caustics in one scattering mode in an *ad hoc* way, which introduces extra non-physical parameters to the model. d'Eon et al. [2011] addressed both problems by introducing a new longitudinal component designed to conserve energy and an azimuthal component based on numerical integration. Our work fixes an energy conservation problem and generalizes their model to arbitrary cross sections.

While the above hair scattering models all assume hair fibers have circular cross sections, Ogaki et al. [2010] modeled a scattering function arising from arbitrary user-specified cross-sectional shapes that can even vary along the length of the fiber. To deal with such geometry, they trace photons through the fiber and tabulate the resulting density estimates. Ogaki et al.'s modeling technique is much more general than ours, but its aim was to model animal furs that have a wide variety of scale structures. On the other hand, our work focuses on human hair and is more experimental in nature. The concurrent work of Yan et al. [2015] models a fur strand as a circular cylinder having two layers with different absorbing materials, reflecting more accurately animal fur's anatomy. Our work assumes a single material throughout the strand and focuses on the effect of cross-sectional shape.

Researchers in the optics community have characterized some aspects of elliptical fibers' light scattering behavior. Marcuse determined the range of azimuthal angles of backward scattered light [Marcuse 1974]. Adler et al. studied the Fourier transform of the angle of backward scattered caustics as a function of incoming angle [Adler et al. 1998]. In comparison, our work is more comprehensive because we characterize both forward and backward scattering lobes.

We are only aware of three other works in the graphics community that measure light scattered from fiberlike structures. The first is Marschner et al. [2003], the second is Sadeghi et al. [2013], and the third is Yan et al. [2015]. All three works used a gonioreflectometer for the task, and so only 1D or coarse 2D slices of the 4D scattering function were feasible to acquire. Our device, on the other hand, can acquire the whole hemisphere of scattered light in one exposure, thus yielding more comprehensive data at greater efficiency than the previous method. However, the measurements are less accurate due to factors we shall describe in Section 6. As such, it is more suitable for observing qualitative behavior of the scattering function than for directly acquiring the function itself.




Fig. 2. Projected diameter $D(\phi)$ and displacement parameter s are quantities that are essential to definition of fiber scattering model that can take into account elliptical cross sections.

3. BACKGROUND

In this section, we review the theory of light scattering from fibers and extend it to handle fibers with arbitrary cross sections. Our treatment is similar to that of Zinke [2008] but simpler. We also review parameters common to previous models and introduce two types of visualization of fiber scattering models that will be used in future sections.

3.1 Angles

We follow the notation of Marschner et al. [2003]. All calculations are done in a coordinate system with right-handed orthonormal basis vectors \mathbf{u} , \mathbf{v} , and \mathbf{w} , and the axis of symmetry of the hair fiber runs along the \mathbf{u} axis. The direction $-\mathbf{u}$ points towards the root, and \mathbf{u} towards the tip. A 3D direction vector ω can be described by two spherical angles: the *longitudinal angle* $\theta \in [-\pi/2, \pi/2]$ and the *azimuthal angle* $\phi \in [0, 2\pi]$:

$$\omega = \begin{bmatrix} \sin \theta \\ \cos \theta \cos \phi \\ \cos \theta \sin \phi \end{bmatrix}.$$

We let ω_i and ω_o denote the incoming and outgoing direction, respectively. The spherical angles of the two directions are denoted by θ_i , ϕ_i , θ_o , and ϕ_o .

3.2 Projected Diameter

We assume that the fiber is a generalized cylinder whose cross section (which might not be circular) is constant along its length. The *projected diameter along azimuthal angle* ϕ is the width of the fiber when the observer views the fiber from azimuthal direction ϕ . More precisely, it is the furthest distance between two lines in the normal plane that are tangent to the cross section and make angle ϕ with the v -axis. (See Figure 2.) We denote the projected diameter by the symbol $D(\phi)$. For example, for a circular cross section of radius 1, $D(\phi_i) = 2$ for all ϕ_i . For an elliptic cross section with minor radius 1 and major radius a , $D(0^\circ) = D(180^\circ) = 2$ and $D(90^\circ) = D(270^\circ) = 2a$.

3.3 Ray Spaces

Define the *fiber ray space* (FRS) as the set of all rays exiting the surface of the fiber (that is, the dot product of the ray's direction with the surface normal at that point is non-negative). Each ray in the FRS can be parameterized by four values: the longitudinal angle θ , the azimuthal angle ϕ , the displacement along the fiber's axis, t , and the displacement along the projected diameter, s , which ranges from $-D(\phi)/2$ to $D(\phi)/2$. We will indicate functions defined on the FRS by a superscript F .

A related space of rays, useful when we wish to ignore the width of the fiber, is the set of rays originating from a specific point along the fiber's axis. We call this space the *conventional ray space* (CRS). A ray in CRS is parameterized by one direction $\omega = (\theta, \phi)$ and the displacement along the fiber's axis, t . We will indicate functions defined on the CRS by a superscript C .

3.4 Radiometry

Radiance incident to or exitant from the fiber's surface is a function of the rays in FRS. Hence, we may write it as $L^F(\theta, \phi, s, t)$ or $L^F(\omega, s, t)$.

Marschner et al. defined two new radiometric quantities: the *curve radiance* \bar{L} and the *curve irradiance* \bar{E} . The curve radiance is power per unit solid angle per unit projected length of the fiber. The curve radiance is a function of ω and t , so we may write $\bar{L}(\omega, t)$. It is related to the radiance in FRS by

$$\bar{L}(\omega, t) = \int_{-D(\phi)/2}^{D(\phi)/2} L^F(\omega, s, t) ds.$$

The curve irradiance is incoming power per unit length of the fiber, which is a function of t and thus denoted $\bar{E}(t)$. It is equal to the integral of the incoming curve radiance scaled by the cosine of the longitudinal angle:

$$\bar{E}(t) = \int_{\mathbb{S}^2} \bar{L}(\omega, t) \cos \theta d\omega = \int_{\mathbb{S}^2} \int_{-D(\phi)/2}^{D(\phi)/2} L^F(\omega, s, t) \cos \theta ds d\omega.$$

The cosine factor is present as a consequence of Lambert's cosine law. Note that the quantity $\bar{L}(\omega, t) \cos \theta d\omega$ can be interpreted as the infinitesimal irradiance created by rays whose directions are within $d\omega$ of ω .

Curve radiance can be related to ordinary radiance by the following *thin fiber assumption*:

$$L^F(\omega, s, t) = L(\omega, 0, t) = L(\omega, t)$$

for all s . That is, we assume the fiber is so thin that the radiance distribution is constant across its width. This means all rays have the same radiance as the ray in the same direction that passes through the fiber axis, so $L(\omega, t)$, as well as $\bar{L}(\omega, t)$ is a function of rays in the CRS. The assumption implies that:

$$\bar{L}(\omega, t) = D(\phi)L(\omega, t).$$

In all the discussion that follows, we will make this assumption. For brevity, we will drop the t parameter from all the radiometric functions, leaving the dependence on t implicit.

3.5 Scattering Model

Assume that the fiber is struck by incoming curve radiance $\bar{L}_i(\omega_i) = D(\phi_i)L_i(\omega_i)$ from a differential solid angle $d\omega_i$ around the direction ω_i , which produces differential curve irradiance $\bar{L}_i(\omega_i) \cos \theta_i d\omega_i$. The fiber scatters this incoming irradiance into a distribution of

differential curve radiance over the outgoing directions, and let us denote its value at direction ω_o by $d\bar{L}_o(\omega_i \rightarrow \omega_o)$. The outgoing curve radiance along a specific direction ω_o , as a result of the whole incoming curve radiance distribution, is given by

$$\bar{L}_o(\omega_o) = \int_{\omega_i \in \mathbb{S}^2} d\bar{L}_o(\omega_i \rightarrow \omega_o).$$

The *bidirectional curve scattering distribution function* (BCSDF), denoted by $S(\omega_i, \omega_o)$ or $S(\theta_i, \phi_i, \theta_o, \phi_o)$ is the ratio of the outgoing differential curve radiance in direction ω_o to the incoming differential curve irradiance from direction ω_i :

$$S(\omega_i, \omega_o) = \frac{d\bar{L}_o(\omega_i \rightarrow \omega_o)}{\bar{L}_i(\omega_i) \cos \theta_i d\omega_i}.$$

This function and the thin fiber assumption imply the following relationship between radiance scattered from hair and the incoming radiance:

$$\begin{aligned} L_o(\omega_o) &= \frac{\bar{L}_o(\omega_o)}{D(\phi_o)} = \frac{1}{D(\phi_o)} \int_{\omega_i \in \mathbb{S}^2} d\bar{L}_o(\omega_i \rightarrow \omega_o) \\ &= \frac{1}{D(\phi_o)} \int_{\mathbb{S}^2} S(\omega_i, \omega_o) \bar{L}_i(\omega_i) \cos \theta_i d\omega_i \\ &= \frac{1}{D(\phi_o)} \int_{\mathbb{S}^2} S(\omega_i, \omega_o) D(\phi_i) L_i(\omega_i) \cos \theta_i d\omega_i, \end{aligned}$$

or, equivalently,

$$L_o(\omega_o) = \int_{-\pi/2}^{\pi/2} \int_0^{2\pi} S(\omega_i, \omega_o) \frac{D(\phi_i)}{D(\phi_o)} L_i(\omega_i) \cos^2 \theta_i d\phi_i d\theta_i, \quad (1)$$

because $d\omega_i = \cos \theta_i d\phi_i d\theta_i$.

A BCSDF is said to be *energy conserving* if

$$\int_{\mathbb{S}^2} S(\omega_i, \omega_o) \cos \theta_o d\omega_o \leq 1$$

for all ω_i . The inequality above implies that the outgoing curve irradiance produced by the BCSDF does not exceed the curve irradiance coming in.

3.6 Factored Scattering Models

Marschner et al. classify scattered light into several “modes”: light that reflects out at first contact (R mode); light that transmits into the cross section and then transmits out when it hits the surface the second time (TT mode); light that transmits, reflects internally, and then transmits out (TRT mode); light that transmits, reflects internally twice, and then transmits out (TRRT mode); and so on. For brevity, we denote each mode by a non-negative integer p , where $p = 0$ means the R mode, $p = 1$ means the TT mode, $p = 2$ means the TRT mode, and so on.

Accordingly, the BCSDF is decomposed into a number of *per-mode scattering functions*:

$$S(\omega_i, \omega_o) = \sum_{p=0}^{\infty} S_p(\omega_i, \omega_o).$$

Each per-mode scattering function is factored into a product of two terms: the *longitudinal scattering function* (LSF) M_p and the *azimuthal scattering function* (ASF) N_p ,

$$S_p(\omega_i, \omega_o) = M_p(\theta_i, \theta_o) N_p(\theta_d, \phi_i, \phi_o), \quad (2)$$

where, depending on the scattering model, θ_d can be equal to the incoming angle, θ_i , or the half angle, $|\theta_i - \theta_o|/2$. The Marschner

model [2003] and the d'Eon model [2011] use the half angle, but the model that will be presented in this article uses the incoming angle.

We say that the whole collection of ASFs is *energy conserving* if, for all θ_d and ϕ_i ,

$$\int_0^{2\pi} \sum_{p=0}^{\infty} N_p(\theta_d, \phi_i, \phi_o) d\phi_o \leq 1.$$

That is, the incoming energy from a single direction should be split among the modes, and no extra energy should be introduced. We say that an ASF $N_p(\theta_d, \phi_i, \phi_o)$ is *reciprocal* if

$$\frac{N_p(\theta_d, \phi_i, \phi_o)}{D(\phi_o)} = \frac{N_p(\theta_d, \phi_o, \phi_i)}{D(\phi_i)} \quad (3)$$

for all θ_d , ϕ_i , and ϕ_o . We choose this definition because it simplifies the outgoing radiance integral in Equation (1):

$$\begin{aligned} \int_{-\pi/2}^{\pi/2} \int_0^{2\pi} \sum_{p=0}^{\infty} M_p(\theta_i, \theta_o) N_p(\theta_d, \phi_i, \phi_o) \frac{D(\phi_i)}{D(\phi_o)} L_i(\omega_i) \cos^2 \theta_i d\phi_i d\theta_i \\ = \int_{-\pi/2}^{\pi/2} \int_0^{2\pi} \sum_{p=0}^{\infty} M_p(\theta_i, \theta_o) N_p(\theta_d, \phi_o, \phi_i) L_i(\omega_i) \cos^2 \theta_i d\phi_i d\theta_i. \end{aligned}$$

We say that an LSF is *energy conserving* if

$$\int_{-\pi/2}^{\pi/2} M_p(\theta_i, \theta_o) \cos^2 \theta_o d\theta_o \leq 1 \quad (4)$$

for all θ_i . This implies that the LSF by itself does not introduce any extra energy. For example, if the ASF is diffuse and energy preserving (that is, if there is only one mode, and $N_p(\theta_d, \phi_i, \phi_o) = 1/(2\pi)$ for all values of θ_d , ϕ_i , and ϕ_o), then the per-mode BCSDF is energy conserving. Moreover, provided that θ_d is equal to θ_i instead of the half angle, we can *prove* that the BCSDF is energy conserving if both the ASFs and the LSF are energy conserving. (The proof is given in the supplementary material.) The full scattering model we shall construct in later sections satisfies all the definitions above and therefore is provably energy conserving.

In this factored scattering model framework, Marschner et al. propose an LSF based on a Gaussian distribution and ASFs up to the TRT mode. While the ASFs for R and TT modes are modeled using circular fibers, the one for the TRT mode features an ad hoc approximation to elliptical cross sections. d'Eon et al. note that Marschner et al.'s LSF is not energy conserving according to Equation (4) and provide a new LSF that satisfies the definition. Using the concept of surface roughness first proposed by Zinke and Weber [2007], they also formulate ASFs based on circular fibers that can be evaluated up to any arbitrary scattering mode. However, because their ASFs use the half angle as θ_d , the resulting BCSDF cannot be proven to be energy conserving.

3.7 Common Model Parameters

Previous hair scattering models [Marschner et al. 2003; d'Eon et al. 2011] and the model we will present share the following parameters:

- The index of refraction of the material η .
- The absorption coefficient of the material σ .
- For each mode p , the longitudinal shift α_p .
- For each mode p , the width of the longitudinal lobe β_p .
- The aspect ratio of the cross section a . If the cross section is an ellipse, then the aspect ratio is the ratio of the length of the major axis to the length of the minor axis. Hence, the circle has $a = 1$.

Note that the Marschner model only covers scattering modes up to $p = 2$. It also contains extra parameters such as the glint scale factor k_G and the azimuthal caustic width w_c in order to support its ad hoc modeling of the TRT mode. The d'Eon model requires that β_p is the same for all p and calls the single β parameter the “surface roughness” of the hair fiber. The d'Eon model also only models circular hair fibers, so a is always 1.

3.8 Visualization

We introduce two types of visualization for parts of the scattering models and measured data.

The first, the (ϕ_i, ϕ_o) -plot, shows $N_p(\theta_d, \phi_i, \phi_o)$ for a fixed θ_d as a scalar field over the (ϕ_i, ϕ_o) -plane. Each point's color indicates the intensity of light scattered to ϕ_o when it arrives from ϕ_i . The (ϕ_i, ϕ_o) -plot allows us to see the complete picture of how the ASF evolves with ϕ_i , and it exposes the inherent symmetry due to reciprocity. It also provides a convenient way of observing how the ASF evolves with changing θ_i . Examples of (ϕ_i, ϕ_o) -plots are given in Figure 3.

The second, the (θ_o, ϕ_d) -plot, shows a slice of the whole scattering function $S(\omega_i, \omega_o)$ for a fixed incoming direction ω_i . The slice is shown as a color map over the (θ_o, ϕ_d) -plane where $\phi_d = \phi_o - \phi_i$. Note that by using the difference angle ϕ_d instead of the outgoing angle ϕ_o , the direction of retroreflection is fixed at $\phi_d = 0^\circ$ and the direction of forward scattering is fixed at $\phi_d = 180^\circ$. In this way, the (θ_o, ϕ_d) -plot does not shift vertically as ϕ_i changes, making it easier to notice changes in the scattering function as ϕ_i changes. Also, as we will see in the next section, it represents the data that are captured by one snapshot of our measurement system. As examples, (θ_o, ϕ_d) -plots of the Marschner and d'Eon models are given in Figure 4.

Inspection of the (ϕ_i, ϕ_o) -plots in Figure 3 provides us with heuristics on how to recognize the first three scattering modes in the (θ_o, ϕ_d) -plots:

- The R mode has a large support in the (ϕ_i, ϕ_o) -plots, covering all but a small area around the lines $\phi_o = \phi_i \pm 180^\circ$. Hence, in (θ_o, ϕ_d) -plots, it should appear as two vertical strips covering the whole range of ϕ_d except a small portion around $\phi_d = 180^\circ$.
- The TT mode has narrower support around $\phi_o = \phi_i \pm 180^\circ$ and is the brightest among the three modes. Hence, it should appear as the brightest lobe, centered vertically around $\phi_d = 180^\circ$.
- The TRT mode has the narrowest support around the line $\phi_o = \phi_i$. Hence, it should appear as blobs located vertically near $\phi_d = 0^\circ$ and $\phi_d = 360^\circ$.

According to Marschner et al. [2003], the modes are also separated in θ due to the tilting of the hair scales. Note that, at the beginning of Section 3, we said that the direction $-\mathbf{u}$ points towards the root and \mathbf{u} points towards the tip. This means that the left side of the (θ_o, ϕ_d) -plot (-90°) should correspond to scattering toward the root, and the right side (90°) should correspond to scattering toward the tip. Marschner et al. observe that the R mode is shifted towards the root (i.e., to the left), but other modes are shifted towards the tip (i.e., to the right). As such, the R mode should appear to the left of the ideal specular line ($\theta_o = -\theta_i$) while other modes appear to the right as can be seen in Figure 4.

4. MODELING

Having reviewed the previous models and establishing our notations, in this section we construct a scattering model for hair fiber with elliptical cross sections. Our model is a factored model like




Fig. 3. The (ϕ_i, ϕ_o) -plots of the first three modes of the Marschner and d'Eon scattering models. Notice that the plots are *shift invariant*, meaning that each vertical slice of the plot is a shifted version of other slices. This property arises from the rotational symmetry of the circular cross section. All observable features in shift-invariant plots are slanted bars along the southwest–northeast direction. One can also see that the d'Eon model is very similar to the Marschner model for the circular fiber. (The relevant parameters are $\eta = 1.55$, $\sigma = 0$, $\beta_0 = \beta_1 = \beta_2 = 5^\circ$, and $a = 1$ (circle). For the Marschner model, we set $k_G = 1$, $w_c = 10^\circ$, $\Delta\eta' = 0.2$, and $\Delta h_M = 0.5$. For all plots, $\theta_d = 0$. The colors between different modes of the same models are in the same scale, but the plots are not in the same scale between the models.)

that of Marschner et al.:

$$S(\omega_i, \omega_o) = \sum_{p=0}^{\infty} M_p(\theta_i, \theta_o) N_p(\theta_i, \phi_i, \phi_o). \quad (5)$$

However, we make the ASF depend on θ_i rather than the half angle $|\theta_i - \theta_o|/2$ for the reason we discussed in Section 3.6.

4.1 Longitudinal Scattering Function

The LSF for Mode p is defined by two parameters: the *width of the lobe* β_p and the *longitudinal shift* α_p . In the original model of Marschner et al., the LSF was simply a Gaussian distribution in θ . As observed by d'Eon et al. [2011] and other researchers, this LSF is not normalized on the sphere and loses energy away from normal incidence. Although d'Eon et al. proposed an elegant and improved alternative, we have noticed some problems that we




Fig. 4. The (θ_o, ϕ_d) -plots of the Marschner and d'Eon scattering models. The R mode manifests as two vertical strips, the TT mode as the brightest lobe at the vertical center of the plot, and the TRT mode as blobs near the top and bottom ends. (The longitudinal tilts are given by $\alpha_0 = -10^\circ$, $\alpha_1 = 5^\circ$, and $\alpha_2 = 15^\circ$, while other model parameters are the same as those in Figure 3. For both plots, $\theta_i = 0^\circ$ and $\phi_i = 0^\circ$. The plots are not to scale with each other.)

discuss in the supplementary material. In this article, we use a simpler LSF that was presented in Khungurn et al. [2015], which we reproduce here for completeness. The function is a normalized version of the original Gaussian,

$$M_p(\theta_i, \theta_o) = \frac{g(\theta_o; \mu(\theta_i), \beta_p)}{G(\mu(\theta_i), \beta_p)},$$

where $g(x; \mu, \sigma) = \exp(-(x - \mu)^2 / 2\sigma^2) / (\sqrt{2\pi}\sigma)$ is the Gaussian distribution, and the normalization factor G is calculated to ensure normalization over the sphere. In order to achieve a lobe shifted by α_p without the lobe becoming excessively concentrated near grazing, we center the Gaussian at

$$\mu(\theta_i) = \text{clamp}(-\theta_i + \alpha_p, [-\pi/2, \pi/2]).$$

The normalization factor G is defined to ensure cosine-weighted normalization over the sphere:

$$G(\mu, \sigma) = \int_{-\pi/2}^{\pi/2} g(\theta; \mu, \sigma) Q(\theta) d\theta \approx \int_{-\pi/2}^{\pi/2} g(\theta; \mu, \sigma) \cos^2(\theta) d\theta,$$

where Q is a polynomial that approximates $\cos^2 \theta$ from above¹:

$$\begin{aligned} Q(\theta) &= 0.002439\theta^8 - 0.04301\theta^6 + 0.3322\theta^4 \\ &\quad - 0.999745\theta^2 + 1.0001 \\ &\geq \cos^2 \theta. \end{aligned}$$

It can be easily shown that this normalization factor implies the LSF being energy conserving in the sense of Equation (4). The indefinite integral of the product of a Gaussian and a polynomial has a closed form, so the integral can be evaluated analytically. (See the appendix.)

We note that Zinke [2008] also defines an LSF that is a scaled Gaussian lobe. The difference between his function and ours is that

¹Alternatively, one can also define G more directly in terms of the antiderivative of $g(\theta; \mu, \sigma) \cos^2 \theta$. However, we found that its closed form involves the imaginary error function, a special function that is not a part of many libraries. We approximated $\cos^2 \theta$ with a polynomial because the antiderivative depends only the real error function, whose implementation is more readily available.

his normalization factor $G(\mu, \sigma)$ does not contain the polynomial Q in the integrand and so is much larger than ours, implying that his function loses more energy.

4.2 Azimuthal Scattering Function

We derive the ASF by tracing rays through the elliptical cross section. Due to Bravais's Law [Tricker 1970; Marschner et al. 2003], the paths of reflected and refracted rays in a cylinder of any cross section behave the same as rays in 2D interacting with the cross section, though with a different refractive index that depends on θ_i . In this 2D world, the fiber ray space is parameterized by an angle ϕ and a displacement s , which still ranges from $-D(\phi)/2$ to $D(\phi)/2$. So, radiance is now written as $L(\phi, s)$. The curve radiance, now parameterized only by ϕ , is defined in a similar way as that of the 3D version:

$$\bar{L}(\phi) = \int_{-D(\phi)/2}^{D(\phi)/2} L(\phi, s) ds. \quad (6)$$

4.2.1 Characterizing a Cross Section. We assume that each interaction of light with the surface results in perfect specular reflection or refraction; we will account for surface roughness later. For our purpose, a cross section is characterized by the following:

- the projected diameter $D(\phi)$, defined in Section 3;
- for each mode p , the *exit direction* $\phi_p^e(\phi, s)$;
- for each mode p , the *exit displacement* $s_p^e(\phi, s)$; and
- for each mode p , the *attenuation* $A_p(\phi, s)$;

where ϕ_p^e and s_p^e describe the direction and displacement of a ray entering at (ϕ, s) after it has undergone $p + 1$ interactions with the cross section's surface. $A_p(\phi, s)$ is the attenuation of the radiance along the ray as it interacts with the surface and the material inside the fiber. Reversibility of paths implies that if $\phi_o = \phi_p^e(\phi_i, s_i)$ and $s_o = s_p^e(\phi_i, s_i)$, then

$$\phi_i = \phi_p^e(\phi_o, s_o), s_i = s_p^e(\phi_o, s_o), \text{ and } A_p(\phi_i, s_i) = A_p(\phi_o, s_o).$$

The relationship between exiting and incident radiance in mode p is then

$$L_{o,p}(\phi, s) = A_p(\phi, s) L_{i,p}(\phi_p^e(\phi, s), s_p^e(\phi, s)). \quad (7)$$

Here $L_{i,p}$ and $L_{o,p}$ are the incoming and outgoing radiance, respectively, restricted to mode p . (Unless otherwise stated, the subscript p denotes restriction to mode p from now on.)

4.2.2 Defining the ASF. Having defined the relationship between $L_{i,p}$ and $L_{o,p}$, we want to derive a suitable relationship between $\bar{L}_{i,p}$ and $\bar{L}_{o,p}$, which is the ASF. Substituting Equation (7) into Equation (6) and assuming $L_{i,p}(\phi, s) \equiv \bar{L}_{i,p}(\phi) / D(\phi)$ gives

$$\bar{L}_{o,p}(\phi_o) = \int_{-D(\phi_o)/2}^{D(\phi_o)/2} A_p(\phi_o, s_o) \frac{\bar{L}_{i,p}(\phi_p^e(\phi_o, s_o))}{D(\phi_p^e(\phi_o, s_o))} ds_o.$$

To define a BCSDf requires a change of variables from s_o to ϕ_i . As observed in prior work, this approach has the problem that the change-of-variables factor $|ds_o/d\phi_i|$ is not always finite. In the model of Marschner et al., this was worked around by explicitly removing the singularities, but for non-circular cross sections it is not simple to compute where the singularities occur. On the other hand, surface roughness prevents these singularities from appearing in practice, so we follow the approach pioneered by Zinke and Weber [2007] and d'Eon et al. [2011] and apply a blur to account for roughness before changing variable.

We will make use of an *angular kernel function* $K^{\gamma_p}(\phi)$ defined with the per-mode width parameter γ_p on the interval $[0, 2\pi)$. The angular kernel function obeys the normalization $\int_0^{2\pi} K^{\gamma_p}(\phi) d\phi = 1$. To maintain symmetry, we convolve $\bar{L}_{i,p}(\phi_i)$ and $\bar{L}_{o,p}(\phi_o)$ each with K^{γ_p} :

$$\bar{L}_{o,p}(\phi_o) = \left[K^{\gamma_p} * \int_{-D(\cdot)/2}^{D(\cdot)/2} A_p(\cdot, s_o) \left[K^{\gamma_p} * \frac{\bar{L}_{i,p}}{D} \right] (\phi_p^e(\cdot, s_o)) ds_o \right] (\phi_o).$$

This equation can be expanded and written as an integral over ϕ_i as follows:

$$\begin{aligned} \bar{L}_{o,p}(\phi_o) &= \int_0^{2\pi} K^{\gamma_p}(\phi_o - \phi'_o) \left[\int_{-D(\phi'_o)/2}^{D(\phi'_o)/2} A_p(\phi'_o, s_o) \right. \\ &\quad \left. \left(\int_0^{2\pi} K^{\gamma_p}(\phi_p^e(\phi'_o, s_o) - \phi_i) \frac{\bar{L}_{i,p}(\phi_i)}{D(\phi_i)} d\phi_i \right) ds_o \right] d\phi'_o \\ &= \int_0^{2\pi} \frac{\bar{L}_{i,p}(\phi_i)}{D(\phi_i)} \left[\int_0^{2\pi} K^{\gamma_p}(\phi_o - \phi'_o) \right. \\ &\quad \left. \left(\int_{-D(\phi'_o)/2}^{D(\phi'_o)/2} K^{\gamma_p}(\phi_p^e(\phi'_o, s_o) - \phi_i) A_p(\phi'_o, s_o) ds_o \right) d\phi'_o \right] d\phi_i \\ &= \int_0^{2\pi} \frac{R_p(\phi_i, \phi_o)}{D(\phi_i)} \bar{L}_{i,p}(\phi_i) d\phi_i, \end{aligned}$$

where

$$R_p(\phi_i, \phi_o) = \int_0^{2\pi} K^{\gamma_p}(\phi_o - \phi'_o) \left(\int_{-D(\phi'_o)/2}^{D(\phi'_o)/2} K^{\gamma_p}(\phi_p^e(\phi'_o, s_o) - \phi_i) A_p(\phi'_o, s_o) ds_o \right) d\phi'_o,$$

and we call this function the *blurred response function*. In the supplementary material, we show that it is reciprocal: $R_p(\phi_i, \phi_o) = R_p(\phi_o, \phi_i)$. So the integral on the right-hand side can be written in terms of ϕ'_i and s_i instead of ϕ'_o and s_o .

With the blurred response function, we could define the ASF as $R_p(\phi_i, \phi_o)/D(\phi_i)$, which obeys reciprocity in the sense of Equation (3). However, in this exact form it is difficult to ensure energy conservation, so we define the ASF for the cross section in a slightly different form,

$$N_p(\phi_i, \phi_o) = \frac{R_p(\phi_i, \phi_o)}{D^{\gamma_p}(\phi_i)}, \quad (8)$$

where the *blurred diameter function* $D^{\gamma_p}(\phi)$ is the convolution of the projected diameter function with the angular kernel:

$$D^{\gamma_p}(\phi) = \int_0^{2\pi} D(\phi') K^{\gamma_p}(\phi - \phi') d\phi'.$$

The ASF defined above is energy conserving, but it is *not* reciprocal in the sense of Equation (3). However, it is close to being reciprocal because it satisfies a similar equation as follows:

$$\frac{N_p(\phi_i, \phi_o)}{D^{\gamma_p}(\phi_o)} = \frac{N_p(\phi_o, \phi_i)}{D^{\gamma_p}(\phi_i)}.$$

That is, instead of the relevant values being equal relative to diameter, they are equal *relative to blurred diameter*. We call this property *approximate reciprocity*. A proof of energy conservation and approximate reciprocity is available in the supplementary material. The property might be exploited in the design of sampling algorithms of ASFs. However, we did not use it in our implementation.

4.2.3 Discussion. Our ASF is a generalization of d'Eon's ASF to non-circular cross section. One difference between their ASF and ours is that their ASF only blurs over the outgoing direction ϕ_o , but ours blurs over both the incoming and outgoing directions. d'Eon's ASF is reciprocal because it is defined only for circular cross section: It is a function of $\phi_i - \phi_o$, so blurring over either ϕ produces the same result. However, with a non-circular cross section, blurring over ϕ_i differs from blurring over ϕ_o , so blurring over just one angle would result in significant non-reciprocity.

Note also that, in this section, we only have defined the ASF $N_p(\theta_i, \phi_i, \phi_o)$ in terms of the functions $D(\phi)$, $\phi_p^e(\phi, s)$, $s_p^e(\phi, s)$, and $A_p(\phi, s)$, which depend on the incoming azimuthal angle θ_i as well as on the specific shape of the cross section.

4.3 Complete Description of the Model

The model is defined by the following parameters:

- The aspect ratio of the elliptical cross section.
- The index of refraction η .
- The absorption coefficient σ of the material inside the fiber.
- For each mode p , the longitudinal shift α_p .
- For each mode p , the width of the longitudinal lobe β_p .
- For each mode p , the width of the angular kernel used in azimuthal scattering γ_p .

The projected diameter function $D(\phi)$ is defined by the shape of the cross section and is unchanged with respect to θ_i . The three other functions that characterize the cross section depend on Bravais's effective index of refraction,

$$\eta' = \sqrt{\eta^2 - \sin^2 \theta_i} / \cos \theta_i.$$

In particular, $\phi_p^e(\phi, s)$ and $s_p^e(\phi, s)$ are defined by tracing the ray (ϕ, s) through the cross section, substituting η' for η , and allowing the ray to intersect the cross section $p + 1$ times before going out. For the attenuation,

$$A_p(\phi, s) = \exp(-\sigma \ell_p(\phi, s) / \cos \theta_i) F_p(\phi, s),$$

where $\ell_p(\phi, s)$ is the length of the portion of the path that lies inside the cross section and $F_p(\phi, s)$ is the product of the Fresnel factors associated with all the intersections with the cross section's surface. If the ray does not emerge from the cross section due to total internal reflection, then we say that $A_p(\phi, s) = 0$.

4.4 Implementation

We have implemented the full scattering model as a part of a path tracer written in Java.

4.4.1 Rendering. For ray tracing, we model each hair fiber as a polyline in 3D. For each polyline segment, there corresponds a circular cylinder whose axis of symmetry coincides with the segment. Cylinders of two segments meeting at a corner are joined with a miter joint so a fiber corresponds to one continuous surface.

Because the actual projected diameter of fibers is a constant, the outgoing radiance integral (1) reduces to

$$\int_{-\pi/2}^{\pi/2} \int_0^{2\pi} S(\omega_i, \omega_o) L_i(\omega_i) \cos^2 \theta_i d\phi_i d\theta_i.$$

We note that, while the actual cross sections are circular, the BCSDf S above can be that of an elliptical cross section. The discrepancy between the actual cross section and the BCSDf's cross section does not cause an accuracy problem: When viewed from a distance, the circular fibers will look as bright as elliptical fibers. To see why this

assertion is true, recall that (1) the BCSDf outputs curve radiance, which is the total outgoing energy across the projected diameter, and (2) outgoing radiance is curve radiance per unit projected diameter. Hence, the outgoing radiance is the curve radiance of an elliptical fiber divided by the projected diameter of a circular fiber. Thus, if we accumulate radiance across the circular fiber's projected diameter (which happens when the viewer is far away from the fiber), we get the elliptical fiber's curve radiance, that is, its total brightness.

Nevertheless, renderings made with circular cylinders can still differ from those made with elliptical cylinders because different geometry results in different paths that light propagate through the scene. (For example, elliptical fibers can cast wider shadows.) Still, we chose to use circular cylinders because it is hard to join two elliptical cylinders in a way that results in one continuous surface.

4.4.2 Cross Sections. We always use an ellipse whose area is π and whose major axis is the v -axis. In other words, the equation of the ellipse in the vw -plane is $v^2/a + aw^2 = 1$. We tabulate the projected diameter function $D(\phi)$ and then evaluate it by linear interpolation.

Unlike with circular cross sections, closed-form formulae for ϕ_p^e , s_p^e , and A_p are unwieldy for elliptical cross sections, so we compute these functions using a simple 2D ray tracing procedure.

4.4.3 Computing the ASF. We only evaluate the first four ASFs (that is, up to $p = 3$). For each mode, we tabulate the ASF $N_p(\theta_i, \phi_i, \phi_o)$. Because the ASF is an even function in θ_i regardless of the cross sectional shape, we only consider θ_i ranging from 0 to $\pi/2$. On the other hand, we did not consider the symmetry of the cross section and let the azimuthal angles ϕ_i and ϕ_o assume their full range of values: $[0, 2\pi)$. Our grid resolution is $128 \times 256 \times 256$. Each entry of the table contains a single-precision floating point number for each of the three color channels. The total space of the tables is about 400MB.

To precompute the ASFs, we evaluate $N_p(\theta_i, \phi_i, \phi_o)$ in slices of fixed θ_i , slice by slice, by kernel density estimation. For a fixed value of θ_i , we randomly generate $N = 1,000,000$ incoming rays $(\phi_{i,1}, s_{i,1}), \dots, (\phi_{i,N}, s_{i,N})$, where ray $(\phi_{i,j}, s_{i,j})$ has power $2\pi D(\phi_{i,j})/N$. For each ray, we compute $\phi_{o,j} = \phi_p^e(\phi_{i,j}, s_{i,j})$, $s_{o,j} = s_p^e(\phi_{i,j}, s_{i,j})$, and $A_j = A_p(\phi_{i,j}, s_{i,j})$. Then, $N_p(\theta_i, \phi_i, \phi_o)$ is approximated as

$$N_p(\theta_i, \phi_i, \phi_o) \approx \frac{2\pi}{ND^{\gamma_p}(\phi_i)} \sum_{j=1}^N A_j D(\phi_{i,j}) K^{\gamma_p}(\phi_i - \phi_{i,j}) K^{\gamma_p}(\phi_o - \phi_{o,j}).$$

We use the 1D Gaussian $g(\phi; 0, \gamma_p)$ as our kernel function K^{γ_p} (thereby making $K^{\gamma_p}(\phi_i - \phi_{i,j})K^{\gamma_p}(\phi_o - \phi_{o,j})$ a 2D Gaussian). When evaluating the above sum, we exclude points that are of 2D Euclidean distance at least $3\gamma_p$ from (ϕ_i, ϕ_o) . Because we also take into account that the angles ϕ_i and ϕ_o wrap around, our kernel function approximates the Gaussian detector function introduced in d'Eon et al. [2011].

4.4.4 Sampling. Efficient evaluation of the outgoing radiance in a path tracer requires importance sampling of (θ_i, ϕ_i) given (θ_o, ϕ_o) . To do so, we drop the incoming radiance term and rewrite the integral as

$$\sum_{p=0}^3 \int_{-\pi/2}^{\pi/2} M_p(\theta_i, \theta_o) \cos^2 \theta_i \left(\int_0^{2\pi} N_p(\theta_i, \phi_i, \phi_o) d\phi_i \right) d\theta_i.$$

The rewritten integral suggests we divide the sampling of θ_i into three steps: (1) sample the mode p ; (2) given p , sample θ_i ; and then (3) given p and θ_i , sample ϕ_i .




Fig. 5. (ϕ_i, ϕ_o) -plots of the first four modes of the ASF at $\theta_i = 0^\circ$ for an elliptical fiber with $a = 1.0$ and $a = 1.5$. According to the model, elliptical fibers give rise to many features not present in circular fibers such as the bright bars in the R mode (\searrow), the extra lobes in the TT mode ($\cdot \cdot$), the eyes in the TRT mode (\circ°), and the winged blobs in the TRRT mode. (All the functions were calculated with $\eta = 1.55$, and $\sigma = 0$. Plots in the same rows are to scale with one another but not across the rows. The darkest color (black) corresponds to the minimum value of plots in the same row, and the brightest color (white) corresponds to the maximum value.)

To carry out the first two steps, for each mode p , we precompute a 3D table T_p indexed by θ_o , ϕ_o , and θ_i with

$$T_p[\theta_o, \phi_o, \theta_i] = M_p(\theta_i, \theta_o) \cos^2 \theta_i \left(\int_0^{2\pi} N_p(\theta_i, \phi_i, \phi_o) d\phi_i \right).$$

In our implementation, T_p has dimension of $128 \times 128 \times 256$. With the table, the power of each mode can be computed by summing the $T_p[\theta_o, \phi_o, \cdot]$ entries. We can then use the power of the modes to sample a mode. After sampling the mode, we can use the entries in the $T_p[\theta_o, \phi_o, \cdot]$ row to sample a θ_i . After sampling p and θ_i , we can use the table for the ASF to sample ϕ_i . We use binary search to locate the sampled table cell. Efficiency, however, may be improved by more sophisticated probability distribution function (PDF) inversion techniques [Cline et al. 2009].

Our implementation entails two space requirements. First, to sample the ASF efficiently, a table of partial sums of $N_p(\theta_i, \phi_i, \phi_o)$ as a function of ϕ_i is required, which incurs 133MB of space in addition to the storage required for N_p itself; thus, the storage pertaining to N_p is 533MB. Second, we must build T_p , which is approximately 67MB in size.

5. PREDICTION

In this section, we compute the ASFs of elliptical fibers and report the predictions of our model for the first four scattering modes: R, TT, TRT, and TRRT. We only consider fibers with $a \leq 1.6$ and $\sigma = 0$. While this is only a small subset of the possible range of parameters, it is sufficient to illustrate the differences between elliptical ASFs and circular ones and to show the qualitative trends that the ASF follows as a and θ_i change.

We provide two types of visualizations of the ASF. The first type (Figure 5 and Figure 7) is the (ϕ_i, ϕ_o) -plots discussed at the end of Section 3. The second type (Figure 6) is a polar plot of the ASF with ϕ_o as the free variable. Displaying the ASFs of different shapes on the same plot makes it easy to compare them. The polar plots in Figure 6 correspond to vertical slices through these images. The (ϕ_i, ϕ_o) -plots allow us to see the complete picture of how the ASF evolves with ϕ_i , and they expose the inherent symmetry due to reciprocity. It also provides a convenient way of observing how the ASF evolves with changing θ_i .

In Figure 5, the four modes of a circular fiber and an elliptical fiber are plotted separately to show the pronounced effect of ellipticity and the distinct features of each mode. In Figure 7, we plot the sums of the four modes for a number of combinations of cross sectional shapes and longitudinal angles to show how the modes evolve with these parameters.

We now make observations of the modes, one by one.

R mode. From the first row of Figure 6, we see that the R mode of elliptical fibers (colored curves) is quite similar to circular ones (black curves). All the plots contain two forward lobes. More eccentric ellipses scatter more forward and less backward than less eccentric ones when $\phi_i = 0^\circ$, and the reverse holds at $\phi_i = 90^\circ$.

The (ϕ_i, ϕ_o) plot in Figure 5 shows that the R mode of an elliptical fiber contains two bands of brighter reflection (). More specifically, this brighter reflection occurs when $\phi_i + \phi_o \approx 180^\circ$, which is to say, the half vector is normal to the flat side of the ellipse.

However, as can be seen from Figure 6, the R mode remains quite dim, even for elliptic fibers, so its visual significance is low when the other components are present.

TT mode. In Figure 6, the plots of the TT mode with $\phi_i = 15^\circ$ and $\phi_i = 30^\circ$ indicate that more elliptical fibers have two lobes in the TT mode. As in the circular case, there is a lobe that peaks at the forward scattering direction $\phi_o \approx \phi_i + 180^\circ$, but there is also an extra lobe, present only when ϕ_i is near 0° or 180° (when the incident light is seeing a narrower cross section) that peaks somewhat to the side. In the (ϕ_i, ϕ_o) -plot, the forward scattering lobe manifests as the two long parallel strips () as their centers are 180° “above” the line $\phi_i = \phi_o$. The extra lobe manifests as the small elongated blobs perpendicular to the bars ().

Figure 6 and Figure 7 suggests that the extra TT lobe is the brightest feature of the ASF (at least for highly transparent fibers). Moreover, as the cross section becomes more elliptical, the extra lobe gets brighter. The behavior of the lobes as θ_i increases is more complicated. The extra lobe gets brighter as θ_i increases and then quickly disappears, but the forward scattering lobes seem to get broader and dimmer slowly as θ_i increases.

In the supplementary material, we argue that the extra lobe arises from a group of incoming rays whose directions become very similar as they exit the cross section. As such, its origin is similar to that of the glints as described by Marschner et al. [2003].

TRT mode. The TRT mode of a circular fiber is dim compared to elliptical fibers. At $\theta_i = 0$, it contains two caustic lobes, which are dim and small in extent. As θ_i increases, the lobes merge into one brighter lobe, which then becomes broader and dimmer as θ_i continues to increase.

The TRT modes of elliptical fibers have much more complex behaviors. In the (ϕ_i, ϕ_o) -plot, the mode appears to have two components: the two “eyes” () and the center blob (). The center blob is very dim and appears as a single lobe only when ϕ_i is near 0° or 180° ; that is, when light strikes the narrow sides of the ellipses. The eyes contain two asymmetric lobes, many times brighter than in a circular fiber, that appear in the range of ϕ_i complementary to the center blob. Our observation is consistent with Marschner et al. [2003], who observed that as ϕ_i changes, the two TRT lobes seem to merge into a single lobe as the cross section rotates so the narrower side faces the light.

Figure 7 suggests that the eyes will become brighter as the fiber becomes more eccentric. The distance between the two eyelids and the curvature also seem to increase as a increases, so we can say that more elliptical fibers have “wider” eyes, meaning that the peaks occur farther from the incident direction. On the other hand, the center blob becomes broader and dimmer.

In fibers with small eccentricity, as θ_i increases, the eyelids get closer to each other, eventually collapse into a line, and then quickly become broader and dimmer. In highly eccentric cross sections ($a = 1.6$ in particular), the eyelids do not collapse but seem to maintain their distance while inverting the direction of curvature.

TRRT mode. In the (ϕ_i, ϕ_o) -plot, the TRRT mode of an elliptical fiber manifests as four “winged blobs” that occur near the line $\phi_i + \phi_o = 0$, when the half vector is aligned with the long axis of the ellipse. The TRRT mode is insignificant in circular fibers, but Figure 7 suggests that the blobs should be observable, particularly since they occur at angles far from where other brighter components dominate. When θ_i is small, the blobs are dimmer than the TT and TRT modes, but they can remain bright and become brighter as θ_i increases and the TT and TRT modes start to disappear. As a result, we include the TRRT mode in our implementation while other works only include up to the TRT mode [Marschner et al. 2003; Sadeghi et al. 2010].

In general, the blobs of the TRRT mode get larger and brighter as θ_i increases, and, after a point, they seem to quickly disappear. The angle θ_i where the blobs disappear seems to decrease as a increases.

To summarize, elliptical cross sections give rise to *much* more structure in the scattering function than is seen in circular fibers, with bright lobes in the TT and TRRT modes, which can appear and disappear quickly as ϕ_i changes. The TRT lobes, while similarly bright (and much brighter than in circular fibers), are more stable, but their locations and brightness vary based on eccentricity and orientation of the fiber. As a result, the model predicts that hair fibers with different aspect ratios will have glints at different angles.

6. MEASUREMENT DEVICE

To investigate how the behavior of real fibers compares to the predictions discussed in the last section, we designed and built a new image-based system for measuring scattering from individual fibers. This new system provides much higher angular resolution than previous methods based on four-axis gonioreflectometers

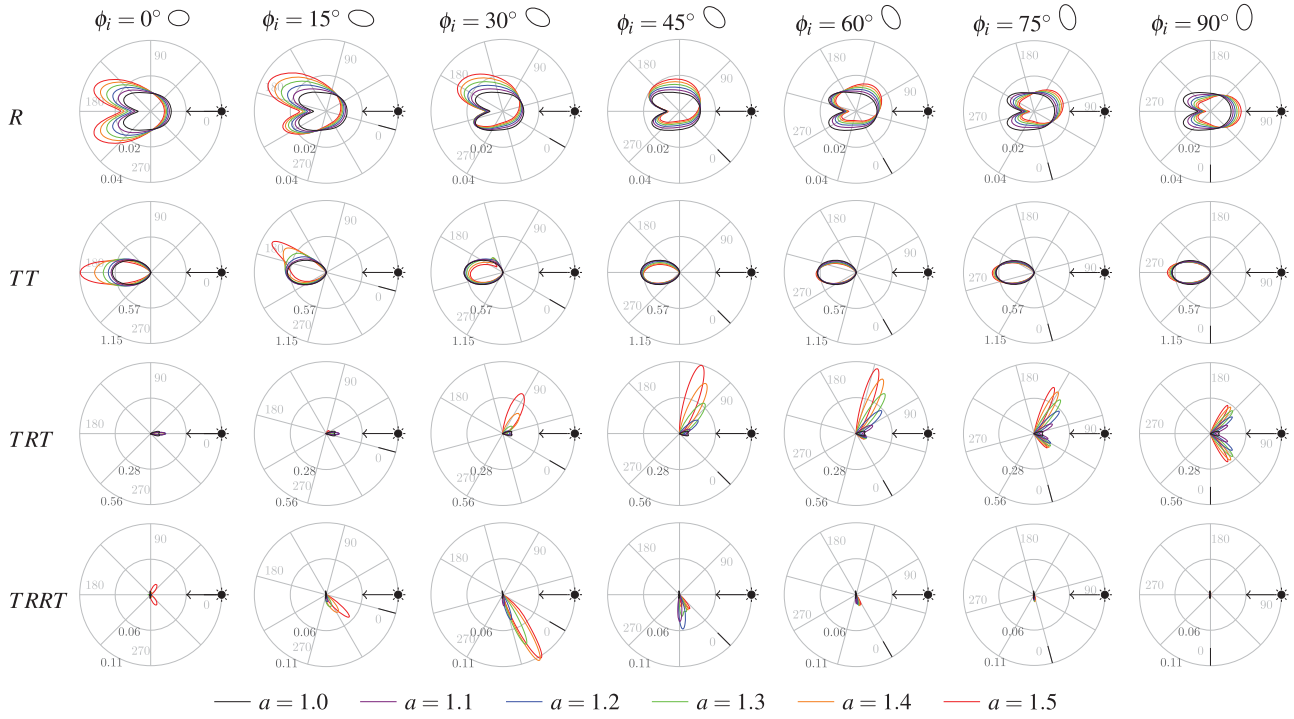


Fig. 6. Polar plots ($r = N_p(0^\circ, \phi_i, \phi_o)$ with ϕ_o being the free variable) of the ASF of elliptical cross sections with $a = 1.0, 1.1, 1.2, \dots, 1.5$. In general, higher aspect ratios are associated with brighter lobes in all modes. The plots are rotated so the light is coming from the right. All ASF values were calculated with $\eta = 1.55$, $\theta_i = 0^\circ$, $\sigma = 0$, and $\gamma_p = 5^\circ$. Plots in the same row are to scale with one another but not across the rows. Due to the symmetry of ellipses, plots of the ASF with ϕ_i being a multiple of 15° from 90° to 360° can be deduced from the plots we show here.

[Marschner et al. 2003] and thereby produces a more complete picture of the scattering function than has been seen before. The device, however, was designed for observing qualitative features of the scattered light and not for accurately measuring the BCSDF.

6.1 Description

Figure 8(a) shows our measurement system. Its main component is a bowl-shaped receiving surface that approximates a hemisphere of radius 7.5cm. The top part of the hemisphere is left out so the height of the bowl is 7.3cm. The bowl was created by 3D printing and painted a diffuse gray.

To perform a measurement, we suspend a hair fiber above the bowl, running along the bowl's diameter so the fiber passes through the center of the hemisphere.² Light from a white LED is channeled through optics to produce a narrow beam that illuminates a short length of the fiber at the center. The illuminated section is between 3mm and 6mm long at normal incidence, and the solid angle of the beam is adjustable between 7×10^{-5} and 2×10^{-4} steradian. The two ends of the fiber are attached to two stepper motors, which rotate in tandem around the hair's axis to change ϕ_i , and the bowl rests on a rotating platform so θ_i can be controlled by rotating the whole apparatus relative to the source.

²Short fibers must be extended—for example, by attaching other fibers to both ends—to be at least 20cm long before mounting. This way, we have used the device to measure a short wool fiber. (See the supplementary material of Khungurn et al. [2015].)

The bowl receives the light scattered by the hair fiber into the lower hemisphere and reflects it approximately diffusely. We measure this reflected light by taking a photograph using a camera mounted on one arm of the gonioreflectometer. The arm was positioned so the camera views the bowl from directly overhead. The relative position between the camera and the bowl was computed by the Caliber camera localization and calibration system [Liu et al. 2015]. Because one photograph only captures a hemisphere of scattered light, we need two photographs to capture a complete sphere. In particular, if we take a picture where the direction of incoming light is (θ_i, ϕ_i) , then we have to take another photograph with light direction $(-\theta_i, \phi_i + 180^\circ)$ to capture the complementary hemisphere. The motors and the rotating platform provide two degrees of freedom and the surface of the bowl provides two more, enabling measurement of the complete 4D scattering function.

A picture taken by the camera is shown in Figure 8(b), annotated to illustrate the angles that correspond to ones in the theoretical setting. The intensity of each pixel of the photograph that corresponds to the bowl's surface is a point sample of the hemisphere of scattered light. These point samples are the data produced by the system. We typically parameterize them with the spherical coordinates (θ_o, ϕ_d) , where θ_o ranges from -90° to 90° and $\phi_d = \phi_o - \phi_i$ ranges from 0° to 360° . The coordinates can be computed easily from the unit vector from the bowl's center to the point on the bowl surface corresponding to the pixel.

In Figure 8(c), we visualize the sphere of scattered light by plotting the point samples of the complete sphere on the plane where the horizontal axis is the azimuthal angle θ_o , and the vertical axis




Fig. 7. (ϕ_i, ϕ_o) -plots of the sum of the R, TT, TRT, TRRT modes of the ASFs of elliptical fibers with various aspect ratios and values of θ_i . While the R mode gets brighter as θ_i approaches 90° , all other modes become dimmer and eventually disappear at grazing angles. All functions were calculated with $\eta = 1.55$, $\sigma = 0$, and $\gamma_p = 0^\circ$. All the plots are to scale with one another. The darkest color (black) corresponds to the value 0, and the brightest color corresponds to the highest value in all the plots, which is roughly 1.8.

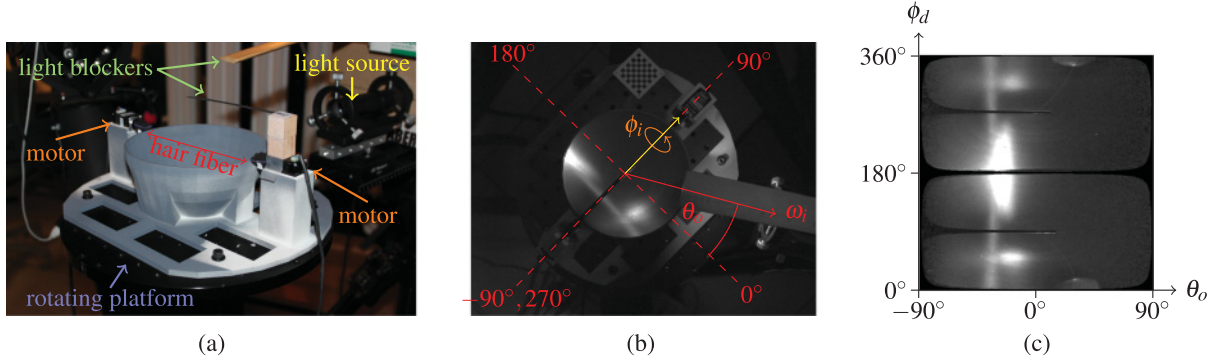


Fig. 8. (a) Our measurement system, (b) a picture taken by a camera located directly above the bowl, and (c) the picture after being mapped into spherical coordinate (θ_o, ϕ_d) , where $\phi_d = \phi_o - \phi_i$.

is the angle $\phi_d = \phi_o - \phi_i$. This is equivalent to the (θ_o, ϕ_d) -plot discussed in the Section 3.

6.2 Practical Design Issues

The device is designed to allow the illuminating beam to clear the edge of the bowl all the way around, so it has a blind spot about 5° wide around the equator of the sphere—all directions for which $\phi_i \approx \phi_o$ or $\phi_i \approx \phi_o + 180^\circ$.

In addition to the desired light, which travels from the source to the hair to the bowl to the camera, other paths contribute to the image as well. The direct path from the source to the hair to the camera produces a bright line at the center of the image, and the path from the source to the bowl to the camera³ produces a bright area on the rim of the bowl right next to the source. To

³The source is designed to miss the bowl, but when its aperture is fully open the beam spills slightly onto the rim, and even when it is closed down further, stray light still produces significant illumination on the rim.

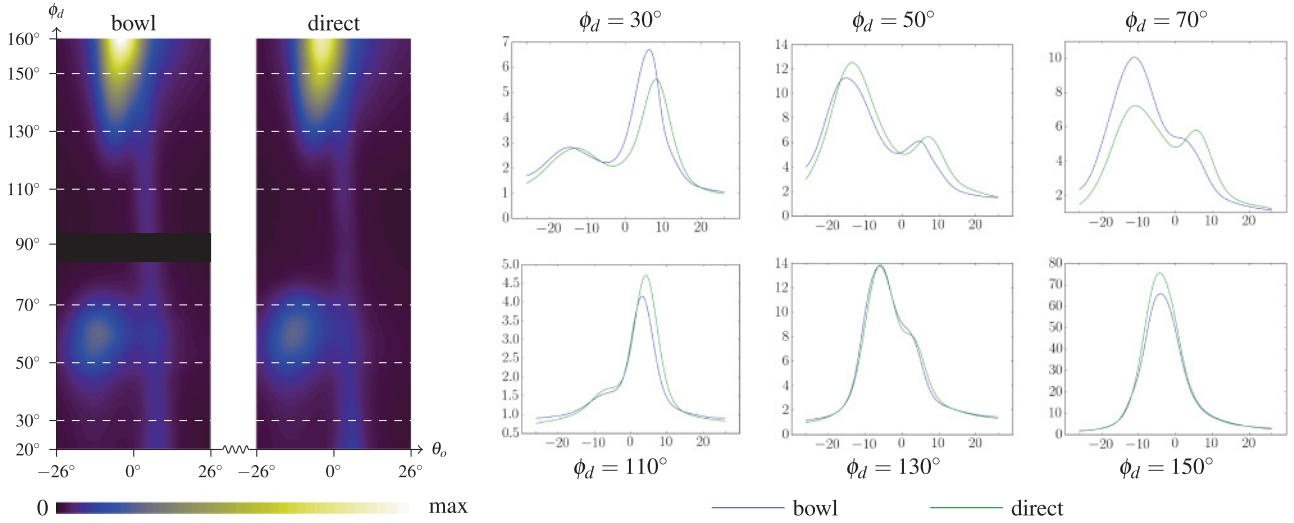


Fig. 9. Comparison between the direct measurement and the bowl measurement, which shows that the device can locate peaks in the BCSDf, making it suitable for qualitatively observing light scattering behavior of hair fibers. However, it is not accurate enough to give definite numerical values of the BCSDf. A segment around $\phi_d = 90^\circ$ of the bowl measurement is blacked out because it contains the blocker, and point samples in the area were removed before the reconstruction of the smooth function. Our optimization yielded the optimal translation with $\Delta\theta_o = 0.62^\circ$ and $\Delta\phi_d = -0.85^\circ$, $k_d = 2.58 \times 10^{-2}$, and $I_a = 24.23$. Here I_a is in the unit of pixel values at the time we performed the direct measurement. (The camera settings between the bowl and the direct measurement were identical, except that, in the direct measurement, we turned the aperture of the camera down so the camera took in 16 times less light, that is, 32 f-stops.)

combat lens flare these paths are blocked by two occluders; this is important to enable the camera to record the very dim illumination on the inside of the bowl, but they create two more narrow blind spots in the data.

Moreover, diffuse interreflection paths from other points on the bowl's surface also contribute; these are mitigated by the relatively low reflectance of the bowl and contribute an approximately constant component that is near the noise floor of our measurements. We will quantify this diffuse interreflection in the next section.

We wrote earlier that we take a “photograph” to capture an outgoing hemisphere of scattered light. In practice, this is a high-dynamic-range image acquired through three exposures typically at 5, 15, and 30s with the QImaging Retiga scientific camera. The exposure times were chosen so the brightest parts of the bowl surface were not overexposed. The image formation model we used to merge these photographs is detailed in the supplementary material.

6.3 Quantitative Characterization

To characterize the accuracy of our device, we compared our results to direct measurements of light scattered from a hair fiber using the following procedure. We mount a Caucasian hair fiber⁴ onto the bowl, rotate the platform so $\theta_i = 90^\circ$, and take a photograph of the bowl surface. We then rotate the fiber by 180° in θ and in ϕ , so the incident direction is the same, but the scattered light just measured using the bowl are visible from above. Subsequently, we acquire 4,048 images of the hair, corresponding to outgoing directions with $-26^\circ \leq \theta_o \leq 26^\circ$ and $20^\circ \leq \phi_d \leq 160^\circ$. From each photograph, we identify the illuminated region, sum up the pixel values inside, divide the sum by $\cos \theta_o$, and interpret the resulting value as the amount of light scattered from the hair fiber in the corre-

sponding direction. This process yields 4,048 unstructured samples of the scattered light.

The data from the above measurement process cannot be directly compared to the bowl measurement for several reasons:

- (1) The numbers of samples are not the same between measurements, and the samples are not at the same (θ_o, ϕ_d) coordinates.
- (2) The bowl measurement is much more noisy because of the roughness of the bowl surface.
- (3) The bowl data represents light reflected off the bowl surface, not the light directly scattered by the fiber.
- (4) Due to uncertainty in the automatic calibration of camera positions when performing the direct measurement, the data of the two measurements are not aligned perfectly. (For the bowl measurement, we manually calibrated the camera so the vertices of 3D model of the bowl aligned with those in the photograph.)

To remedy the first two issues, we reconstructed a function from the (θ_o, ϕ_d) -plane to the light intensity I by Gaussian kernel smoothing: If the point samples are $(\theta_{o,1}, \phi_{d,1}, I_1), \dots, (\theta_{o,N}, \phi_{d,N}, I_N)$, then the intensity value at (θ_o, ϕ_d) is given by

$$I(\theta_o, \phi_d) = \frac{\sum_{j=1}^N K(\theta_o, \phi_d; \theta_{o,j}, \phi_{d,j}) I_j}{\sum_{j=1}^N K(\theta_o, \phi_d; \theta_{o,j}, \phi_{d,j})}$$

where K is a 2D circular Gaussian function with standard deviation of 2° .

To remedy the third issue, we recognize that the intensity in the direct measurement I_{direct} is related to the intensity in the bowl measurement I_{bowl} by

$$I_{\text{bowl}} = k_d(I_{\text{direct}} + I_a), \quad (9)$$

where k_d is the diffuse reflectance of the bowl surface and I_a is the intensity of ambient light, which comes from multiple reflection off

⁴We used a fiber from Sample 3, details of which are given in Section 7.

the bowl surface. We optimized for k_d and I_a by minimizing the square error between the right-hand side and left-hand side of (9) at evenly spaced discrete sample points in the rectangle $-26^\circ \leq \theta_o \leq 26^\circ$ and $20^\circ \leq \phi_d \leq 160^\circ$.

To remedy the last issue, we optimized for a translation in the (θ_o, ϕ_d) -plane that, when applied to the direct sample points, minimizes the square error mentioned in the previous paragraph. The shift in ϕ_d accounts for error in the rotation of the turntable, and the shift in θ_o partially accounts for error in the vertical position of the hair fiber.

The plots of the direct and the bowl measurement after all the above adjustments are given in Figure 9. We can see that both measurements contain peaks that appear at approximately the same locations. That I_a is positive is consistent with the fact that inter-reflection can only add to the brightness of the reflected light from the bowl. Now, consider the plot with $\phi_d = 110^\circ$, which contains a peak of the R mode whose value is approximately 4. Because $k_d I_a \approx 0.63$, we infer that interreflection can be as bright as 15% of the R mode, implying that the R mode and other brighter modes can be clearly distinguished from background. However, while the peaks can be located and observed easily, their amplitudes in the two measurements do not match. Taking the direct measurement as ground truth, relative error can be as large as 101%. (This might be due to the fact that, even after applying the translations, the measurements did not align well enough.) In conclusion, we demonstrated that our device is adequate for observing and locating peaks in the BCSDf but not suitable for quantitatively measuring the function.

Nevertheless, we would like to point out the expediency of our device: the whole hemisphere of outgoing light can be acquired in less than a minute. On the other hand, the direct measurement took about 6h to perform but yielded only a small portion of the hemisphere.

6.4 Repeatability

We next test the repeatability of our measurements. To do so, we used the device to capture the scattered light from the same portion⁵ of a human hair fiber 3 times. For each measurement, we set $\theta_i = 90^\circ$, rotated the step motors so ϕ_i takes 180 equally spaced values in the 360° range and acquired the full sphere of scattering light for each of the ϕ_i values. This process results in 180 photos that constitute a 3D slice of the BCSDf with θ_i hold fixed.

We visualize 1 of 180 spheres from the three measurements in Figure 10. It can be seen that, while all three measurements contain the same features at roughly the same locations, the brightness and shapes of the features do not match exactly: those of Measurement 2 have the highest magnitude, followed by those of Measurement 1 and then Measurement 3. We hypothesize that the differences in brightness is due to differences in positioning the fiber in the beam, which is not uniform across its vertical extent.

Nevertheless, the measurements generally agree up to a constant factor. To see this, we also show graphs of a 1D slice of the measurements for several angles in the Figure 11. Here, the brightness values are divided by the sum of the values in the same measurements to account for the difference in scale between the measurements. We can see that the graphs have generally the same shapes.

⁵We marked a point 5cm from the root end of a fiber with black ink. When mounting the fiber on the bowl, we made sure that the marked part just cleared the end of the 3D printed part we attached to one of the motors' shaft. The ink mark was about 2mm wide, so the illuminated part in each measurement could differ, but they will be no further than 2mm apart.




Fig. 10. Visualization of one of the 180 spheres acquired in the three measurements of the same hair fiber. While the photographs appear similar, there are differences in brightness. In creating the visualization, we rendered all pixels whose brightness is greater than one-tenth of the maximum value to be white. This was done to make the brightness difference in the TRT blobs near the bottom easy to see. The red lines in the pictures are the 1D slices graphed in Figure 11.

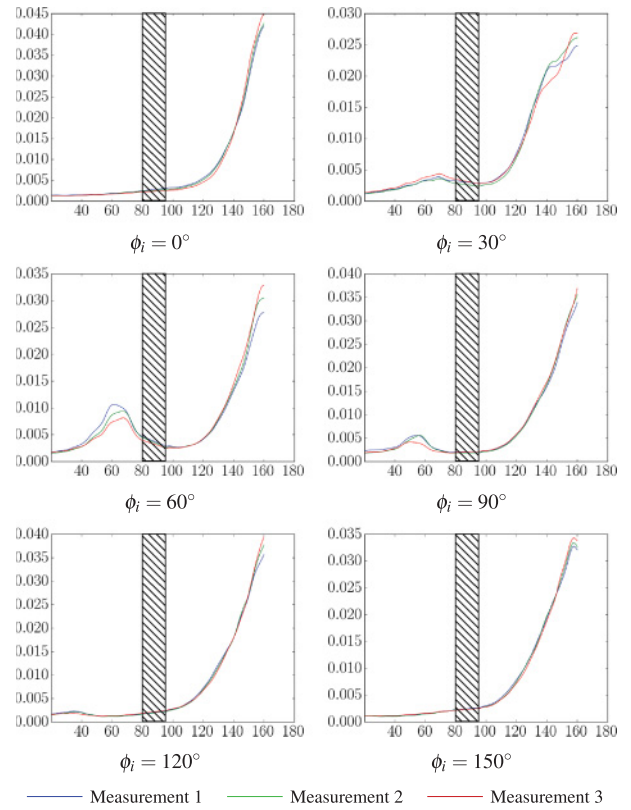


Fig. 11. Graphs of 1D slices of three repeated measurements for six different values of ϕ_i . The fact that all graphs (after correcting for scale differences) roughly coincide implies that the device is reliable enough to observe qualitative features in hair BCSDf. The slices' locations are indicated by the red lines in Figure 10. The x-axis represents ϕ_d , in degrees. We throw away data from $\phi_d = 80^\circ$ to $\phi_d = 95^\circ$ because they are contaminated by the blocker. To account for the difference in scale between the measurements, we made the y-axis represent the ratio of brightness at each point to the sum of brightness across all points in the same measurement. (That is, we divide the brightness values by their sum.)

From the measurements, we can conclude that our device is not reliable at recording absolute brightness of the scattered light. This is the direct consequence of our not controlling many factors that can affect the recorded brightness such as the vertical position of the fiber, the uniformity of the light beam, and the reflectance properties of the bowl itself. However, because the same features consistently appeared in all measurements, the device is suitable for observing large-scale features of the BCSDf. We believe that the shortcomings of the device can be overcome by better engineering, but, as it stands, the device is adequate for the purpose of this article.

7. MEASUREMENTS

In this section, we present measurements of a broad range of hair fibers and discuss features that we observed.

7.1 Hair Samples and Measurement Setups

We measured hair fibers taken from seven hair samples, which are identified by letters from A to G. We took 2 hair fibers from each hair sample, for a total of 14 fibers. The details of each fiber, along with its scanning electron microscope (SEM) image, are given in Figure 12.

The setups for all the measurements are almost the same, but with differences in the setting of the aperture in the light source optics. We measured all fibers that are numbered “1” first with the aperture fully open. As noted previously, this setup results in increased stray light. To reduce contamination of the data, we measured the fibers that are numbered “2” with the aperture closed to roughly one-third of the maximum area. This resulted in less stray light and sharper scattering patterns but required longer exposures to capture the dimmer signals.

For each fiber, we performed three types of measurements.

7.1.1 Aspect Ratio Estimation. As this article seeks to study the effect of ellipticity on the scattering behavior of hair fibers, we also made measurements to directly determine the fibers’ ellipticity. After the scattering measurements for each fiber, we cut the fiber at two to six locations near the illuminated spot and used a scanning electron microscope to image the cut ends. This procedure images the fiber’s cross sections at locations up to 3cm from the illuminated spot. For each of the micrographs, we identified 20–60 points on the boundary of the cross section and fit an ellipse,⁶ taking the ratio between the major and minor axes of this ellipse to be the aspect ratio of the cross section. All the resulting aspect ratios and their averages can be found in Figure 12. The full set of micrographs and fitting data is given in the supplementary material.

This process produces several samples of the aspect ratio of each fiber, but the data show that the aspect ratio can vary considerably along the length of a single fiber. Therefore they give only a rough indication of the aspect ratio at the scattering point. Several other factors also contribute to error in the measurements, including (1) the cut might not be exactly perpendicular to the fiber’s axis; (2) the cutting might have distorted the cross section’s shape; and (3) in some images, the cross sections were viewed at an angle 10° to 20° from the fiber axis, which had to be estimated manually. Despite these uncertainties, the data still provide useful information about how elliptic each fiber is.

⁶We used the fitting algorithm described in Fitzgibbon et al. [1999] and the code from van Foreest [2012].

7.2 (θ_o, ϕ_d) -Measurements

We start with (θ_o, ϕ_d) -measurements, as they are easily derived from a single photograph of the bowl. The measurements allow us to identify the observable features in the scattered light. To study how these features evolve as the incoming scattering angle θ_i changes, we varied the incoming longitudinal angle θ_i from 80° to −80° with the resolution of 2°. For each fiber sample, we chose the azimuthal angle ϕ_i so the features characteristic of the TRT mode can be observed. For each θ_i , we captured a (θ_o, ϕ_d) -plot, resulting in 81 plots. From these plots, we created videos (available in the supplementary material) that show the evolution of the scattering function as a function of θ_i . We also provide a sampling of these plots in the supplementary documents.

7.2.1 Observable Features. We found four different features in the generated plots, not all present in all fibers, which we attribute to different modes of scattering from the fibers. The features are identified for an example plot of Fiber B2 in Figure 13(a). The features are as follows:

- A bright, horizontally narrow, forward scattering line that always occurs at $\theta_o = -\theta_i$ (within the accuracy of our image registration). The feature is surrounded by the red rectangle in Figure 13.
- Two vertical strips; one originating from the top edge of the plot and the other from the bottom edge. Both strips curve towards the previous forward scattering line. (Green.)
- One or more forward scattering lobes that are horizontally wider than the previous features. (Cyan.)
- Two small blobs, one appearing near the bottom edge, and the other near the top edge. (Magenta.)

According to the heuristics discussed at the end of Section 3, we believe the green and magenta features are the R and TRT modes, respectively.

The four features are not present in all the plots. As can be seen in Figure 13(b) and the supplementary document, the cyan and the magenta features cannot be observed in plots of fibers from Samples E, F, and G, all of which are strongly pigmented fibers that appear black. This fact suggests that these two features are caused by light being transmitted through the fibers while the other two are not. As such, we believe that the cyan feature is the TT mode.

There are some disagreements between the data and the predictions made by previous models. The first is that previous models predict that the TT mode should appear more or less as one contiguous lobe. This prediction matches the wider forward scattering blobs of all light colored fibers except those from Sample C where there seem to be two separated blobs: one in the top hemisphere and the other in the bottom. On the other hand, the elliptical model allows the TT mode to manifest as two lobes: one for the slanted bar and one for the extra lobe discussed in the last section.

The second is that previous models predict that TRT blobs should appear only for values of ϕ_d that are within 20° of 0°—that is, quite close to the top or bottom of the plots. However, in the measurements such blobs appear as far as almost 90° from the plane of incidence. We showed through modeling of ASFs of elliptical fibers in Section 4 that with an elliptical cross section these blobs can still be explained by the TRT mode.

However, all models, ours included, do not predict the red feature.

7.2.2 The E Mode. The feature marked in red appears even for black hair, and its intensity increases at grazing angles. As such, it seems to be caused by reflection from the surface; therefore we consider it part of the R mode. However, it is distinct enough to need its own name, so we refer to it as the “E” mode. Here E stands for


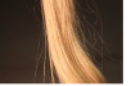

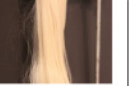

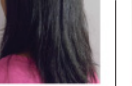

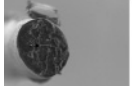


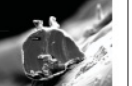
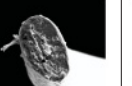
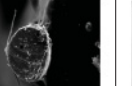
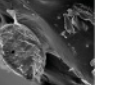

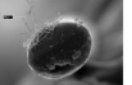

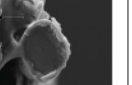
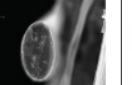
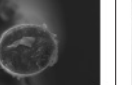
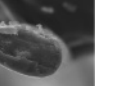
| | A | B | C | D | E | F | G |
|----------------------------------|---|---|---|---|---|---|---|
| Source | Caucasian | Caucasian | Caucasian | artificial wig | African | Chinese | Indian |
| Photographs |  |  |  |  |  |  |  |
| SEM images (Fiber 1) |  |  |  |  |  |  |  |
| Measured aspect ratios (Fiber 1) | 1.136 1.213 1.469 ($\bar{x} = 1.273$) | 1.446 1.501 ($\bar{x} = 1.473$) | 1.603 1.717 ($\bar{x} = 1.660$) | 1.127 1.385 ($\bar{x} = 1.256$) | 1.834 1.923 1.933 ($\bar{x} = 1.897$) | 1.308 1.361 ($\bar{x} = 1.334$) | 1.310 1.482 ($\bar{x} = 1.396$) |
| SEM images (Fiber 2) |  |  |  |  |  |  |  |
| Measured aspect ratios (Fiber 2) | 1.256 1.268 1.411 1.429 ($\bar{x} = 1.341$) | 1.077 1.195 1.246 1.353 1.436 1.552 ($\bar{x} = 1.310$) | 1.548 1.553 1.746 1.762 ($\bar{x} = 1.605$) | 1.147 1.159 1.233 1.310 ($\bar{x} = 1.212$) | 1.561 1.647 1.789 ($\bar{x} = 1.666$) | 1.035 1.123 1.142 1.183 ($\bar{x} = 1.121$) | 1.503 1.524 1.609 1.616 1.780 1.839 ($\bar{x} = 1.645$) |

Fig. 12. Hair samples.

“equi-angle peak” (EAP), a feature of light scattered from human hair reported by Stamm et al. [1977] that we believe is the same as our red feature.

In particular, Stamm et al. observed a faint peak at the ideal specular angle that grows brighter as θ_i increases in magnitude. Nevertheless, they only measured light scattered backward ($\phi_d = 0^\circ$) while we measured the whole range of ϕ_d . This enables us to see that the E mode is the brightest around forward directions and is very sharp in the θ_o direction. It is clearly separated from the rest of the R mode, and it grows stronger as θ_i approaches grazing angles, just like Stamm’s EAP. Hence, Stamm’s EAP was the part of the E mode that crosses the $\phi_d = 0^\circ$ line. Note that Marschner et al. also measured a full hemisphere of scattered light but not with high-enough resolution to distinguish the E mode from other modes.

The cause of the E mode is unclear. Stamm et al. provided an explanation in terms of reflection from subsurface interfaces, but it does not fit our full 3D observations. We conjecture that the E mode may be a wave optics effect related to the phenomenon of emerging specularly on flat surfaces [He et al. 1991], but a proper explanation will have to await further investigation.

7.2.3 Longitudinal Shifts. With the exception of fibers from the artificial wig, we found that all fibers exhibit longitudinal shifts of the R, TT, and TRT modes, consistent with what has been observed in previous works. That is, the R mode shifts to one side of the $\theta_o = -\theta_i$ line while other modes shift to the other side. On the other hand, as can be seen in Figure 14(a), all the modes of fibers from the artificial wig are centered horizontally at $\theta_o = -\theta_i$. This observation can be explained by the fact that fibers from the artificial wig have




Fig. 13. Observable features in the (θ_o, ϕ_o) -plot of (a) Fiber B2 and (b) Fiber E2. The features marked by colored rectangles: R mode is green, TT mode is cyan, and TRT mode is magenta. Fiber E2 has a dark color, so the TT and TRT mode are not present. The red rectangle surrounds the E mode, a phenomenon not predicted by any hair scattering models, including ours. The two plots were generated from measurements at $\theta_i = -30^\circ$.

no scales (see Figure 14(b)) and behave like smooth cylinders. Consequently, all modes of scattered light are contained in the cone defined by $\theta_o = -\theta_i$.

7.2.4 Evolution with Respect to θ_i . With more detailed measurements than previously available, we can more completely characterize how the features change with respect to θ_i . For brevity, we discuss the evolution in detail in the supplementary document. The general trend is that the R mode and the E mode are bright




Fig. 14. Fibers without scales such as those from an artificial wig yield no longitudinal shifts. Here we show (a) a (θ_o, ϕ_d) -plot of Fiber D2, which is artificial, and (b) an SEM micrograph showing the surface of the artificial hair, which has no scales. Notice that all the modes' centers are aligned horizontally.

near grazing angles but dim where θ_i is low in absolute value. In contrast, other modes behave in the exact opposite way.

7.3 (ϕ_i, ϕ_o) -Measurements

To get a more complete picture of the azimuthal scattering behavior, we created (ϕ_i, ϕ_o) -plots like the ones in Section 3.8. To do so, we set the rotating platform to a fixed position to hold θ_i constant and then rotate the fiber to achieve 180 equally spaced ϕ_i values from 0° to 360° . For each ϕ_i value, we captured the two photographs to cover the complete sphere and warped the photos into spherical coordinates as explained previously. To counteract the contribution of interreflection, we subtracted from each pixel 0.5% of the maximum intensity of all images taken with the same θ_i . From each of these 180 processed images, we selected a fixed column, corresponding to a particular value of θ_o , and arranged the columns to form a (ϕ_i, ϕ_o) -plot.

When conducting the measurements, we did not have means to determine the absolute ϕ_i angle. As such, the generated (ϕ_i, ϕ_o) -plot is a shifted version of the plots in Section 3. To make comparison with scattering models easier, we selected this unknown ϕ_i offset to align the observed features across the measurements and with the model in Section 4. Then, for convenience, we consider the left side of the generated plots to correspond to $\phi_i = 0^\circ$.

For each fiber, we created three sets of (ϕ_i, ϕ_o) -plots corresponding to $\theta_i = 0^\circ$, -30° , and -60° for each fiber to see the effect of longitudinal angle on the ASF. We chose negative angles instead of positive ones because the area with most of the scattered energy (around $\theta_o = 30^\circ$ and 60° , respectively) does not contain one of the light blockers.

We choose four different values of θ_o so they align with the features we believe to be the E, R, TT, and TRT modes. We identified the R, TT, and TRT modes using the heuristics discussed at the end of Section 3. For the E mode, we use $\theta_o = -\theta_i$. As we did not generate a plot when the corresponding mode is not present or when it is not well separated from other brighter modes, this process resulted in at most 12 plots being generated for each fiber. The plots are available in Figure 15, Figure 16, and the supplementary material.

The above measurements, which are intended to explore the ASF, actually do not measure the ASF directly. According to the separable model discussed in Section 4, these plots should show $M(\theta_i, \theta_o)N_p(\theta_d, \cdot, \cdot)$, which differs from $N_p(\theta_d, \cdot, \cdot)$ only by a constant, with additive contamination from other modes. However, in reality, the shifts and widths of the lobes modeled by M also change

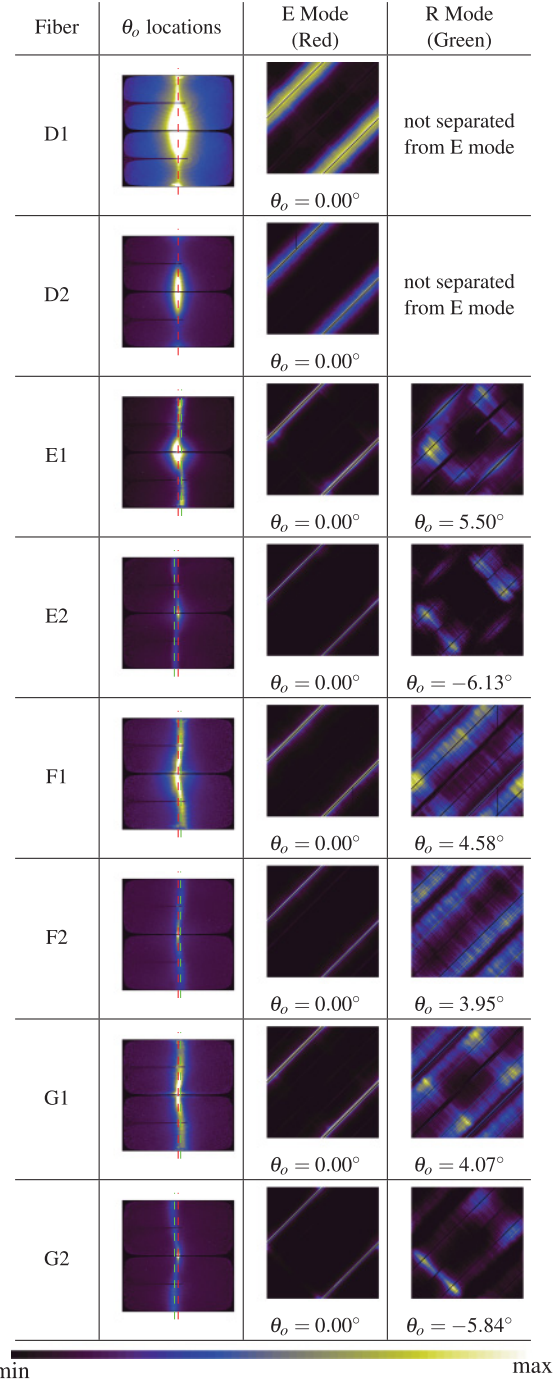


Fig. 15. (ϕ_i, ϕ_o) -plots of fibers from the artificial wig (Sample D) and dark-colored human hair (Samples E to G). Only the E and R modes are shown because other modes are either (1) not separated from the two in case of the wig fiber or (2) not visible at all in case of dark hair fibers. All plots are generated with $\theta_i = 0^\circ$. The complete set of plots with $\theta_i = -30^\circ$ and -60° is given in the supplementary materials. None of the plots are to-scale with one another.

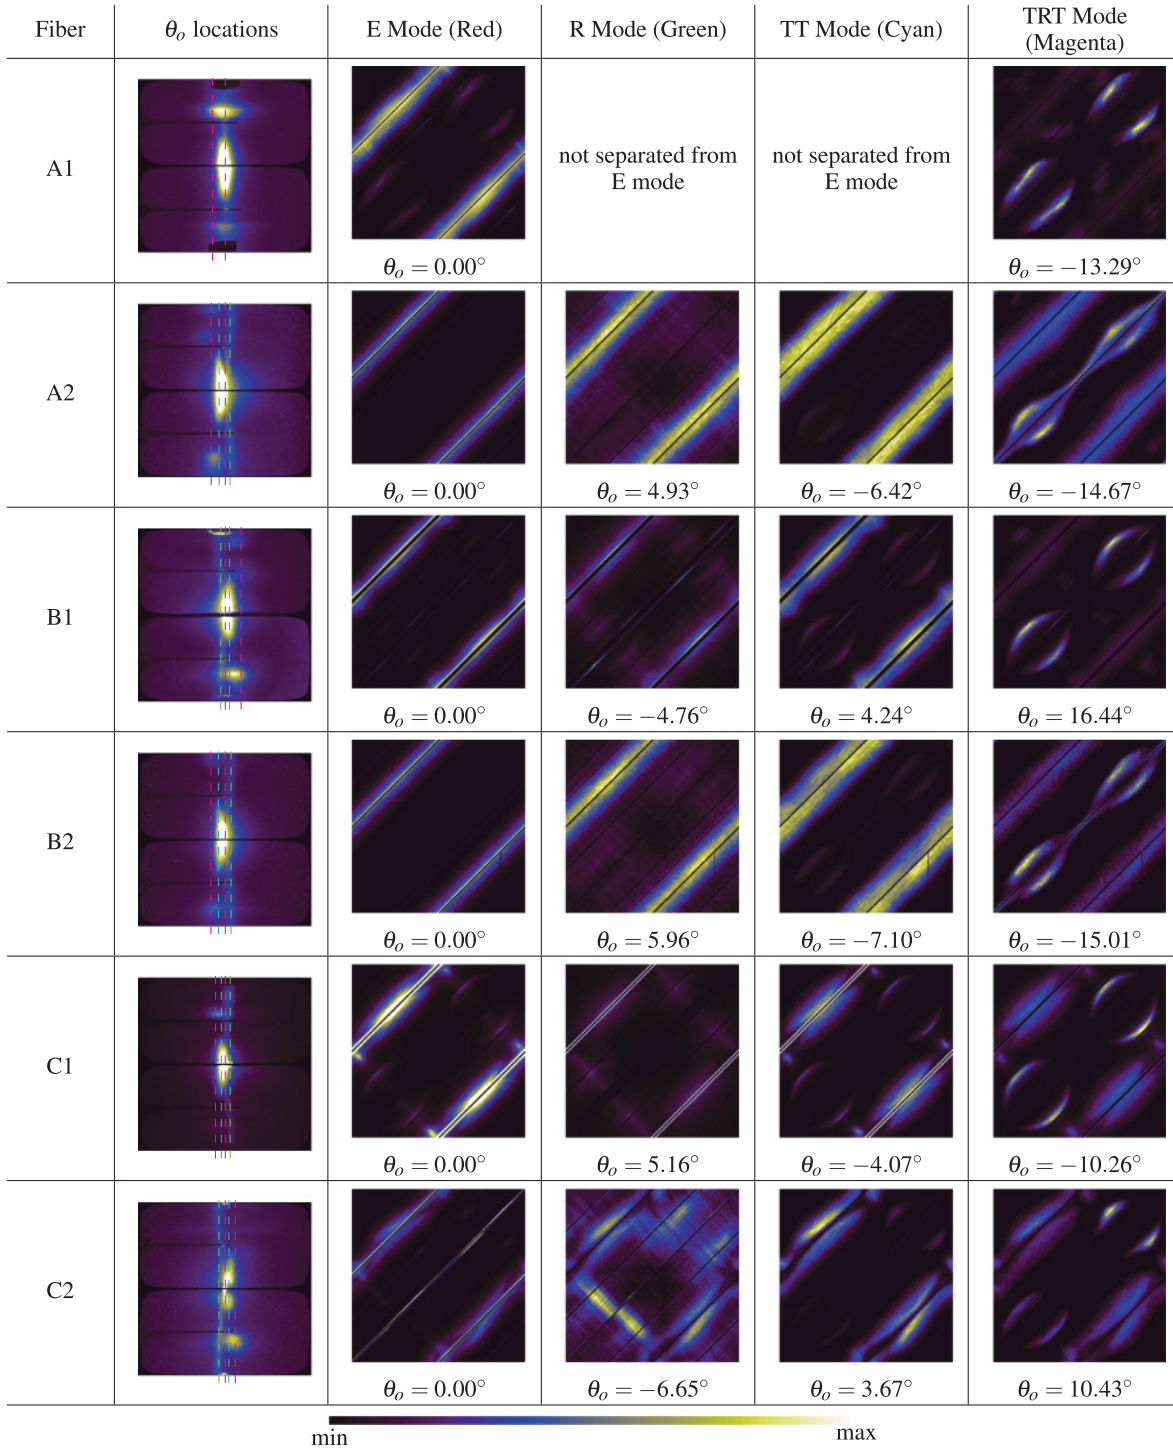


Fig. 16. (ϕ_i, ϕ_o) -plots of light-colored human hair fibers (Samples A to C) generated at $\theta_i = 0^\circ$. The plots clearly show that the TRT eyes ($^\circ$) predicted by the elliptical model exist in real hair fibers. Behaviors of other modes are harder to observe because they are not well separated from one another. Evidence of the brighter bands in the R mode (\setminus) can be found in Samples B and C. The extra lobes in the TT mode ($\cdot\cdot$) can be clearly observed in Fiber B1 and C1. The complete set of plots with $\theta_i = -30^\circ$ and -60° is given in the supplementary materials. None of the plots are to scale with one another. For the first three modes of Fiber C1, due to the large spill on the far rim of the bowl, we clamp the maximum pixel value to about 50% of the real maximum pixel value to make it possible to observe the features more clearly. For the E mode of Fiber C2, we excluded the bright narrow band around the backward scattering direction when computing the maximum pixel value because it was also caused by spilled light.




Fig. 17. The (ϕ_i, ϕ_o) -plots of the TRT mode of the Marschner model for a number of aspect ratios. The model with low aspect ratio approximates the TRT mode of Fiber A2 and B2 well though the variations in brightness are incorrect. At higher aspect ratios, it cannot produce the bright, separated eyelids seen in real fibers. In all plots, $\eta = 1.55$, $\theta_d = 0$, $k_G = 1$, $w_c = 10^\circ$, $\Delta\eta' = 0.2$, and $\Delta h_m = 0.5$. The plots are not to scale with one another. The original article only recommends using the model for only $0.85 \leq a \leq 1/0.85 = 1.176$, but we show plots with a outside the range as well.

as a function of ϕ_i : The non-constant shifts cause some parts of the ASF to appear dimmer or even not to be observable in the slices we took, and the changing width produces a change in intensity. These effects are especially significant in the R mode, which, as we previously observed, curves quite noticeably in the (θ_i, ϕ_d) -plots.

It is important to keep these considerations in mind when comparing the measurements to the theoretical N function: Qualitative agreement in location and shape of features can be expected, but intensities are not expected to match, and some features can be missing from the measured ASFs.

We now make observations of the scattering modes. While we list the E mode first in the figures, we will describe it last because it is usually contaminated by other modes.

7.3.1 R Mode. All the fibers except those from Samples D and F exhibit two bands of brighter reflection that are predicted by the elliptical model but not previous circular models.

7.3.2 TT Mode. Fibers A1, B1, C1, and C2 exhibit both the two slanted bars (\swarrow) and the small elongated blobs (\circ). Note that previous circular models predict the slanted bar but not the small blobs, so the elliptical model is more accurate. Fiber A2 and B2 exhibit only the strips, so their cross sections might be close to circular. As with the R mode, contamination from the TRT mode is sometimes visible in these plots.

7.3.3 TRT Mode. The two eyes (\circ) predicted by the elliptical model can be observed. In fibers where the eyes open narrowly (A2 and B2), there exists a band along the $\phi_i = \phi_o$ line joining the two eyes. In other fibers, the eyes are completely separated.

On the other hand, circular scattering models predict that the TRT mode appears as a narrow band of uniform width around the $\phi_i = \phi_o$ line. The Marschner model, with its approximation of elliptical fibers, can predict the TRT modes of Fiber A2 and B2 reasonably well (though it puts too much energy at retroreflection and not enough in the “eyelids”), but not those of other fibers. (See Figure 17.)

7.3.4 E Mode. The E mode consists of narrow and bright bands around the forward scattering lines, which get broader with increasing θ_i . The E mode plots also show features resembling the TT mode, including slanted bars (particularly fibers C1 and C2) and perpendicular blobs (especially for $\theta_i = -60^\circ$ in more elliptical fibers; see supplementary material), which we interpret as contamination from the TT mode.

7.3.5 Conclusions. From these observations, it is clear that real hair fibers exhibit features that are better explained by elliptical cross sections. These features include (1) the brighter bands in the

R mode, (2) the small blobs in the TT mode, and (3) the eyes in the TRT mode.

8. MODEL EVALUATION

In this section, we further evaluate the elliptical scattering model by comparing the model (ϕ_i, ϕ_o) -plots against the measured data. We also present renderings produced by the model so its visual impact can be assessed.

8.1 Comparison with Measured Data

To validate our model’s ability to predict scattering behaviors, we fit the model to the measured (ϕ_i, ϕ_o) -plots and then compared the plots generated from the fitted parameters with the original data. We

- used the θ_i values from the measurements,
- fixed σ to be 0 (perfectly transparent fibers),
- fixed η to be 1.55, and
- fixed γ_p to be 5° for all p .

This leaves the aspect ratio as the only parameter to be determined.

For each fiber, we pick a mode that is visible and varies the most as a changes⁷

- the TRT mode for light-colored human hair fibers (Samples A, B, and C),
- the TT mode for artificial fibers (Sample D), and
- the R mode for dark-colored fibers (Sample E, F, and G).

Selecting a mode limits us to work with three measured plots: one for each of three values of θ_i . We then find a for which the three corresponding (ϕ_i, ϕ_o) -plots are the most similar to the three measured plots. To do so, we iterate through values of a ranging from 1.00 to 1.99 by increments of 0.01. For each a , we generate three plots with θ_i being 0° , -30° , and -60° . Next, all plots, including the measured plots, are normalized so the pixel values add up to 1. (That is, they are regarded as probability distributions over the (ϕ_i, ϕ_o) -plane.) The score associated with this a value is the sum of the L_1 -distances between the three pairs of generated and measured plots. The a value with the lowest score is regarded as the best fit for the fiber. The plots of the mode being fitted at $\theta_i = 0^\circ$, generated

⁷It is possible to fit the aspect ratio using a (ϕ_i, ϕ_o) -plot that contains more than one mode. However, the problem to solve becomes more complicated because the contribution of each mode is scaled by an unknown factor, which has to be optimized for. We chose to involve only one mode because it yields a much simpler optimization problem.

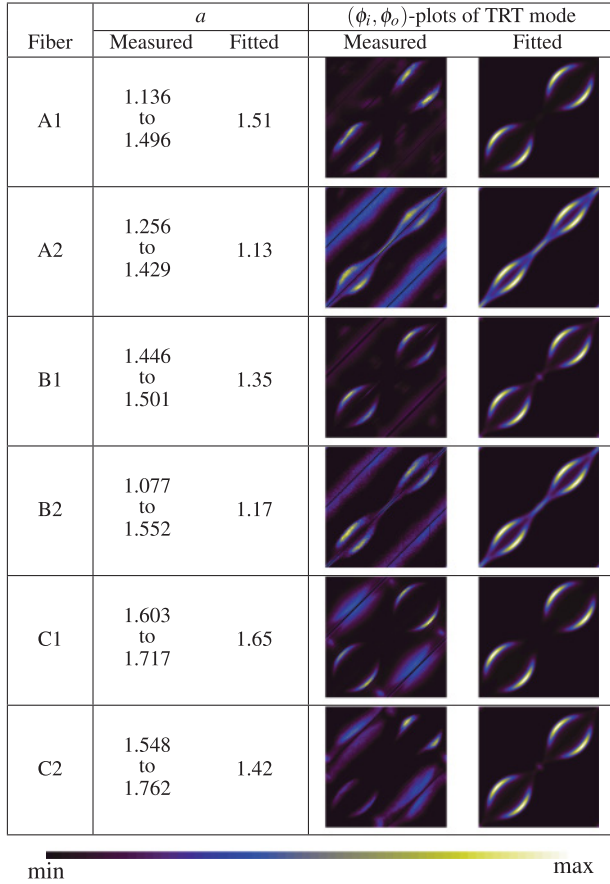


Fig. 18. Results of fitting fiber aspect ratios for Fiber A to C to match the captured (ϕ_i, ϕ_o) -plots of the TRT mode. The generated plots match well to the measured plots, and the fitted aspect ratios are close to the measured ones. Here we only show the plots of one mode at $\theta_i = 0^\circ$. Plots of other modes at other θ_i values are shown in the supplementary material. The plots' colors are not to scale with one another.

from the best aspect ratios, are shown in Figure 18 and Figure 19. The supplementary material contains plots of other modes at other values of θ_i .

For light-colored human hair fibers, the model can be fitted well to the TRT mode. We invite the reader to consult the supplementary material to see that the fitted aspect ratios yield reasonable plots for the TT mode, too. Moreover, the fitted aspect ratios either fall within or close to the range of measured aspect ratios. The successes of matching to the TRT mode tells us that, while fibers are not perfectly elliptical, their TRT modes usually behave like those of elliptical fibers.

As fibers from Sample D do not separate the scattering modes in the θ_o direction, we could only match the parameter against the TT mode, which is the brightest and completely obscures other modes. Because the TT modes of fibers from Sample D behave much like a perfectly circular fiber, the matching process results in $a = 1.00$ being the best fit aspect ratio. However, SEM micrographs indicated that the fibers are not perfectly circular, and the (ϕ_i, ϕ_o) -plots generated using the average measured aspect ratio would display perpendicular blobs, which are not present in the measured plots. This discrepancy might be caused by the artificial fiber's index of refraction being lower than that of human hair fibers.

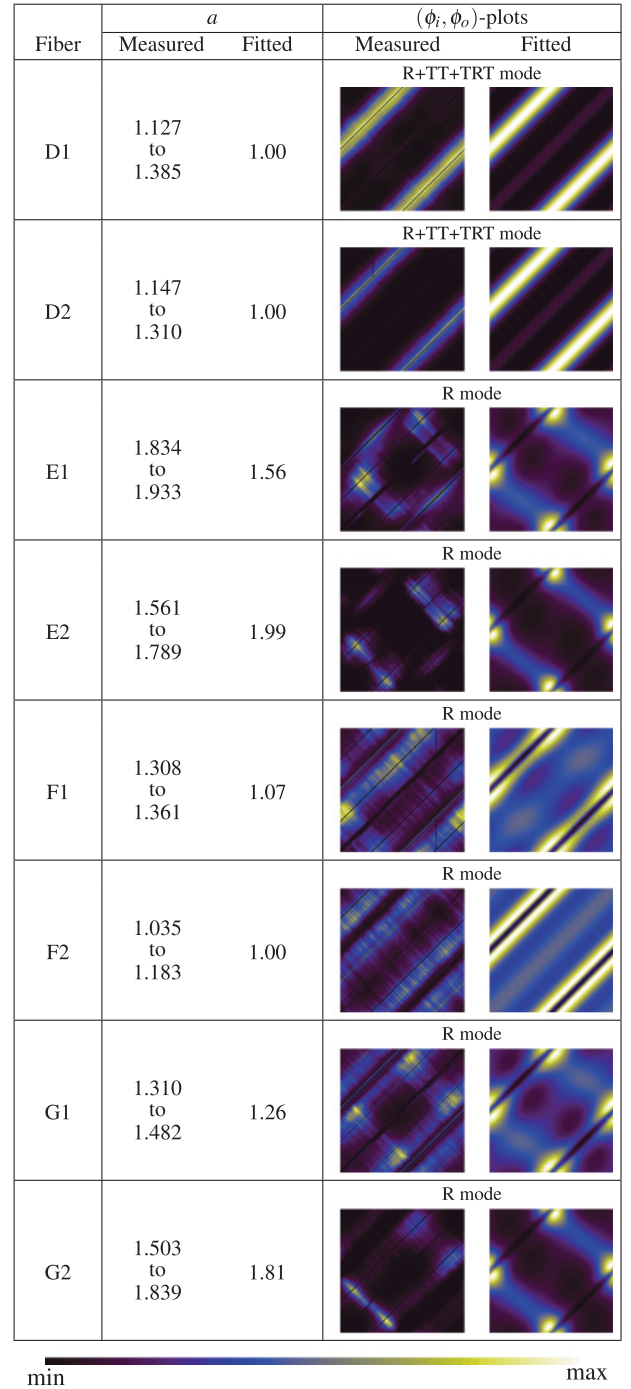


Fig. 19. Results of fitting fiber aspect ratios for Fiber D to G (those that do not show clear TRT modes) to match the captured (ϕ_i, ϕ_o) -plots. The generated plots from fitted aspect ratios contain the same features as the measured plots, but matching the aspect ratios is harder as some elliptical fibers only show features that resemble circular fibers. For fibers from Sample D, we show the sum of the first three modes because all the modes are present in all the measured (ϕ_i, ϕ_o) -plots. For other dark-colored fibers, we only show the R mode as it is the only visible mode that the model covers. All plots are for $\theta_i = 0^\circ$. Plots of other modes at other θ_i angles are given in the supplementary material. The plots' colors are not to scale with one another.

For dark-colored fibers, we could only match against the R mode. While our model can capture the two bright bands and can somewhat match the sizes of the bands, the best-fit aspect ratios can be far from the measured values, especially for Fibers E1, E2, and F2. Moreover, the generated plots generally do not look like the measured data. The two bands are not symmetric to each other as predicted, indicating that the R mode is sensitive to the asymmetry in the shape of the cross sections. We also observe dark bands around the forward scattering lines in the measured plots, but this is caused by the R mode curving away from the fixed longitudinal angle we used to create the (ϕ_i, ϕ_o) -plot, a phenomenon not accounted for by our model, and so does not indicate a disagreement in the ASF.

In conclusion, our elliptical ASF model can predict the azimuthal behavior in the TT and TRT modes well and can capture some important features of the R mode, indicating that changing from a circular to an elliptical model for the cross section can produce models much more faithful to real fibers. However, some important effects seen in the measurements, including curving of the R mode in θ_o direction and the existence of the E mode, require additional improvements in scattering models beyond upgrading the ASF to account for elliptical cross sections.

8.2 Rendering Results

To compare azimuthal scattering behavior of our models to older models, we show single scattering behavior of a planar array of vertical hair strands whose cross sections are rotated by 0° to 360° from left to right in Figure 20. We also provide, in Figure 21, renderings with full multiple scattering of hair geometries illuminated by a constant environment light source and an area light source positioned at various horizontal angles around them. The renderings were done at roughly half a megapixel resolution with 512 samples per pixel. An image typically took 15 to 30min on a 192-core compute cluster, depending on hair geometry. Videos comparing different models being lit by revolving an area source are available in the supplementary material.⁸

From the single-fiber measurements and model predictions, we know that the most important effects of eccentricity on the azimuthal distribution are (1) that the TRT highlight occurs at a larger difference in ϕ , (2) that the TRT highlight occurs over a wider range of angles, and (3) that the TRT and TT highlights vary substantially in brightness with the orientation of the fiber.

The most obvious result of these effects is that the TRT component, which is a relatively subtle effect in circular hair, produces a bright, glittery highlight that occurs at a large scattering angle. In Figure 20, this effect manifests as bright blobs produced by our elliptical models when the light angles are 39° and 66° away from the camera. Moreover, in Figure 21, elliptic fibers with aspect ratio 1.2 become much brighter at 39° , and those with aspect ratio 1.6 at 66° . These bright, colorful glints create a strong texture in the image, with an almost metallic appearance, that is not seen in circular hair. Previous models for hair scattering have never reproduced glints with any serious attempt at accuracy. (While the Marschner model does produce some brightness variation in the TT mode, the glints it produces are generally too dim to notice in real scenes.) The ability to render them correctly will enable more natural looking and realistic results.

This glint angle for aspect ratio 1.6 is confirmed by a simple experiment of illuminating the swatch that was the source of two of our hair samples (Sample 3) from the same angles (bottom row

in Figure 21). The photographs illustrate the occurrence of glints at angles between 60° and 80° . The glints in photographs do not increase the overall brightness as much as in renderings, and they occur over a wider range of angles. They show up in short segments in the photographs while whole hair strands seem to light up in the renderings. We surmise that this effect is caused by real hair having natural variation that is lacking in our models: Both eccentricity and orientation likely vary from fiber to fiber and along fibers. We randomized the orientation of each hair strand in our models at the root, but eccentricity is perfectly constant and the fibers do not twist.

A second difference can be seen in the TT component. One can observe that, in Figure 20, some fibers with our elliptical models become much brighter than surrounding fibers when the light source is at 135° from the camera. Such an effect is also present at 180° in the model with aspect ratio 1.6. It is the result of the TT component's focusing light strongly for a narrow range of orientations (that is, when one of the perpendicular blobs appears), leading to another, subtler glint effect. These TT glints have not been reported in previous measurements or modeled by any previous scattering models. However, they are hard to observe in less contrived rendering situations, as we see minimal differences between forward scattering highlights between the models in Figure 21.

When the hair appearance is dominated by the R component and by multiple scattering, the effects of eccentricity are not as dramatic as the glints. They amount to moderate change in the overall hair color because eccentricity changes the distribution of path lengths through the fibers, resulting in different colors of modes other than the R mode. This suggests that approximate methods for multiple scattering that have been developed for the circular case [Zinke et al. 2008; Shinya et al. 2010] can likely be used with elliptical fibers.

9. CONCLUSIONS

This article has presented a study of the scattering behavior of hair fibers, focused on azimuthal effects. We designed a new measurement device that enabled us to observe a more complete picture of the light scattered by a fiber, which showed that real hair fibers behave differently from what was predicted by previous scattering models. Features not previously predicted include the two bands in the R mode, the perpendicular blobs in the TT mode, the eyes in the TRT mode, and the E mode itself. Using geometric optics to model light scattered from ideal rough elliptical fibers, we learned that the features in the R, TT, and TRT modes are products of the geometry of the cross section. However, the geometric optics analysis does not predict the E mode, and we conjecture that this mode is a wave optics phenomenon. Rendered images made using a model built on the new cross-section analysis better reproduce the glittery appearance that can be observed in real hair fibers.

9.1 Limitations and Future Work

While this article does not model the E mode, light scattering behavior of small circular cylinders has been studied by physicists [van de Hulst 1957], and their work might inform us on how to model the E mode of elliptical fibers.

This article is primarily concerned with azimuthal effects, but the measured data also motivate improvements to the longitudinal component of the model. In particular, the curving of the R mode needs to be addressed, and many components become more

⁸The frame video frames are rendered with 128 samples per pixel instead of 512.




Fig. 20. Renderings of an array of vertical hair fibers whose cross sections are rotated by 0° to 360° from left to right. We compare (1) our model with aspect ratio, 1.0, 1.2, and 1.6; (2) the circular model from d'Eon et al. [2011]; and (3) the classic model from Marschner et al. [2003] with aspect ratios 1.0 and 1.2. (We do not show the Marschner model with aspect ratio parameter of 1.6 because the model was claimed to work with aspect ratio up to about $1/0.85 \approx 1.18$.) The scenes all have one-directional light source with a horizontal direction. The angle that this direction makes with the line from the eye to the center of the image plane is indicated in the leftmost column of the table. In this way, 0° indicates that the light's direction is the same as the viewing direction, and 180° indicates that the light's direction is anti-parallel to the viewing direction. We chose the tilt angles of the four modes so they separate into distinct horizontal bands whose altitudes are indicated by the labels to the right of the table. The renderings only take into account single scattering.

longitudinally blurred than the simple constant- β model predicts. The recent work of d'Eon et al. [2014] on non-separable BCSDf can potentially handle these phenomena, but it only deals with circular cross sections.

Our implementation currently requires big tables to be built for each combination of hair parameters. This makes it inconvenient to render a head of hair where fibers have different parameters or

where parameters vary along each fiber. A scheme to compactly compress the ASFs or efficient algorithms to analytically evaluate and sample them would enable rendering of more complex and realistic hair.

Last, our measurement techniques can be used with any type of fiber. As a result, they can be used to study fibers such as cloth fibers and animal fur.



Fig. 21. Renderings with full multiple scattering of three pieces of hair geometry lit by an area source rotating around them. The leftmost column contains the angle, in the horizontal plane, that the light makes with the camera’s view direction. Observe that our model with $a = 1.2$ becomes brighter than other models at 39° and so does our model with $a = 1.6$ at 66° . The final column shows photographs of Sample C for reference; note glints appearing at 66° . The supplementary material contains the parameters for the models and the uncropped renderings. (While we attempted to find absorption parameters that would match the overall color of Sample C, the real hair is considerably more absorbing.)

APPENDIX

Let $P(x) = p_0 + p_1x + \dots + p_kx^k$. Then,

$$\int P(x)g(x; \mu, \sigma) dx = \frac{A}{2} \operatorname{erf}\left(\frac{x - \mu}{\sqrt{2}\sigma}\right) - \sigma^2 B(x)g(x; \mu, \sigma) + C,$$

where $B(x) = b_0 + b_1x + \dots + b_{k-1}x^{k-1}$,

$$b_j = \begin{cases} p_k, & j = k - 1, \\ p_{k-1} + \mu b_{k-1}, & j = k - 2, \\ p_{j+1} + \mu b_{j+1} + (j + 2)\sigma^2 b_{j+2}, & 0 \leq j < k - 2, \end{cases}$$

and $A = p_0 + \mu b_0 + \sigma^2 b_1$.

The formula can be verified by simply taking the derivative. We also provide detailed derivations in the supplementary material.

ACM Transactions on Graphics, Vol. 36, No. 2, Article 13, Publication date: March 2017.

ACKNOWLEDGMENTS

The authors thank Saba Weatherspoon, Siman Peng, Kavita Bala, and Uniliver R&D for donating hair fiber samples. We are also grateful to Albert Liu for the Caliber geometric calibration system, to Cem Yuksel for the hair models, and to Wenzel Jakob for the Mitsuba physically based renderer, which we used to initially implement our model.

REFERENCES

- Charles L. Adler, James A. Lock, and Bradley R. Stone. 1998. Rainbow scattering by a cylinder with a nearly elliptical cross section. *Appl. Opt.* 37, 9 (Mar. 1998), 1540–1550. DOI: <http://dx.doi.org/10.1364/AO.37.001540>
- Helen K. Bustard and Robin W. Smith. 1991. Investigation into the scattering of light by human hair. *Appl. Opt.* 30, 24 (Aug. 1991), 3485–3491. DOI: <http://dx.doi.org/10.1364/AO.30.003485>

- D. Cline, A. Razdan, and P. Wonka. 2009. A comparison of tabular PDF inversion methods. *Comput. Graph. Forum* (2009). DOI : <http://dx.doi.org/10.1111/j.1467-8659.2008.01197.x>
- Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. 2011. An energy-conserving hair reflectance model. *Comput. Graph. Forum* 30, 4 (2011), 1181–1187. DOI : <http://dx.doi.org/10.1111/j.1467-8659.2011.01976.x>
- Eugene d'Eon, Steve Marschner, and Johannes Hanika. 2014. A fiber scattering model with non-separable lobes. In *ACM SIGGRAPH 2014 Talks (SIGGRAPH'14)*. ACM, New York, NY, Article 46, 1 pages. DOI : <http://dx.doi.org/10.1145/2614106.2614161>
- Andrew Fitzgibbon, Maurizio Pilu, and Robert B. Fisher. 1999. Direct least square fitting of ellipses. *IEEE Trans. Pattern Anal. Mach. Intell.* 21, 5 (May 1999), 476–480. DOI : <http://dx.doi.org/10.1109/34.765658>
- Dan B. Goldman. 1997. Fake fur rendering. In *SIGGRAPH '97 Proceedings*. 127–134. DOI : <http://dx.doi.org/10.1145/258734.258807>
- Xiao D. He, Kenneth E. Torrance, Francois X. Sillion, and Donald P. Greenberg. 1991. A comprehensive physical model for light reflection. *SIGGRAPH '91* 25, 4 (July 1991), 175–186.
- Christophe Hery and Ravi Ramamoorthi. 2012. Importance sampling of reflection from hair fibers. *J. Comput. Graph. Tech.* 1, 1 (22 June 2012), 1–17.
- J. T. Kajiya and T. L. Kay. 1989. Rendering fur with three dimensional textures. In *SIGGRAPH '89 Proceedings*. 271–280. DOI : <http://dx.doi.org/10.1145/74333.74361>
- Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. (2015). To appear in *ACM Transactions Graphics*.
- Tae-Yong Kim. 2002. *Modeling, Rendering and Animating Human Hair*. Ph.D. Dissertation. Los Angeles, CA. UMI Order Number: AAI 3093778.
- Albert Liu, Steve Marschner, and Noah Snavely. 2015. Caliber: Camera localization and calibration using rigidity constraints. *Int. J. Comput. Vis.* (2015), 1–21. DOI : <http://dx.doi.org/10.1007/s11263-015-0866-1>
- D. Marcuse. 1974. Light scattering from elliptical fibers. *Appl. Opt.* 13, 8 (Aug. 1974), 1903–1905. DOI : <http://dx.doi.org/10.1364/AO.13.001903>
- Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light scattering from human hair fibers. In *SIGGRAPH 2003 Proceedings*. 780–791. DOI : <http://dx.doi.org/10.1145/1201775.882345>
- Shinji Ogaki, Yusuke Tokuyoshi, and Sebastian Schoellhammer. 2010. An empirical fur shader. In *ACM SIGGRAPH ASIA 2010 Sketches (SA'10)*. Article 16. DOI : <http://dx.doi.org/10.1145/1899950.1899966>
- J. Ou, F. Xie, P. Krishnamachari, and F. Pellacini. 2012. ISHair: Importance sampling for hair scattering. In *Computer Graphics Forum*, Vol. 31. 1537–1545.
- Clarence R. Robbins. 1994. *Chemical and Physical Behavior of Human Hair* (3rd ed.). Springer-Verlag, New York.
- Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.* 32, 2, Article 14 (April 2013), 12 pages. DOI : <http://dx.doi.org/10.1145/2451236.2451240>
- Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. 2010. An artist friendly hair shading system. In *SIGGRAPH 2010 Proceedings*. Article 56. DOI : <http://dx.doi.org/10.1145/1833349.1778793>
- M. Shinya, M. Shiraishi, Y. Dobashi, K. Iwasaki, and T. Nishita. 2010. A simplified plane-parallel scattering model and its application to hair rendering. In *2010 18th Pacific Conference on Computer Graphics and Applications*. IEEE, 85–92.
- Robert F. Stamm, Mario L. Garcia, and Judith J. Fuchs. 1977. The optical properties of human hair I. Fundamental considerations and goniophotometer curves. *J. Soc. Cosmet. Chem.* 28 (Sep. 1977), 571–599.
- R. A. R. Tricker. 1970. *Introduction to Meteorological Optics*. Mills & Boon, London.
- H. C. van de Hulst. 1957. *Light Scattering by Small Particles*. Dover Publications.
- Nicky van Foreest. 2012. Fitting an Ellipse to a Set of Data Points. Retrieved September 2012 from <http://nicky.vanforeest.com/misc/fitEllipse/fitEllipse.html>.
- Kun Xu, Li-Qian Ma, Bo Ren, Rui Wang, and Shi-Min Hu. 2011. Interactive hair rendering and appearance editing under environment lighting. In *SIGGRAPH Asia 2011 Proceedings (SA'11)*. Article 173. DOI : <http://dx.doi.org/10.1145/2024156.2024207>
- Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. Physically-accurate fur reflectance: Modeling, measurement and rendering. *ACM Trans. Graph. (Proceedings of SIGGRAPH Asia 2015)* 34, 6 (2015), 185.
- Arno Zinke. 2008. *Photo-Realistic Rendering of Fiber Assemblies*. Ph.D. Dissertation. Bonn, Germany.
- Arno Zinke, Tomas Lay Herrera, Anton Andriyenko, Martin Rump, Andreas Weber, and Reinhard Klein. 2009. A practical approach for photometric acquisition of hair color. *ACM Trans. Graph. (Proceedings of SIGGRAPH Asia)* 28, 5 (2009), 165. DOI : <http://dx.doi.org/10.1145/1618452.1618511>
- Arno Zinke and Andreas Weber. 2007. Light scattering from filaments. *IEEE Trans. Visualiz. Comput. Graph.* 13, 2 (2007), 342–356. DOI : <http://dx.doi.org/10.1109/TVCG.2007.43>
- A. Zinke, C. Yuksel, A. Weber, and J. Keyser. 2008. Dual scattering approximation for fast multiple scattering in hair. *ACM Trans. Graph.* 27, 3 (2008), 1–10.

Received April 2015; revised October 2016; accepted November 2016