

A Practical Wave Optics Reflection Model for Hair and Fur

MENGQI (MANDY) XIA, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

BRUCE WALTER, Cornell University, USA

CHRISTOPHE HERY, Meta Reality Labs, USA

OLIVIER MAURY, Meta Reality Labs, USA

ERIC MICHIELSSEN, University of Michigan, USA

STEVE MARSCHNER, Cornell University, USA

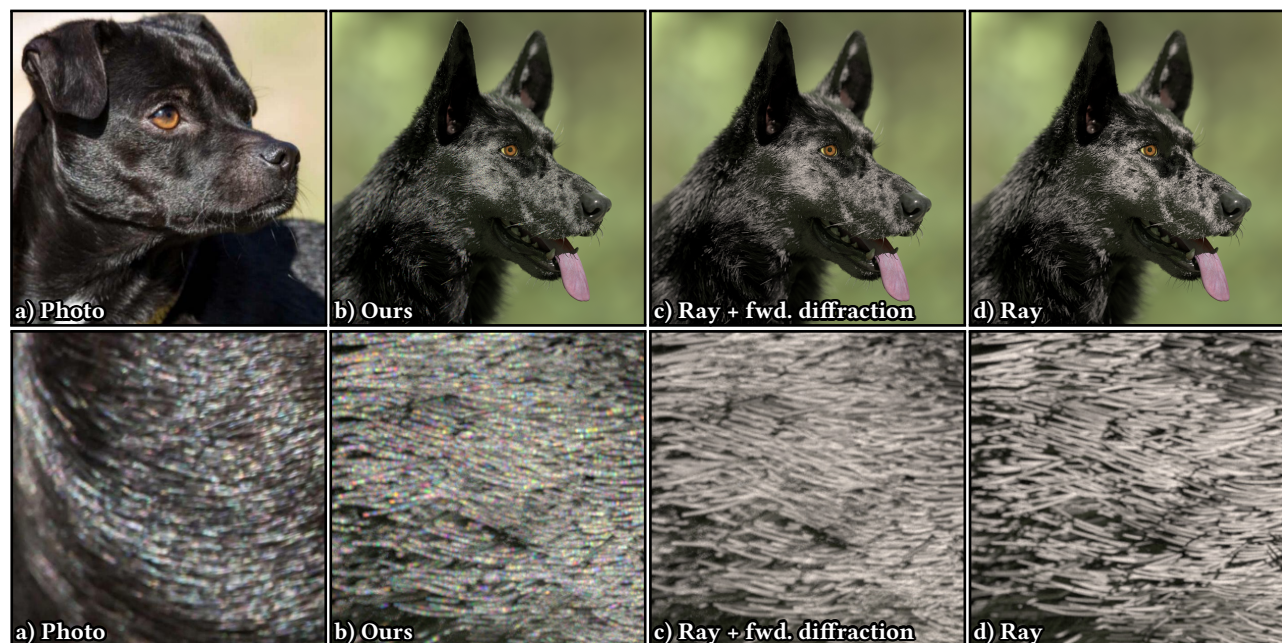


Fig. 1. Rendering of black dog fur with aperture blur, compared to a photograph. Our model is able to produce realistic colorful glints, which were missing from all previous models. Our model includes forward diffraction around the fiber. While not contributing to the colorful glints in this scene, the diffraction component produces softer highlights on the fur compared to the ray-based model [Pharr et al. 2016]. The macroscopic appearance of the dog fur generated from our model contains a hue shift, which is also observed in the photograph. Photograph by Dietrich Zawischa.

Traditional fiber scattering models, based on ray optics, are missing some important visual aspects of fiber appearance. Previous work [Xia et al. 2020] on wave scattering from ideal extrusions demonstrated that diffraction produces strong forward scattering and colorful effects that are missing from ray-based models. However, that work was unable to include some important surface characteristics such as surface roughness and tilted cuticle scales, which are known to be important for fiber appearance. In this work, we take an important step to study wave effects from rough fibers with

Authors' addresses: Mengqi (Mandy) Xia, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland; Bruce Walter, Cornell University, USA; Christophe Hery, Meta Reality Labs, USA; Olivier Maury, Meta Reality Labs, USA; Eric Michielssen, University of Michigan, USA; Steve Marschner, Cornell University, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2023/8-ART1 \$15.00

<https://doi.org/10.1145/3592446>

arbitrary 3D microgeometry. While the full-wave simulation of realistic 3D fibers remains intractable, we developed a 3D wave optics simulator based on a physical optics approximation, using a GPU-based hierarchical algorithm to greatly accelerate the calculation. It simulates surface reflection and diffractive scattering, which are present in all fibers and typically dominate for darkly pigmented fibers. The simulation provides a detailed picture of first order scattering, but it is not practical to use for production rendering as this would require tabulation per fiber geometry. To practically handle geometry variations in the scene, we propose a model based on wavelet noise, capturing the important statistical features in the simulation results that are relevant for rendering. Both our simulation and practical model show similar granular patterns to those observed in optical measurement. Our compact noise model can be easily combined with existing scattering models to render hair and fur of various colors, introducing visually important colorful glints that were missing from all previous models.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**.

ACM Reference Format:

Mengqi (Mandy) Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. 2023. A Practical Wave Optics Reflection

Model for Hair and Fur. *ACM Trans. Graph.* 42, 4, Article 1 (August 2023), 15 pages. <https://doi.org/10.1145/3592446>

1 INTRODUCTION

Light scattering from fibers is important for rendering human hair and animal fur for visual effects and other applications. Many successful models for fiber scattering have been built using geometric optics, but rays cannot tell the whole story.

In recent work, Xia et al. [2020] turned to wave optics in search of a more accurate description of the interaction between light and fibers, using Boundary Element Method (BEM) calculations to predict scattering by cylinders of any cross-section. Their simulations agreed with the traditional ray models in some respects but showed important differences that are borne out in practice. Fibers exhibit strong forward scattering due to diffraction, scattering considerably more light than is predicted by ray models; small fibers have wavelength-dependent scattering that produces strong colors; and singularities like caustics are softened by the limited bandwidth of the scattered wave field.

However, Xia et al. relied on a crucial assumption to make the BEM simulations tractable: fibers were not allowed to deviate from ideal extrusions. It has long been known [Marschner et al. 2003] that surface structure, including the asymmetry due to cuticle scales, is responsible for important effects in hair appearance. Xia et al. left the implications of wave optics for these effects unknown.

There is evidence for effects of 3D structure that are missed by ray models. Measurements of fiber scattering have shown sharp features not predicted by ray reflection from rough fibers, likely explained by diffraction. Hair and fur also exhibit glinty color noise under sharp illumination. For example, in Figure 1, black dog fur fibers appear colorful when observed closely under directional light; they can exhibit a slight hue shift even from a distance.

To model these effects and produce the first fully 3D wave scattering model for rough fibers, we developed a 3D wave simulator based on a physical optics (PO) approximation to predict surface reflection and diffraction from rough fibers. Our simulator is much more efficient than a full-wave solver because PO assumes single-scattering, which eliminates the linear system solve that is the main bottleneck. We accelerate the remaining computation with an octree-based GPU implementation. Our simulator is general and can handle arbitrary 3D geometry while achieving high accuracy in the low curvature regime.

We present simulation results for explicit fiber geometry models that capture the overall elliptical shape of fur and hair fibers and also include random micro-scale details, including roughness and cuticle scales. The results of these simulations show two major differences when compared to the state-of-the-art ray-optics models: a forward diffraction lobe that was also described by Xia et al. [2020], and intricate and colorful random structures in the specular reflection lobe. This random structure is very different from the smooth scattering patterns that existing models produce. We compare our simulation result against a measurement of speckle on a real hair fiber.

The predictions of the model are very useful in diagnosing what is missing from existing models, but tabulating these patterns for rendering is impractical. We analyze the highlight noise in our simulations using the theory of optical speckle to characterize them in




Fig. 2. This figure shows a scanning electron micrograph (SEM) with surface details of hair [Nanjundaswamy 2016], as well as four fiber instances that we generated with varying rough surface parameters.

terms of the statistics of a random process. We then propose a practical model suitable for rendering that starts with an existing fiber scattering model, adds a diffraction lobe based on elementary diffraction theory, and modulates the reflection highlight by a random process designed to resemble optical speckle and generated using procedural noise tuned to fit the speckle statistics in our simulations. We validate our noise-based representation against the simulation results. We render a variety of hair and fur scenes using our new practical model, producing realistic colored glints that resemble those seen in real-world hair of different colors under sunlight. Our source code is available at <https://github.com/mandyxmq/WaveFiber3d>.

2 RELATED WORK

We will first review the previous ray- and wave-based fiber scattering models in computer graphics, relevant computational electromagnetics tools, and related surface models based on physical optics. We will also review the previous studies on speckle and noise in our community that are related to our effort in fitting noise functions to speckle patterns in fiber scattering.

Ray based fiber models. Most of the hair and fiber models used in computer graphics are based on geometric optics. The widely used model from Marschner et al. [2003] analyzed ray paths in dielectric cylinders and cones, and splits the scattering into R, TT and TRT modes which are represented as separable products of azimuthal and longitudinal functions. Numerous extensions and improvements to this model have been proposed such as adding a diffuse component [Zinke et al. 2009], artist-friendly parameterizations [Sadeghi et al. 2010], improved energy conservation [d'Eon et al. 2011], non-separable representations [d'Eon et al. 2014; Huang et al. 2022], and adaptations for production rendering [Chiang et al. 2016]. Generalizations for other fiber types have also been proposed such as for elliptical hairs [Khungurn and Marschner 2017], animal hairs with interior medullas [Yan et al. 2017, 2015], and textile fibers with more general cross sections [Aliaga et al. 2017].

Wave optics fiber models. Linder [2014] used an analytic solution to study cylindrical fibers with perfectly circular cross-sections. Xia et al. [2020] conducted 2D wave simulation on cylinders with arbitrary cross-sections and demonstrated some important differences as compared to geometric optics based fiber models. However, they

assume perfect extrusions and are thus unable to handle geometric features such as general surface roughness and tilted cuticle scales, which are the main focus of this paper. Benamira and Pattanaik [2021] proposed a faster hybrid model that uses wave optics solutions only for the forward scattering diffraction component and geometric optics for the rest of the fiber scattering distribution. The ray plus the forward-diffraction component in our model is similar to theirs, but we also take into account the dependence on incident longitudinal angle.

Computational Electromagnetics. Computational electromagnetics (CEM) is a discipline that develops computational methods to understand electromagnetic phenomena, and it applies to optics since light is an electromagnetic wave. The most frequently used algorithms in CEM are the finite-difference time-domain (FDTD) method [Kane Yee 1966; Taflov et al. 2005], the finite element method (FEM) [Jin 2015], and the boundary element method (BEM), which is also referred to as the method of moments (MoM) [Gibson 2021; Huddleston et al. 1986; Wu and Tsai 1977]. Although fast algorithms such as the multilevel fast multipole algorithm (MLFMA) [Song et al. 1997] exist, full 3D wave simulation is still very expensive for hair and fur fiber simulations. Therefore, we propose to apply the physical optics approximation and handle the important surface reflection and diffraction components first.

Physical optics planar models. Physical optics approximations, such as Beckmann-Kirchhoff [Beckmann and Spizzichino 1987] and Harvey-Shack [Harvey 1979], have been widely used in graphics to compute scattering from a variety of roughly planar surfaces that can be well modeled as height fields. Examples include Gaussian random [He et al. 1991; Kajiya 1985], stationary periodic [Stam 1999], tabulated [Dong et al. 2015], and scratched [Werner et al. 2017] surfaces. Previous work combines one-bounce Kirchhoff scalar diffraction theory with path tracing for secondary bounces [Falster et al. 2020]. Recent advances in the treatment of diffraction from rough surfaces have been introduced by Krywonos et al. [Krywonos 2006; Krywonos et al. 2011]. Holzschuch and Pacanowski [2017] introduce a two-scale microfacet model that combines reflection and diffraction. Yan et al. [2018] utilize PO to render specular microgeometry on rough surfaces. Clausen et al. [2023] examine rough surface scattering through measurement comparisons. However, unlike the planar case, fiber geometry is a closed surface that only partially blocks the incident wave. Thus, it requires handling additional effects such as forward diffractive scattering and extensive shadowing that could be safely neglected in these planar models. Therefore, while the physical optics approximations are similar, applying them to our fiber geometry required developing new methods.

Speckle in rendering. Speckle is a granular structure that appears in images and diffraction patterns produced by optically rough objects. Speckle theory [Goodman 2007] studies the statistics of these patterns and has uses in imaging applications, such as tissue imaging, motion tracking, and non-line-of-sight imaging. Previously, in the computer graphics community, researchers have applied Monte Carlo methods to simulate speckle from volume scattering [Bar et al. 2019, 2020]. However, their models are for homogeneous media and do not apply to fibers. Steinberg and Yan [2022] rendered speckle




Fig. 3. An illustration of the commonly adopted longitudinal-azimuthal parameterization for fiber scattering models. Each of the directions ω_i and ω_r in 3D is parameterized using the polar angle θ , defined as the angle between ω and the plane perpendicular to the cylinder axis, and the azimuthal angle ϕ , defined in that plane.

from planar rough surfaces. We observed that the statistics of the fiber speckles deviate from those of planar surfaces. Our contribution is to faithfully capture the important statistics of fiber speckle patterns and develop a practical model for rendering.

Procedural noise. Procedural noise functions are widely used in computer graphics because they offer many advantages. They are often very fast to evaluate, adding complex and intricate details to the rendering on-the-fly, and they have low memory usage [Lagae et al. 2010]. These noise functions have many applications in synthesizing complex textures for rendering natural phenomena such as rocks, clouds, and ocean waves. One category of procedural noise functions is lattice gradient noise, such as Perlin noise [Perlin 1985, 2002], simplex noise [Olano et al. 2002], flow noise [Perlin and Neyret 2001], and curl noise [Bridson et al. 2007]. These methods generate noise by interpolating or convolving random values or gradients defined at the integer grid points. Explicit noise functions pre-generate noise and store a tile for rendering; variation is achieved by offsetting the noise tile. They have also shown great potential in reproducing complex textures for rendering. Wavelet noise [Cook and DeRose 2005] is an explicit noise that is almost perfectly band-limited. Moreover, the statistical distribution of wavelet noise can be calculated and controlled, which makes it a great fit for our problem. We leverage wavelet noise to construct a practical model that can handle fiber micro-geometry variations.

3 OVERVIEW AND BACKGROUND

In this section, we will provide an overview of how we construct a practical wave optics fiber scattering model and provide background knowledge on the key parts of the pipeline.

The model construction works as follows: We first create 3D micro-geometry of the fibers, inspired by SEM images of hair and fur. We then develop an efficient 3D wave simulation (Section 4) to predict surface reflection and diffraction from rough fibers. Next, we apply speckle theory to study the statistics of the scattering patterns and construct a compact noise representation of the speckle patterns (Section 5). We compare our simulation result with a measurement, which also helps us infer reasonable fiber parameters such as size, cuticle angle, and surface roughness. Finally, we integrate our model into a modern rendering system, with related implementation details in Section 7 and results in Section 8.

3.1 Fiber scattering models

Fiber scattering models describe how light interacts with a single fiber using the Bidirectional Curve Scattering Distribution Function (BCSDF). These models are essential for rendering hair, fur, and cloth fibers. Similar to the BSDF, the BCSDF describes outgoing radiance L_r as an integration of incident radiance L_i multiplied by the BCSDF S :

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

We will write the BCSDF in spherical coordinates as $S(\theta_i, \theta_r, \phi_i, \phi_r, \lambda)$ using the angles illustrated in Figure 3; λ denotes wavelength. Most fiber scattering models write the BCSDF as a sum of reflective and transmissive modes S_p

$$S(\theta_i, \theta_r, \phi_i, \phi_r, \lambda) = \sum_{p=0}^{\infty} S_p(\theta_i, \theta_r, \phi_i, \phi_r, \lambda). \quad (2)$$

The first scattering mode S_0 describes the surface reflection, which in previous models often either represents reflection from a smooth fiber [Marschner et al. 2003] or can be seen as the surface reflection statistically averaged over rough fiber instances [Chiang et al. 2016; Huang et al. 2022]. We use our 3D wave simulator to simulate the surface reflection and diffraction lobes in fiber scattering to more accurately estimate S_0 . Following the Marschner hair model [Marschner et al. 2003], many fiber scattering models assume each mode S_p to be factored into a longitudinal function M_p and an azimuthal function N_p :

$$S_p(\theta_i, \theta_r, \phi_i, \phi_r, \lambda) = M_p(\theta_i, \theta_r) N_p(\theta_i, \phi_i, \phi_r, \lambda). \quad (3)$$

However our simulation does not use this approximation and provides much more accurate modeling of S_0 .

We develop a practical fiber scattering model, which captures the statistical properties in the simulation of the first scattering mode. Our simulator computes the reflection, denoted as $S_{0,\text{sim}}$, from particular instances of the hair geometry, but these can be expensive to generate and store. For practical rendering, our idea is to approximate $S_{0,\text{sim}}$ as a product of the average component $S_{0,\text{avg}}$ and a noise component $f(\theta_h, \phi_h, \lambda)$.

$$S_{0,\text{sim}}(\theta_i, \theta_r, \phi_i, \phi_r, \lambda) \approx S_{0,\text{avg}}(\theta_i, \theta_r, \phi_i, \phi_r, \lambda) f(\theta_h, \phi_h, \lambda). \quad (4)$$

$S_{0,\text{avg}}$ can be seen as averaged over fiber instances whose surface roughness follows the same random distribution. The noise component $f(\theta_h, \phi_h, \lambda)$ represents the ratio between the first scattering mode of a specific fiber instance and that of the averaged value. θ_h and ϕ_h are the longitudinal and azimuthal angles for the half vector. We can interpret $S_{0,\text{avg}}(\theta_i, \theta_r, \phi_i, \phi_r, \lambda) f(\theta_h, \phi_h, \lambda)$ as drawing a specific example from a rough fiber distribution and evaluating its first scattering mode. We denote this component as $S_{0,\text{prac}}$ and it is our new practical model. We can combine $S_{0,\text{prac}}$ with existing models to account for all full scattering orders:

$$\begin{aligned} S_{\text{prac}}(\theta_i, \theta_r, \phi_i, \phi_r, \lambda) \\ = S_{0,\text{prac}}(\theta_i, \theta_r, \phi_i, \phi_r, \lambda) + \sum_{p=1}^{\infty} S_p(\theta_i, \theta_r, \phi_i, \phi_r, \lambda). \end{aligned} \quad (5)$$

Compared to the previous wave optics fiber scattering model [Xia et al. 2020], our simulation considers cuticles and other rough surface

details on fibers, which are important for hair and fur. Moreover, our practical fiber scattering model is much more concise than the tabulation in the previous wave optics fiber work, allowing us to handle variations in geometry in the scene.

3.2 Speckle theory

Fiber surfaces are optically rough, giving rise to random fluctuations in the scattered fields, which superpose and form speckles. Speckles appear when random phasors are summed together. A random phasor sum of N terms can be written as [Goodman 2007].

$$\mathbf{A} = \frac{1}{\sqrt{N}} \sum_{n=1}^N \mathbf{a}_n = \frac{1}{\sqrt{N}} \sum_{n=1}^N a_n e^{i\phi_n},$$

where \mathbf{A} represents the resultant phasor, \mathbf{a}_n represents the n th component phasor in the sum, and a_n and ϕ_n are the amplitude and phase of \mathbf{a}_n . One important case is fully developed speckle, where 1) the amplitudes and phases of different phasors are statistically independent; 2) the amplitude and phase of one phasor are independent of each other; 3) the phases are uniformly distributed on the interval $[-\pi, \pi]$. According to our simulation results, the fiber scattering problem satisfies these assumptions well, as long as the illumination beam size is sufficiently large compared to the correlation length of the rough surface. Thus, we apply the theory of fully developed speckle to guide us in developing a practical model.

4 WAVE SIMULATION WITH 3D FIBER MICROGEOMETRY

Since the micro-geometry of the fiber is at a scale comparable to visible light wavelengths, we need to consider the electromagnetic wave nature of light and conduct wave optics simulations to accurately predict light scattering from rough fibers. In this section, we introduce our 3D wave simulator that predicts electromagnetic wave scattering from rough fibers and explain how we make the simulation efficient.

4.1 Wave optics

In wave optics simulation, we compute electric and magnetic fields and analyze how the scattering object affects them. Fields with sinusoidal time variation are called time-harmonic fields, and their mathematical analysis can be simplified by using complex quantities. Phasors of the electric field \mathbf{E} and the magnetic field \mathbf{H} are defined as:

$$\mathbf{E}_{\text{inst}} = \text{Re}(\mathbf{E}e^{j\omega t}), \quad \mathbf{H}_{\text{inst}} = \text{Re}(\mathbf{H}e^{j\omega t}). \quad (6)$$

\mathbf{E}_{inst} and \mathbf{H}_{inst} are the instantaneous electric and magnetic fields, and ω is the angular frequency. In what follows, we assume time-harmonic fields and suppress the time dependence $e^{j\omega t}$ unless specified. The wave fields obey the time-harmonic Maxwell's equations, which relate corresponding currents and fields:

$$\begin{aligned} \nabla \times \mathbf{E} &= -\mathbf{M} - j\omega\mu\mathbf{H} \\ \nabla \times \mathbf{H} &= \mathbf{J} + j\omega\epsilon\mathbf{E}. \end{aligned} \quad (7)$$

ϵ and μ are the permittivity and permeability. Here \mathbf{J} and \mathbf{M} are time-harmonic electric and magnetic current densities. In our problems,

they are fictitious currents that make the problem easier to solve mathematically.

The object is illuminated by an incident wave, and the incident electric and magnetic fields are denoted as \mathbf{E}_i and \mathbf{H}_i . The presence of the scatterer alters the fields, and we call the resulting fields the *total fields*. We denote the total fields outside the object as \mathbf{E}_1 and \mathbf{H}_1 . We can further write \mathbf{E}_1 and \mathbf{H}_1 as the sums of incident fields and the *scattered fields* \mathbf{E}_s and \mathbf{H}_s .

$$\mathbf{E}_1 = \mathbf{E}_i + \mathbf{E}_s, \quad \mathbf{H}_1 = \mathbf{H}_i + \mathbf{H}_s. \quad (8)$$

The scattered fields propagate outward from the scatterer, and we can compute the energy flow from them. The scattered fields will be the key to computing scattering functions.

Full-wave simulation in 3D is very expensive, and it is not practical to apply it to compute fiber scattering of real hair sizes without fast algorithms. To conduct a full-wave simulation, we need to discretize the object into a mesh, and the rule of thumb for resolution is 10 elements per wavelength, which corresponds to millions of mesh elements for even a very short hair segment that is tens of microns long. There are two steps in a full 3D wave simulation: The first step is to solve a linear system that describes the interaction between all surface points. The size of the matrix scales quadratically with respect to the number of elements in the mesh, resulting in a huge matrix in the brute force wave simulation in 3D, even when simulating a short hair segment. From the first step of the calculation, we obtain surface currents \mathbf{J} and \mathbf{M} , which we can think of as secondary sources radiating the scattered wave. The second step is to compute the scattered waves (in the far field in our application) from these surface currents. The bottleneck for full 3D wave simulation at the size of hair and fur fibers lies in the first step. We propose using a physical optics approximation (PO) (Section 4.2) and computing surface reflection with single scattering. PO makes the surface current computation so efficient that computing the far-field radiation becomes the new bottleneck of our problem. We then develop an octree-based algorithm to accelerate the far-field scattering computation (Section 4.3). For a more complete treatment of the topic of electromagnetic wave scattering from objects, see [Bondeson et al. 2012; Huddleston et al. 1986; Poggio and Miller 1970; Wu and Tsai 1977].

4.2 Physical Optics Approximation

Physical optics assumes single scattering, so that the surface currents can be locally computed based on the incident fields and the material properties. PO also assumes that the local geometry is flat enough to be approximated locally as a plane. Thus, we can relate the reflected field and the incident field via surface reflection coefficients in the flat surface reflection computation. Although PO ignores multiple scattering within a single fiber geometry, it allows us to efficiently analyze surface reflection and diffraction from fibers. Our PO simulator is general and works for arbitrary 3D objects. It provides accurate predictions for most scattering directions of hair fibers as the roughness of these fibers is often relatively low, where some measured data can be found in [LaTorre and Bhushan 2005]. The algorithm is illustrated in Figure 4.

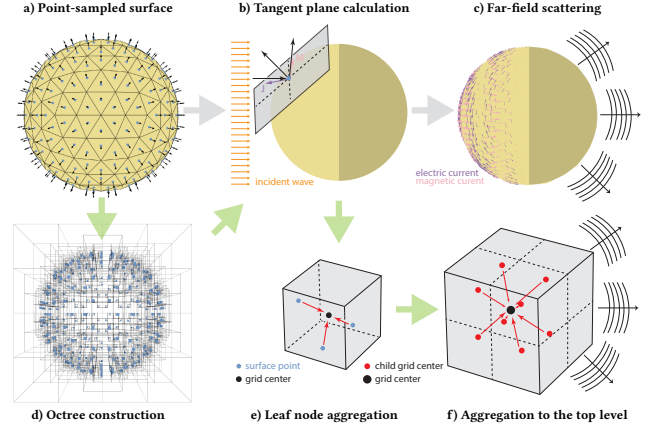


Fig. 4. Illustration of the PO algorithm without the octree structure (following the grey arrows) and with the octree structure (following the green arrows).

Surface current calculation. The surface geometry of the object is represented as a densely point-sampled surface, where each point corresponds to an element in a mesh representation. For each sampled point \mathbf{r}' , we also store the face normal $\mathbf{n}(\mathbf{r}')$ and the area of the element. At each sampled point \mathbf{r}' , we compute the surface currents by considering a virtual tangent plane perpendicular to $\mathbf{n}(\mathbf{r}')$. We compute the surface currents $\mathbf{J}(\mathbf{r}')$ and $\mathbf{M}(\mathbf{r}')$ using the analytic solution for reflection from a flat surface. Specifically, we first build a local polarization frame using the incoming wave direction \mathbf{e}_i and the normal $\mathbf{n}(\mathbf{r}')$, and decompose the incident fields into the sum of parallel and perpendicular polarized fields $\mathbf{E}_i = \mathbf{E}_i^p + \mathbf{E}_i^s$. Then the reflected field and the total field at the boundary can be computed via

$$\begin{aligned} \mathbf{E}_r &= \mathbf{E}_r^p + \mathbf{E}_r^s = F^p \mathbf{E}_i^p + F^s \mathbf{E}_i^s, \\ \mathbf{E}_1 &= \mathbf{E}_i + \mathbf{E}_r. \end{aligned} \quad (9)$$

In the above equations, F^p and F^s are the reflection coefficients in Fresnel's equations for parallel and perpendicular polarization. The reflected field \mathbf{E}_r here is an approximation of the actual scattered field at the surface based on the assumption that the current depends only on the incident field but not on scattering from other parts of the object. Similarly, we can compute the scattered and total magnetic fields. By applying the known relationship between the currents and the total fields, we can compute the surface currents.

$$\mathbf{M} = -\mathbf{n} \times \mathbf{E}_1, \quad \mathbf{J} = \mathbf{n} \times \mathbf{H}_1. \quad (10)$$

The incident field can be arbitrary in theory. In practice, for fiber scattering, we use a Gaussian-windowed plane wave. The amplitude of the field follows a Gaussian distribution perpendicular to the incident wave propagation direction and in the incident plane defined by the propagation direction and the fiber axis.

Far-field radiation in 3D. After computing the surface currents, we effectively transform the original scattering problem into a radiation problem using Huygens's principle: The surface currents serve as the equivalent sources to the original scattering problem and radiate the same scattered fields that we aim to compute. We can mathematically write down the radiation from the surface currents \mathbf{J} and \mathbf{M} . In the




Fig. 5. We compare our model with [Xia et al. 2020] and a photograph. We demonstrate that our model is able to handle the cuticle tilt and the highlight shifts downwards along the fiber when the cuticle angle changes from 0° to -5° , while [Xia et al. 2020] does not handle cuticles, thus no shift in the highlight. The two models produce a similar overall highlight for the no cuticle case, but our model additionally produces colorful glints on the hair. A similar color effect is observed in a photo of human hair under sunlight.

far field, the scattered electric field at point \mathbf{r} is [Gibson 2021]:

$$\mathbf{E}_s(\mathbf{r}) = j\omega\mu_0 \frac{e^{-jk_0R}}{4\pi R} \hat{\mathbf{r}} \times \int_{\Gamma} \left[\hat{\mathbf{r}} \times \mathbf{J}(\mathbf{r}') + \frac{1}{Z_0} \mathbf{M}(\mathbf{r}') \right] e^{jk_0\mathbf{r}' \cdot \hat{\mathbf{r}}} d\mathbf{r}', \quad (11)$$

where k_0 is the wave number, $\hat{\mathbf{r}}$ is the scattering direction, $R = |\mathbf{r}|$, Γ is the surface of the scatterer, \mathbf{r}' is a point on the surface, and $Z_0 = \sqrt{\frac{\mu_0}{\epsilon_0}}$ is the free space impedance. \mathbf{H}_s can be easily computed once we know \mathbf{E}_s , since waves propagate radially and locally follow a planar wavefront in the far field [Jackson 2021]. We define $\mathbf{E}_s^{\text{far}}(\hat{\mathbf{r}})$ and $\mathbf{H}_s^{\text{far}}(\hat{\mathbf{r}})$ by

$$\mathbf{E}_s(\mathbf{r}) = \frac{e^{-jk_0R}}{R} \mathbf{E}_s^{\text{far}}(\hat{\mathbf{r}}), \quad \mathbf{H}_s(\mathbf{r}) = \frac{e^{-jk_0R}}{R} \mathbf{H}_s^{\text{far}}(\hat{\mathbf{r}}), \quad (12)$$

where $\mathbf{E}_s^{\text{far}}$ and $\mathbf{H}_s^{\text{far}}(\hat{\mathbf{r}})$ only depend on the scattering direction $\hat{\mathbf{r}}$ but not R .

We compute this integral (Equation 11) numerically using the pointwise current computed from the local tangent plane calculation. For each incident direction and wavelength, to compute the scattering distribution for all outgoing directions, the brute force calculation requires computing the contribution from each point on the surface to each outgoing direction. Assuming there are M points on the surface and N scattering directions, the time complexity of the brute force far field radiation calculation is $O(MN)$. In the following subsection, we will describe an octree-based algorithm that achieves a lower time complexity of $O(M + \log(M)N)$ and significantly accelerates the computation.

4.3 Multilevel fast Physical Optics

We accelerate the brute force far-field calculation using an octree-based algorithm. The algorithm is illustrated in Figure 4, following the green arrows. We borrowed the idea from the multilevel fast

multipole algorithm (MLFMA) [Chew et al. 2001], which is originally used to accelerate the first step in the full wave simulation that computes the surface currents. The algorithm works by first constructing an octree that includes all the sampled points on the surface of the scatterer. Then, it computes the surface currents associated with the surface points in the same way as it does in the non-tree-based PO algorithm. Next, from the bottom (leaf) level of the tree to the top level, we accumulate the far-field contribution along the way, computing all the scattering directions at once. This accelerated algorithm is general and can be applied to accelerate the far-field calculation in exact wave simulations as well.

This method has three key components. The first is the translation of the far-field scattering kernel. This kernel is an exponential function that can be expressed as a product of exponential functions.

$$e^{jk_0\mathbf{r}' \cdot \hat{\mathbf{r}}} = e^{jk_0(\mathbf{r}' - \mathbf{c}_L) \cdot \hat{\mathbf{r}}} \prod_{i=1}^L e^{jk_0(\mathbf{c}_i - \mathbf{c}_{i-1}) \cdot \hat{\mathbf{r}}}, \quad (13)$$

where \mathbf{r}' is a surface point, $\mathbf{c}_0, \dots, \mathbf{c}_L$ are the centers of the tree nodes containing that point, from the bottom to the top level; $e^{jk_0(\mathbf{c}_i - \mathbf{c}_{i-1}) \cdot \hat{\mathbf{r}}}$ are translation kernels that transform the far-field contribution from one reference point to another. The second key idea is to make use of the band-limited property of far-field scattering [Bucci and Franceschetti 1987]. In theory, the ideal sampling rate in scattering angles θ_r and ϕ_r scales linearly with the length of the bounding box that contains all the surface points. Thus, we need far fewer sampled directions at the bottom level of the tree to accurately represent the scattering contribution of the surface points contained in each leaf node. To complete the method, we use Fast Fourier Transform (FFT) to perform interpolation and upsampling [Sarvas 2003] as we go up the tree.

To put these three components together, the acceleration algorithm works as follows: We first define the set of directions (θ_r and ϕ_r angles) for each level of the tree. Starting from the leaf level with

the coarsest set of directions, for each direction and each surface point \mathbf{r}' , we compute $e^{jk_0(\mathbf{r}'-\mathbf{c})\cdot\hat{\mathbf{r}}}$, where \mathbf{c} is the center of the grid containing the surface points. Then, we multiply it with the surface currents and compute the sum of the far-field contribution for each leaf node. From the second lowest level to the top level, we have the same set of operations: For each node on the current level, we first upsample its children's far-field contribution from the previous set of directions to the current set. Specifically, we perform a forward FFT, zero-padding the transformed array, and then perform an inverse FFT on each component of the electric and magnetic fields, in both θ and ϕ dimensions. Then, we translate the contribution to the current grid center by multiplying it with the translation kernel $e^{jk_0(\mathbf{c}_{\text{child}}-\mathbf{c})\cdot\hat{\mathbf{r}}}$, where $\mathbf{c}_{\text{child}} - \mathbf{c}$ is the vector from child to self, and $\hat{\mathbf{r}}$ is each direction in the current directions set. This is followed by aggregating all children's contributions. This process continues to the top level of the tree, where we translate and accumulate the contribution to the origin and compute the desired total scattering distribution.

Performance. The octree-based algorithm significantly improves the performance of far-field calculation and achieves a two-order-of-magnitude speedup compared to the brute force implementation on the CPU. We include the simulation time of three different fiber segments in Table 1. The first two fibers are circular, and the last one is elliptical, with different major radius r_1 , minor radius r_2 , and length l . We specify the corresponding number of points represented on the fiber surface and the number of directions for computing the scattering function. We report the simulation's running time on CPU and GPU, with brute force implementation and tree-based implementation at different levels. We ran the GPU simulations on an Nvidia RTX 3090 GPU, and we ran the CPU simulations on an AMD Ryzen Threadripper 3970X 32-core processor with multithreading using OpenMP. The running time corresponds to a simulation for one wavelength and a single incident direction. We implement the GPU-based PO using CUDA kernels and the Thrust library.

Table 1. Simulation time.

	fiber #1	fiber #2	fiber #3
(r_1, r_2, l) in μm	(20, 20, 20)	(40, 40, 50)	(40, 35, 400)
#points \times #dirs	1.4M \times 0.4M	4.7M \times 1.3M	35M \times 10M
CPU	811s	9352s	259200s*
GPU (brute force)	83s	785s	40905s
GPU (2 level tree)	126s	428s	12482s
GPU (3 level tree)	52s	257s	4320s
GPU (4 level tree)	63s	284s	2421s
GPU (5 level tree)	102s	374s	2202s
GPU (6 level tree)	147s	431s	2062s

*CPU run time is estimated using a subset of directions, assuming that the run time for the brute force approach is proportional to the number of directions.

Fiber microgeometry and scattering patterns. We run our wave simulation on 3D microgeometry of hair and fur fibers (Figure 2). We generate these surfaces using a model that displaces the surface of an elliptical cylinder by a Gaussian random height field, optionally adding a cuticle tilt to the surface. The model parameters include the major and minor radii of the elliptical cross-section, the cuticle

length, the cuticle angle, and the parameters of the 2D Gaussian process. There is an additional 1D Gaussian random structure that describes the irregular edge of the cuticle. Our 3D wave simulation results resemble the predictions of ray-based models in many respects, but they differ in two important ways. First, as observed by Xia et al. (2020), wave optics predicts a strong forward-scattering diffraction component. Second, although ray models can predict the overall envelope of the specular highlight, our simulations show intricate, wavelength-dependent granular patterns, and they generate colorful structures when converted from spectral to RGB, as shown in Figure 6. We observe that:

- Fibers that share the same set of geometric parameters produce statistically similar granular patterns with a similar shape and size. Fibers with different geometric parameters, on the other hand, can produce granular patterns with very different statistics.
- As the incident angle changes, the speckle pattern shifts following the half vector direction.
- As the wavelength increases, the size of the speckles also increases. A similar phenomenon has also been studied in planar surface scattering [Goodman 2007]: If w represents the average width of an individual speckle, then for small changes in wavelength from λ_1 to $\lambda_1 + \Delta\lambda$, the change in the average width of an individual speckle lobe in one dimension is from w to $w(1 + \Delta\lambda/\lambda)$.

These behaviors can be better observed in our supplementary video. In the next section, we will develop a practical model that aims to replicate the statistics observed in the simulation.

5 A PRACTICAL FIBER SCATTERING MODEL

So far, we have described our 3D wave simulator, which we can use to compute scattering distributions for given fiber micro-geometries. Similar to the previous wave optics fiber scattering model work [Xia et al. 2020], we can tabulate the scattering distribution and render fibers. However, we run into the problem of handling geometry variations in the scene. For each fiber, and also along a single fiber, we should have variations in micro-geometry, but the memory-intensive tabulation cannot handle such variations. This problem is more severe in 3D, as the table size for the full scattering function is much larger than that in the previous work, which only tabulates the azimuthal scattering functions. To address this problem, in this section, we propose a novel practical fiber scattering model based on a wavelet noise representation.

Rough fibers contain microscopic facets, and when light scatters from them, randomly phased elementary contributions are superposed to form the scattered field. These contributions interfere with each other, and the resultant scattering intensity distribution varies spatially and produces granular patterns. Our goal is to construct a compact representation to capture the important statistics of these phenomena that are relevant for fiber rendering so that we can reproduce colorful glints with high fidelity but without unnecessary cost.

We assume that these granular patterns can be characterized by second-order one- and two-point statistics, namely means, variances, and autocorrelation functions (ACFs). Furthermore, we leverage




Fig. 6. We compare the speckle patterns computed from our wave simulator with the synthesized speckle patterns based on wavelet noise. We have three fiber examples, and we plot their scattering function over the outgoing directions. The first example has higher roughness, and the last example has a cuticle angle of -3 degrees. We are able to produce the memory effect when changing the incident angle in both the simulation and the synthesized results. This effect can be observed in the red rectangles in each row. The speckle pattern generated from the noise function captures the speckle size and color distributions well. The average components used here for the synthesized results are the ensemble averages of the simulation results for each fiber distribution.

speckle theory to inform us about the one- and two-point statistics. In the following, we will introduce speckle statistics and explain how we construct our noise-based representation.

5.1 Speckle statistics

In fiber scattering, the correlation length of the rough surface is usually small compared to the illuminated area and the coherence area. Therefore, we can assume fully developed speckle. For fully developed speckle, the real part (\mathcal{R}) and the imaginary part (\mathcal{I}) of the field at a single point in space follow a complex Gaussian distribution [Goodman 2007]. The joint probability density function for \mathcal{R} and \mathcal{I} of the field is:

$$p_{\mathcal{R},\mathcal{I}}(\mathcal{R},\mathcal{I}) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\mathcal{R}^2 + \mathcal{I}^2}{2\sigma^2}\right). \quad (14)$$

and the intensity (\mathbf{I}) computed from the field follows an exponential distribution

$$\mathbf{I} = \mathcal{R}^2 + \mathcal{I}^2, \quad p_{\mathbf{I}}(\mathbf{I}) = \frac{1}{2\sigma^2} \exp\left(-\frac{\mathbf{I}}{2\sigma^2}\right). \quad (15)$$

In our representation, we would like the intensity to follow this distribution so that we preserve the statistical property for a single direction. We would also like to approximate the two-point statistics of the speckle intensity. The autocorrelation function (ACF) of the scattering intensity $\mathbf{I}_{\mathbf{p}_1}$ and $\mathbf{I}_{\mathbf{p}_2}$ at two points \mathbf{p}_1 and \mathbf{p}_2 is defined as

follows:

$$C(\mathbf{I}_{\mathbf{p}_1}, \mathbf{I}_{\mathbf{p}_2}) = \frac{\overline{(\mathbf{I}_{\mathbf{p}_1} - \overline{\mathbf{I}_{\mathbf{p}_1}})(\mathbf{I}_{\mathbf{p}_2} - \overline{\mathbf{I}_{\mathbf{p}_2}})}}{\sigma(\mathbf{I}_{\mathbf{p}_1})\sigma(\mathbf{I}_{\mathbf{p}_2})}, \quad (16)$$

where the overline symbol denotes the ensemble average. The statistics of the scattering intensity from rough fibers is different from rough planar surfaces, and we propose a novel representation based on wavelet noise to capture the single-point and two-point statistics. By doing so, we are able to produce faithful color variations.

5.2 Wavelet noise representation of the speckles

Wavelet noise is a widely used noise function and it provides a way for us to sample instances of 3D random fields with narrow-band ACFs. We utilize it as a compact representation for $f(\theta_h, \phi_h, \lambda)$.

Similar to Perlin and other procedural noise functions in computer graphics, a wavelet noise function can be written as a weighted sum of different frequency bands. In our problem, we write the final noise component $f(\theta_h, \phi_h, \lambda)$ as a weighted sum of the intensity computed from each single band:

$$f(\mathbf{x}) = \sum_{b=0}^{n-1} w_b(\mathbf{x}) \mathbf{I}(2^b g_\lambda(\mathbf{x})), \quad (17)$$

where $g_\lambda(x) = (s_\theta(\lambda)\theta_h, s_\phi\phi_h, s_\lambda\lambda)$, s_θ is a wavelength-dependent linear scaling; s_ϕ and s_λ are constants per fiber distribution. We will use \mathbf{I}_b as a shorthand for $\mathbf{I}(2^b g_\lambda(\mathbf{x}))$. In our problem, we restrict $w_b(\mathbf{x})$ to be non-negative to avoid negative intensity. We would like to compute the optimal weights $w_b(\mathbf{x})$ for frequency band $b = 0$

to $b = n - 1$, so that the autocorrelation function of the intensity computed from the noise function approximates that of the fiber speckles.

The fitting process works as follows: for a given distribution of rough fiber surfaces, we instantiate different examples of rough fibers from the given distribution. We then run our wave simulator and compute the scattering distribution on half-vector direction patches around a set of half-vector directions. For each direction in the set, we compute the target autocorrelation function $C_t(\mathbf{x}_1, \mathbf{x}_2)$ using the corresponding patches. Then we generate wavelet noise instances for every single band within the same range of directions. We aim to compute corresponding weights for different bands of the noise so that the ACF of the final noise $C_f(\mathbf{x}_1, \mathbf{x}_2)$ is approximately equal to the target.

Following the Wiener-Khinchin theorem, the autocorrelation function of the final noise can be computed using the Fourier Transform:

$$C_f(\mathbf{x}_1, \mathbf{x}_2) = \mathfrak{F} \left[\left\| \mathfrak{F} \left[\sum_{b=0}^{n-1} w_b \mathbf{I}_b \right] \right\|^2 \right]. \quad (18)$$

Instead of directly fitting the weights for each noise band, we transform the problem into fitting the weights for the ACF of each band, making use of the following lemma.

LEMMA 5.1. *The autocorrelation function (ACF) of the weighted sum of bands is approximately equal to the weighted sum of the ACFs of the bands, with the latter weights being the square of the former. That is, when $w_b = \sqrt{v_b}$, we have: $C_f(\mathbf{x}_1, \mathbf{x}_2) \approx \sum_{b=0}^{n-1} v_b C_b(\mathbf{x}_1, \mathbf{x}_2)$.*

PROOF.

$$\begin{aligned} C_f(\mathbf{x}_1, \mathbf{x}_2) &= \mathfrak{F} \left[\left\| \mathfrak{F} \left[\sum_{b=0}^{n-1} \sqrt{v_b} \mathbf{I}_b \right] \right\|^2 \right] \\ &= \sum_{i=0}^{n-1} v_i \mathfrak{F} [|\mathfrak{F} [\mathbf{I}_i]|^2] + \sum_{b \neq d} 2\sqrt{v_b v_d} \mathfrak{F} [|\mathfrak{F} [\mathbf{I}_b] \mathfrak{F} [\mathbf{I}_d]|] \\ &\approx \sum_{b=0}^{n-1} v_b C_b(\mathbf{x}_1, \mathbf{x}_2). \end{aligned}$$

The term $\sum_{b \neq d} 2\sqrt{v_b v_d} \mathfrak{F} [|\mathfrak{F} [\mathbf{I}_b] \mathfrak{F} [\mathbf{I}_d]|]$ has a small contribution that can be ignored (validated in the supplemental material), as the different bands of the wavelet noise are approximately separated in frequency space. \square

Our goal now is to find a set of non-negative coefficients v_b so that the weighted sum of the autocorrelation function $C_b(\mathbf{x}_1, \mathbf{x}_2)$ of each band approximately equals the target autocorrelation function:

$$C_t(\mathbf{x}_1, \mathbf{x}_2) \approx \sum_{b=0}^{n-1} v_b C_b(\mathbf{x}_1, \mathbf{x}_2). \quad (19)$$

We form a least squares problem by first evaluating the target autocorrelation function for a set of points, denoting the values using a vector \mathbf{y} . Then we evaluate the same points for each single band and form a matrix \mathbf{A} , where A_{bi} is the value for the b th band at the i th point. We then solve for the weights \mathbf{v} by minimizing $\|\mathbf{y} - \mathbf{A}\mathbf{v}\|_2$, subject to $\mathbf{v} \geq 0$. Scaling the weights \mathbf{v} together does not change the ACF of the noise. We decide the scaling factor by

ensuring $\mathbb{E}[f(\mathbf{x})] = 1$ on the fitted patch, which ensures that we preserve energy on average:

$$\begin{aligned} &\mathbb{E} [S_{\text{avg}}(\theta_i, \theta_r, \phi_r, \phi_r, \lambda) f(\theta_h, \phi_h, \lambda)] \\ &= S_{\text{avg}}(\theta_i, \theta_r, \phi_r, \phi_r, \lambda) \mathbb{E} [f(\mathbf{x})] \\ &\approx S_{\text{avg}}(\theta_i, \theta_r, \phi_r, \phi_r, \lambda). \end{aligned} \quad (20)$$

Then we let $w_b = \sqrt{v_b}$ and the final noise component is

$$f(\mathbf{x}) = \sum_{b=0}^{n-1} \sqrt{v_b} \mathbf{I}_b. \quad (21)$$

Implementation details and limitations. To represent the speckle pattern using wavelet noise, we generate two noise tiles $M_r(\mathbf{x})$ and $M_i(\mathbf{x})$, analogous to the real and imaginary parts of the field. Each of $M_r(\mathbf{x})$ and $M_i(\mathbf{x})$ is a 3D noise function, with the three dimensions being the longitudinal angle of the half-vector direction, the azimuthal angle of the half-vector direction, and the wavelength dimension, i.e., $\mathbf{x} = (\theta_h, \phi_h, \lambda)$. Each tile has a size of $128 \times 128 \times 50$, corresponding to the longitudinal, azimuthal, and wavelength dimensions. For each rough surface distribution, we use four frequency bands. To handle the geometry variations of fibers, we use the texture coordinates along the fiber and each fiber's unique ID to offset the noise tile in the longitudinal and azimuthal dimensions.

By construction, the value at a given point on $M_r(\mathbf{x})$ and $M_i(\mathbf{x})$ follows a Gaussian distribution, and the intensity computed from them is $\mathbf{I}(\mathbf{x}) = M_r(\mathbf{x})^2 + M_i(\mathbf{x})^2$, which naturally follows an exponential distribution. This guarantees that we have the correct single-point statistics. We observe that the speckle size changes slowly over angles. In practice, we find that fitting around 3 to 5 half-vector directions already achieves good results. The fitting is done using the `nnls` function in `scipy.optimize`. The fitting process, including generating individual bands and computing weights for each band, takes about 5 minutes after we acquire the simulation results.

The fitting has its limitations: It becomes less accurate at grazing incidence angles. It should be noted that the simulation itself is also less accurate towards grazing angles, as the finite segment we simulate will introduce artifacts at the end of the geometry. Additionally, the half-vector direction degenerates in the exact forward direction. Currently, we do not explicitly fit the speckle pattern around the forward direction. An alternative approach would be to use a different scheme around the forward direction, such as fitting with respect to the outgoing direction instead of the half-vector direction.

6 VALIDATION

We first validate our 3D wave simulation and demonstrate that our approximated 3D wave solver has high accuracy in computing surface reflection and diffraction. Then, we conduct a measurement of human hair fiber scattering and show that similar speckle patterns appear in both the captured and simulated results. Furthermore, we validate our noise representation against the simulation results.

6.1 Wave simulation validation

In this work, we develop a 3D wave simulator based on a physical optics (PO) approximation and accelerate the far-field calculation

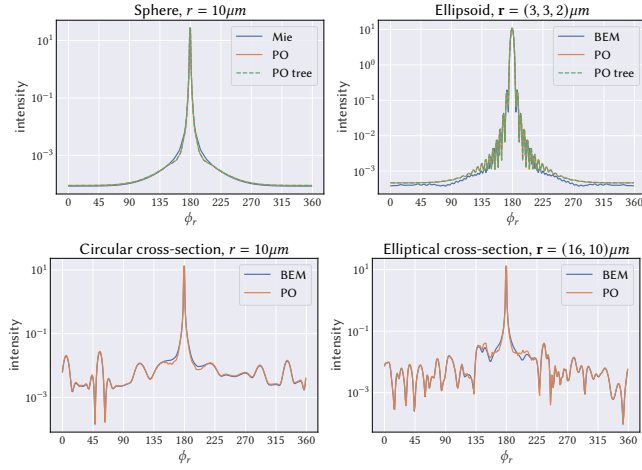


Fig. 7. The top left image compares our 3D wave simulator to Mie theory on sphere scattering. The top right image compares our simulator with exact wave scattering computed using the Boundary Element Method (BEM) in 3D. Here, we compute scattering from an ellipsoid using our simulator and Bemp [Betcke and Scroggs 2021]. The second row compares our simulator with BEM in 2D [Xia et al. 2020], where we compute scattering from a rough circular cross-section and a rough elliptical cross-section.

using a tree-based algorithm. We validate PO and our implementation of the tree-based algorithm by comparing them to the exact wave optics solutions. When absorption is high, surface reflection and diffraction dominate the scattering, so we use highly absorptive materials in all the validation tests. In all the tests, we assume that light travels in the $+x$ direction, and we compute the scattered intensity in the x - y plane for 3600 ϕ_r angles, which we then average down to 360 angles. We compute both polarizations and average them to obtain the unpolarized results.

Comparison with Mie Scattering. We first compare with Mie scattering in the first plot in Figure 7. Mie scattering provides an exact solution to Maxwell’s equations for scattering from homogeneous spheres [van de Hulst 1957]. For a sphere with a radius of $10\mu\text{m}$ and a refractive index of $1.55 - 0.1j$, we show that our brute force implementation and tree-based implementation match exactly as expected, and they are very accurate in most angles compared to Mie, including the exact forward direction ($\phi_r = 180^\circ$) and most backward scattering directions. However, PO lacks energy in some directions. This might be due to the hard cut-off at grazing angles in the PO approximation, whereas in reality, the field smoothly dies off and can creep around grazing angles [Andronov and Bouche 1995]. We study the effect of object size on the PO approximation in the supplemental material. The PO approximation is more accurate when the radius is large compared to the wavelength because the local tangent plane assumption is more accurate when the curvature is smaller.

Comparison with BEM simulation in 3D. We also compare PO with the 3D BEM simulator Bemp [Betcke and Scroggs 2021] (Figure 7 top right image). We simulate wave scattering from a small ellipsoid. The three axes of the ellipsoid align with the x , y , and z axes, with the radii of the three axes being $3\mu\text{m}$, $3\mu\text{m}$, and $2\mu\text{m}$. The refractive

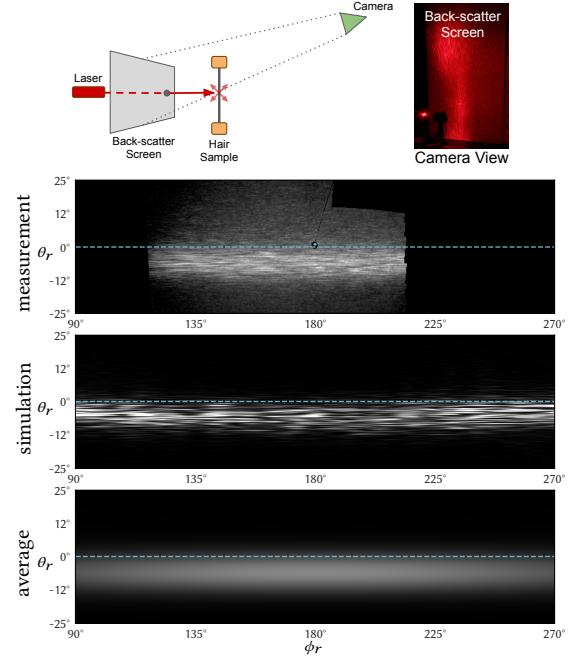


Fig. 8. Optical measurements reveal speckle patterns resulting from rough fiber scattering. In this example, we illuminate a human hair fiber with a laser beam and observe elongated speckle patterns around the backscattering direction. Our 3D simulator produces similar patterns for the estimated parameters obtained from the experiment. Previous models that consider scattering from a smooth fiber or scattering statically averaged over rough fibers only produce a smooth highlight. In all figures, we observe a shift in the θ_r direction due to the cuticle angle. The cyan dot in the measurement image shows the beam direction.

index of the ellipsoid is $1.55 - 0.3j$. The PO approximation captures the overall shape of the scattering well. We expect the accuracy to increase as the size of the object increases. Note that the simulation using Bemp took 3 hours on our machine, while PO took 2 seconds for 3600 ϕ_r angles.

Comparison with BEM simulation in 2D. We compare PO with the exact wave simulation using a 2D Boundary Element Method (BEM) solver that was published by [Xia et al. 2020] in the second row of Figure 7. We construct rough circular and elliptical cross-sections in the x - y plane by wrapping 1D Gaussian height fields around circular and elliptical cross-sections. The roughness of the 1D Gaussian height fields is 0.02 and 0.05 respectively. We use a refractive index of $1.55 - 0.1j$. This comparison shows that our simulator achieves high accuracy in the backward scattering direction and exact forward scattering direction ($\phi_r = 180^\circ$) when handling geometry with roughness.

6.2 Measurement

For comparison, we also measured scattering from a human hair using the setup shown in Figure 8. A laser beam was aligned to illuminate a small spot on a hair sample. Light scattered from the hair falls on a diffuse screen located around the backward scattering directions. A camera images the screen using high dynamic range captures, and the results are converted to relative measurements of



Fig. 9. A head of blonde hair is rendered using our practical model, demonstrating the ability to combine our model with existing higher-order scattering modes to render colored hair. Subtle colored glints can be observed in the zoom-in image on the top right.

the scattered light as a function of direction. Only directions covered by the screen and unobstructed from the camera are captured in this setup; the missing directions are shown as black in this data. The laser is a HeNe with a wavelength of 633nm and a beam spot size of roughly 0.7mm (along the hair) by 3mm (perpendicular to the hair) and passes through a small hole in the screen.

In Figure 8, we compare the measurement data to our simulation and the backscatter from previous models. The measured fiber’s diameter was estimated to be roughly 80 μ m. The cuticle angle was estimated, by the R component shift in the longitudinal direction, to be about -3 degrees. All the results show a shift in the longitudinal direction that is caused by the cuticle angle. However, in the measurement and our simulation, we also observe speckle patterns, similarly elongated in the azimuthal direction, that were missing from previous hair scattering models.

6.3 Noise representation validation

In addition to validating our wave simulator and comparing simulation results with measurements, we further validate our noise representation by comparing our synthesized speckle patterns to the simulated ones. In Figure 6, we show patterns of three fibers corresponding to the first, middle, and bottom two rows. The first one is a circular fiber with a diameter of 80 μ m, surface roughness of 0.1 (extended in Figure 2 a) for visualization). The second one is an elliptical fiber with a major diameter of 80 μ m and a minor diameter of 70 μ m and roughness of 0.02 (Figure 2 b); the last case is an elliptical fiber with a major diameter of 80 μ m and a minor diameter of 70 μ m, and with a cuticle angle of -3 degrees and roughness of 0.02 (Figure 2 c). To synthesize the speckle patterns, for each case, we run the simulation for 50 different instances and fit wavelet noise with 4 frequency bands to the autocorrelation function of the speckle pattern from the simulation. To synthesize the pattern, we multiply the averaged scattering intensity computed from the simulation with the speckle component represented by wavelet noise. We show that our noise representation is able to capture the speckle statistics well, showing similar speckle size and color variations. Moreover, we are able to reproduce the well-known memory effect. As the

incident ϕ angle changes (in the first and second columns), we see the speckle pattern translate in both the simulated and synthesized results. As the wavelength increases (the second and third columns), the speckle size increases in both the simulated and synthesized results.

In both the simulated and synthesized results, we observe that there are two trends of motion as we change the incident ϕ_i angle. The first trend is that the forward scattering and the average component shift in the same direction as the incident direction changes. The second trend is that the noise pattern shifts in the direction opposite to the incident direction. We achieve these two trends naturally by design. These motions can be observed in our supplemental video.

7 RENDERING

The rendered images are generated by PBRT-v3 [Pharr et al. 2016] using the spectral rendering mode. In the following, we will explain how we render each model.

Ray model plus forward diffraction. The ray-based model, denoted as S_{ray} , is the default model in PBRT. The diffraction model, denoted as S_{diffract} , includes an additional tabulated component based on single-slit diffraction. The diffraction from the fiber can be approximated using single-slit diffraction, with the slit width equal to the diameter of the cylinder. The diffraction depends on the ratio of the diameter of the cylinder to the wavelength, denoted as a , the incident longitudinal angle θ_i , and the difference between the azimuthal angle of the outgoing direction and the incoming light propagating direction, $\phi_d = \phi_r - (\phi_i + \pi)$, where $+\pi$ converts the direction in the graphics convention to the light propagating direction.

$$f_{\text{diffract}}(\theta_i, \phi_d, a) = a \cos \theta_i \text{sinc}^2(a \cos \theta_i \sin \phi_d). \quad (22)$$

$f_{\text{diffract}}(\theta_i, \phi_d, a)$ integrates to 1 for $\phi_d \in [-\pi/2, \pi/2]$, which satisfies the energy conservation criterion. In a previous work [Benamira and Pattanaik 2021], single-slit diffraction was also used to approximate the diffraction from a hair fiber, but the dependence on θ_i was missing. We tabulate the diffraction component using a table of size $50 \times 50 \times 200$ and a same-size table for importance sampling the diffraction component. The single-slit diffraction can be used to approximate diffraction from an infinitely long cylinder. To incorporate it into a BCSDf, we multiply it with the longitudinal function used in PBRT [d’Eon et al. 2011].

As explained in [Xia et al. 2020], the extinction cross-section of a fiber is often larger than its geometric cross-section. It is well known that as the particle size increases, the extinction cross-section approaches twice the geometric cross-section and roughly half of the energy goes into diffraction [van de Hulst 1957]. When dealing with human hair and relatively large fur fibers, we can assume that we are in the large particle regime. To include the diffraction component, we enlarge the hair fibers by a factor of two and assign half of the energy to diffraction. The diffraction model for a fiber diameter D is:

$$\begin{aligned} S_{\text{diffract}}(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) \\ = \frac{1}{2} [S_{\text{ray}}(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) + f_{\text{diffract}}(\theta_i, \phi_d, D/\lambda)]. \end{aligned} \quad (23)$$



Fig. 10. We render dog fur using our practical fiber scattering model, demonstrating the ability to handle fibers with textured color.

The importance sampling is performed by randomly sampling the diffraction lobe and the original ray optics model with equal probability.

Our practical model. We multiply a noise component $f(\theta_h, \phi_h, \lambda)$ by a first-order scattering component $S_{0,\text{avg}}$ that represents the surface reflection and diffraction averaged over fiber instances coming from the same distribution. $S_{0,\text{avg}}$ can be any of the existing models that represents a statistical average, combined with the diffraction function f_{diffract} introduced above. We have

$$\begin{aligned} S_{0,\text{prac}}(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) &= S_{0,\text{avg}}(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) f(\theta_h, \phi_h, \lambda) \\ &= \frac{1}{2} [S_{0,\text{ray}}(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) + f_{\text{diffract}}(\theta_i, \phi_i, D/\lambda)] f(\theta_h, \phi_h, \lambda). \end{aligned} \quad (24)$$

The new practical model requires tabulating the diffraction function and generating the noise tiles as the first step of rendering. Together, they take up 10 MB of memory and can support arbitrary geometry variations. In practice, generating the noise tiles as the first step in rendering takes about one second.

Tabulated BCSDf from 3D wave simulation. As mentioned earlier, it is possible to tabulate the 3D simulation results and render them. We will explain how we compute BCSDf from 3D wave simulation and detail the tabulation process.

After running the 3D wave simulation, we obtain $\mathbf{E}_s^{\text{far}}(\hat{\mathbf{r}})$ and $\mathbf{H}_s^{\text{far}}(\hat{\mathbf{r}})$, which define the scattered fields \mathbf{E}_s and \mathbf{H}_s via (12). To define a BCSDf for rendering, we need to relate the scattered intensity to the incident power. In electromagnetics, the time-averaged Poynting vector [Jackson 2021]

$$\langle \mathbf{S} \rangle = \frac{1}{2} \text{Re}(\mathbf{E} \times \mathbf{H}^*) \quad (25)$$

plays a role analogous to vector irradiance in radiometry: given a differential area dA at location \mathbf{r} with a unit normal vector $\hat{\mathbf{n}}$, then the net radiant flux through dA is $\langle \mathbf{S}(\mathbf{r}) \rangle \cdot \hat{\mathbf{n}} dA$. To compute the far-field intensity, consider irradiance on the inside of a sphere of radius R where $kR \gg 1$. Since waves propagate radially [Jackson 2021] in the

far field, $\langle \mathbf{S} \rangle$ and $\hat{\mathbf{n}}$ are both parallel to $\hat{\mathbf{r}}$ and the scattered intensity $I_s(\theta_i, \phi_i, \theta_r, \phi_r, \lambda)$ can be computed from $\mathbf{E}_s^{\text{far}}(\hat{\mathbf{r}})$ and $\mathbf{H}_s^{\text{far}}(\hat{\mathbf{r}})$ as:

$$\begin{aligned} I_s(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) &= (\langle \mathbf{S}(\mathbf{r}) \rangle \cdot \hat{\mathbf{n}}) R^2 = |\langle \mathbf{S}(\mathbf{r}) \rangle| R^2 \\ &= \left| \frac{1}{2} \text{Re}(\mathbf{E}_s^{\text{far}}(\hat{\mathbf{r}}) \times \mathbf{H}_s^{\text{far}}(\hat{\mathbf{r}})^*) \right| \\ &= \frac{1}{2} \sqrt{\frac{\epsilon_0}{\mu_0}} |\mathbf{E}_s^{\text{far}}(\hat{\mathbf{r}})|^2. \end{aligned} \quad (26)$$

By integrating I_s over the sphere, we can compute the scattered power

$$P_s = \int_{\Omega} I_s(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) d\omega, \quad (27)$$

and $P_s \geq 0$. Besides scattering the incoming light, the fiber segment can also absorb light. Absorbed power can be calculated by integrating the normal component of the total field's Poynting vector over the boundary surface Γ , as the net flow at the boundary is the absorption:

$$\begin{aligned} P_a &= \int_{\Gamma} \frac{1}{2} \text{Re}(\mathbf{E}_1 \times \mathbf{H}_1^*) \cdot \hat{\mathbf{n}}_1(A) dA \\ &= \int_{\Gamma} \frac{1}{2} \text{Re}(\mathbf{J}^* \times \mathbf{M}) \cdot \hat{\mathbf{n}}_1(s) ds, \end{aligned} \quad (28)$$

where dA is a differential area on the boundary and $P_a \leq 0$. The second equation is derived by applying Equation (10). We can also compute the energy incident on the fiber segment using

$$P_i = \int_{\Gamma} \frac{1}{2} \text{Re}(\mathbf{E}_i \times \mathbf{H}_i^*) \cdot \hat{\mathbf{n}}_1(A) dA. \quad (29)$$

Similar to [Xia et al. 2020], we define the BCSDf as

$$S(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) = \frac{I_s(\theta_i, \phi_i, \theta_r, \phi_r, \lambda)}{|P_a - P_s|}. \quad (30)$$

In the wave sense, how much light is perturbed by the fiber is related to the extinction cross-section C_{ext} . The extinction efficiency Q_{ext} is the ratio between the extinction cross-section C_{ext} and the geometry cross-section C_g :

$$Q_{\text{ext}} = \frac{C_{\text{ext}}}{C_g} = \frac{|P_a - P_s|}{|P_i|}. \quad (31)$$

To account for the extinction cross-section, we enlarge the fiber in the scene similar to [Xia et al. 2020].

The tabulation is in 5D (1D for the wavelength, 2D for the incoming direction, and 2D for the outgoing direction). We make use of the memory effect of the phenomenon to reduce the table size. Specifically, we observed that as θ_i changes by $\Delta\theta$, the first-order change in the scattering pattern is a shift in θ_o by $\Delta\theta$. Also, as ϕ_i changes by $\Delta\phi$, the scattering pattern shifts in ϕ_r by $\Delta\phi$ for angles that are close to the forward scattering direction, while the rest of the scattering pattern shifts in ϕ_r by $-\Delta\phi$. Thus, we can tabulate the full scattering with less resolution on the incident direction. At rendering time, we query the neighboring tabulated incident directions, apply the corresponding shifts, and linearly interpolate the scattering functions on both θ and ϕ dimensions. We end up with a BSDF tabulation of size $25 \times 32 \times 72 \times 180 \times 360$, which is still very memory-intensive, with each fiber instance costing 15GB




Fig. 11. We compare surface reflection in the back-scattering setting. In the scene, there is one row of fibers rendered using different models. Image a) is rendered by concatenating fiber segments and querying scattering functions from 900 fiber instances. It handles geometry variations in a brute-force way by concatenating fiber segments in the scene, which serves as a reference result. Image b) is our new practical model based on a noise representation. We are able to produce color variations very similar to the reference. Image c) is rendered using one tabulated instance from our 3D simulation. Since no geometry variation is handled in the scene, we see that the color is much more correlated than in the reference. Image d) is rendered using the 2D wave optics fiber scattering model from [Xia et al. 2020]. There is only mild color produced, and since it applies a heuristic longitudinal blur, the color is coherent along the fiber. Image e) is rendered using a ray optics fiber model, which does not produce color.

of memory. Therefore, it is impractical to use tabulation to handle fiber variations.

8 RESULTS

In this section, we demonstrate that our practical wave optics fiber scattering model is capable of producing colorful glints on hair and fur fibers, adding important detail to their appearance. We will also discuss how and why previous models have been unable to successfully produce such effects. Finally, we will report on the performance of our wave simulator and the rendering cost of the new practical model.

8.1 Appearance effects

When looking closely at hair and fur fibers under sunlight, we can easily spot colorful glints, as shown in the dog photo in Figure 1 and the human hair photos in Figures 5 and 9. This color effect, although subtle from a distance, enhances the appearance and can sometimes produce a slight hue shift in the fibers.

Our work considers the full 3D microgeometry of fibers. By explicitly fitting the statistics of the wave simulation results, we are able to realistically reproduce this effect. As shown in Figure 1 and Figure 5, our model produces natural color variation across and along fibers, with similar hues as shown in the photographs. Our model can be easily combined with existing models to render fibers of different colors, as demonstrated by Figures 9 and 10. In Figure 9, we show that we can also reproduce colored glints on light-colored hair. The colored glints are more subtle than those of the dark fibers, as multiple scattering tends to average out the color, and the color contrast is lower.

We compare our results to previous models, which cannot successfully produce the colorful glint effect realistically for various reasons. The ray-based models do not account for wavelength dependence in the surface reflection term, resulting in a colorless appearance, as seen in Figure 1 d) and Figure 11 e). The previous wave optics fiber scattering model by Xia et al. [2020] assumes that the fiber is a perfect extrusion. Although it can handle arbitrary cross-sections, their work lacks the random structural detail along the fiber axis. In practice, their model produces an unnatural correlation of color along the fiber. For example, in Figure 12, the previous model produces a blue color that extends along the fiber, which is very different from what we observe in real life. Additionally, we

compare our results with rendering using a tabulated scattering distribution that comes from our 3D wave simulator. This is the most straightforward way to utilize our 3D simulation for rendering, but it has significant limitations because it cannot handle geometry variations. In reality, each fiber instance is different, and the micro-geometry also varies along the fiber, giving rise to color variations. Rendering using the tabulation of a single instance often results in obtrusive and unnatural color effects on fibers, as shown in Figure 12.

Besides the difference in the produced colors, another difference with the work of [Xia et al. 2020] is that we are able to handle cuticle angles. In Figure 5, we show that our model and the previous model produce similar overall highlights for the no cuticle case. As the cuticle angle changes, we are able to produce the shift of highlight shifts along the fiber caused by the cuticle tilt, which is not considered in the previous model. Moreover, our model is more practical as we can handle geometry variations with a small memory cost (10MB), while the previous model requires about 2.6GB memory to tabulate an azimuthal scattering function of perfectly extruded fiber with a rough cross-section.

It is challenging to obtain a ground truth result for the colorful glint appearance. Besides comparing to photographs, we also compare with a brute force way of handling geometry variations in the scene. We instantiate 900 fiber instances from one rough fiber distribution and tabulate their scattering distributions for one particular viewing direction. In Figure 11, we compare surface reflection in the back-scattering setting. Image a) is rendered using the tabulated instances, handling geometry variations by concatenating fiber segments. We consider it a reference solution. Note that the brute force way is feasible in this artificial scene as we use an orthographic camera and only one viewing direction is required. Image b) is rendered using our practical model, which produces color variations similar to the reference. Image c) is rendered using one tabulated instance from the 3D simulation. We observe that the color is much more correlated than the ground truth result. Image d) is rendered using a tabulated scattering function generated by the 2D simulation provided by [Xia et al. 2020]. The 2D cross-section is one slice of the 3D micro-geometry in c). The color is faint and coherent along the fiber. Image e) is rendered using the ray optics fiber model, producing no color. In our supplemental video, we compare the wave-based models in this setting with a rotating directional

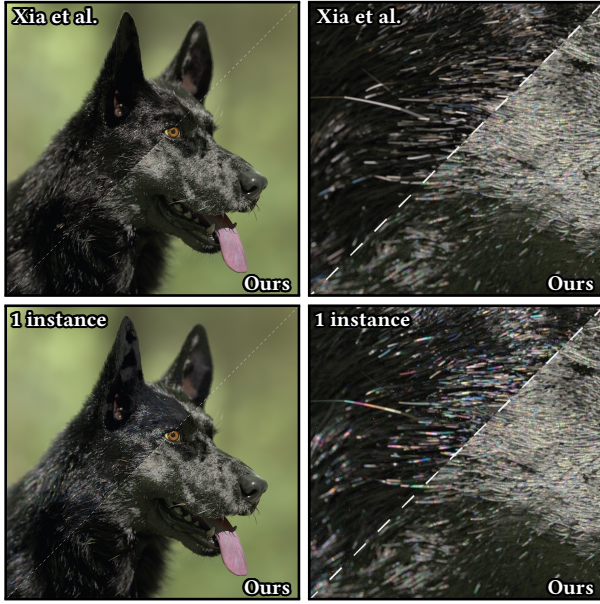


Fig. 12. We compare the color effects on the fur produced by our model to that produced by [Xia et al. 2020]. We also compare with tabulating the scattering function of a single fiber instance after running our 3D wave simulation. The previous model produces faint color with an unrealistic continuation along the fiber. The single instance result produces colorful fibers but with color being unnaturally correlated along the fiber. Our model produces colorful glints with more realistic hues and variations.

light. We observe similar speckle shifts over time in the brute force result and in our model, showing that we are capturing important statistics of the speckles. The tabulated function from the 3D simulation produces vivid color but too much correlation, which is not seen in the brute force result. The tabulated function from the 2D simulation produces much less color and unrealistic coherence. We also generated animation sequences of black hair and blonde hair lit with directional light, changing horizontally and vertically. We observe temporal correlation as the light changes, where colored glints move along the fiber in a continuous manner.

These results demonstrate that our new model adds important appearance effects to the existing models and enhances the realism of rendering hair and fur.

8.2 Rendering time

We report the rendering time in Table 2. Compared to the baseline ray optics model, rendering takes longer when using wave-based methods. We also report the per-sample run time by dividing the number of pixels and the spp count so that it is easier to compare across methods that used different spp counts. The reported time is converted to the run time on our AMD Ryzen Threadripper 3970X 32-core processor. Our method is more efficient than both [Xia et al. 2020] and tabulating the 3D simulation results, as our memory usage is much lower.

Table 2. Rendering time.

Fig.	Method	Resolution	Spp	Time	Per sample
1	Ray	4k×4k	256	0.9h	0.75μs
1	Diffraction	4k×4k	256	1.0h	0.88μs
1	Ours	4k×4k	256	4.8h	4.23μs
12	Xia et al.	4k×4k	64	11.6h	40.78μs
12	Tabulated 3D	4k×4k	32	8.3h	58.35μs
5	Ours	3.2k×3.2k	128	0.5h	1.51μs
5	Xia et al.	3.2k×3.2k	64	1.5h	10.55μs
9	Ours	3.6k×3.6k	16k	466.7h	8.1μs
10	Ours	1440×1080	512	2.1h	9.5μs
11	Ray	1k×1k	64	1.5s	0.02μs
11	Brute force	1k×1k	64	2.0s	0.03μs
11	Ours	1k×1k	64	6.0s	0.09μs
11	Xia et al.	1k×1k	64	10.0s	0.16μs
11	Tabulated 3D	1k×1k	64	19.6s	0.31μs

9 DISCUSSION AND CONCLUSION

This paper takes an important step in wave-based fiber scattering for graphics: we have built the first practical 3D simulation of this phenomenon, finally allowing the full surface geometry to be included in the model. The simulation results predict important features for fiber scattering that cannot be handled using previous models: The highlight is inseparable in the longitudinal and azimuthal directions, which is inaccurately handled in most previous fiber models that assume separability, including the first wave optics fiber scattering model [Xia et al. 2020]. Moreover, the simulation predicts speckle patterns that are missed by all previous wave and ray-based fiber scattering models. These features are observed in real measurements and are essential to the appearance of hair and fur.

To leverage the accurate prediction of surface reflection and diffraction from the simulation, one could tabulate the 5D scattering distribution. However, tabulation is memory-intensive and cannot handle geometry variations in the scene, which greatly limits the applicability of the model. To overcome this obstacle, we develop a practical wave optics fiber model based on a noise representation that captures the statistical properties of the speckle patterns generated by rough fibers. Our practical model makes it possible to incorporate the important colorful glint effect for production hair and fur rendering. In this paper, we only simulate speckle in the first reflection mode, but preliminary studies indicate that similar noise patterns are also present in higher-order modes and likely cause additional colored glints in light hair. Our noise approach could be extended to handle these as well, although it will require developing new techniques to predict or measure the higher mode statistics, which is left as future work. We show that our model can be easily combined with previous fiber models to render hair and fur with various colors. This work brings us closer to a complete wave optics fiber scattering model that can be used for production rendering.

ACKNOWLEDGMENTS

We thank Ningwei Ma for helping with Figure 10, and Peimeng Sui and Anastasia Remizova for the hair photos. This work was supported by the National Science Foundation under grant IIS-1909467.

REFERENCES

- Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. 2017. An appearance model for textile fibers. *Computer Graphics Forum* 36, 4 (2017), 35–45.
- I Andronov and D Bouche. 1995. The creeping and whispering gallery waves on the surface of a transparent body. *Journal of electromagnetic waves and applications* 9, 4 (1995), 503–520.
- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–22.
- Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2020. Rendering near-field speckle statistics in scattering media. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–18.
- Petr Beckmann and Andre Spizzichino. 1987. The scattering of electromagnetic waves from rough surfaces. *Norwood* (1987).
- Alexis Benamira and Sumanta Pattanaik. 2021. A combined scattering and diffraction model for elliptical hair rendering. *Computer Graphics Forum* 40, 4 (2021), 163–175.
- Timo Betcke and Matthew Scroggs. 2021. Bempp-cl: A fast Python based just-in-time compiling boundary element library. *Journal of Open Source Software* 6, 59 (2021), 2879–2879.
- Anders Bondeson, Thomas Rylander, and Pär Ingelström. 2012. *Computational electromagnetics*. Springer.
- Robert Bridson, Jim Hourihane, and Marcus Nordenstam. 2007. Curl-noise for procedural fluid flow. *ACM Transactions on Graphics (ToG)* 26, 3 (2007), 46–es.
- OVIDIOM Bucci and Giorgio Franceschetti. 1987. On the spatial bandwidth of scattered fields. *IEEE transactions on antennas and propagation* 35, 12 (1987), 1445–1455.
- Weng Cho Chew, Eric Michielssen, JM Song, and Jian-Ming Jin. 2001. *Fast and efficient algorithms in computational electromagnetics*. Artech House, Inc., USA.
- Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. 2016. A Practical and Controllable Hair and Fur Model for Production Path Tracing. *Computer Graphics Forum* 35, 2 (2016), 275–283. <https://doi.org/10.1111/cgf.12830> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12830>
- O Clausen, Y Chen, A Fuhrmann, and Ricardo Marroquim. 2023. Investigation and Simulation of Diffraction on Rough Surfaces. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, 245–260.
- Robert L Cook and Tony DeRose. 2005. Wavelet noise. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 803–811.
- Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. 2011. An Energy-conserving Hair Reflectance Model. In *Proceedings of the Twenty-second Eurographics Conference on Rendering* (Prague, Czech Republic) (EGSR '11). Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 1181–1187. <https://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. *SIGGRAPH Talks* 1 (2014), 46–1.
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P Greenberg. 2015. Predicting appearance from measured microgeometry of metal surfaces. *ACM Transactions on Graphics (TOG)* 35, 1 (2015), 1–13.
- Viggo Falster, Adrian Jarabo, and Jeppe Revall Frisvad. 2020. Computing the bidirectional scattering of a microstructure using scalar diffraction theory and path tracing. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 231–242.
- Walton C Gibson. 2021. *The method of moments in electromagnetics*. Chapman and Hall/CRC, USA.
- Joseph W Goodman. 2007. *Speckle phenomena in optics: theory and applications*. Roberts and Company Publishers.
- James E Harvey. 1979. Fourier treatment of near-field scalar diffraction theory. *American Journal of Physics* 47, 11 (1979), 974–980.
- Xiao D He, Kenneth E Torrance, Francois X Sillion, and Donald P Greenberg. 1991. A comprehensive physical model for light reflection. *ACM SIGGRAPH computer graphics* 25, 4 (1991), 175–186.
- Nicolas Holzschuch and Romain Pacanowski. 2017. A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.
- Weizhen Huang, Matthias B Hullin, and Johannes Hanika. 2022. A Microfacet-based Hair Scattering Model. , 79–91 pages.
- P Huddleston, L Medgyesi-Mitschang, and J Putnam. 1986. Combined field integral equation formulation for scattering by dielectrically coated conducting bodies. *IEEE transactions on antennas and propagation* 34, 4 (1986), 510–520.
- John David Jackson. 2021. *Classical electrodynamics*. John Wiley & Sons.
- Jian-Ming Jin. 2015. *The finite element method in electromagnetics*. John Wiley & Sons.
- James T Kajiya. 1985. Anisotropic reflection models. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*. 15–21.
- Kane Yee. 1966. Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE Transactions on Antennas and Propagation* 14, 3 (May 1966), 302–307.
- Pramook Khungurn and Steve Marschner. 2017. Azimuthal Scattering from Elliptical Hair Fibers. *ACM Trans. Graph.* 36, 2, Article 13 (April 2017), 23 pages. <https://doi.org/10.1145/2998578>
- Andrey Krywonos. 2006. Predicting surface scatter using a linear systems formulation of non-paraxial scalar diffraction. (2006).
- Andrey Krywonos, James E Harvey, and Narak Choi. 2011. Linear systems formulation of scattering theory for rough surfaces with arbitrary incident and scattering angles. *JOSA A* 28, 6 (2011), 1121–1138.
- Ares Lagae, Sylvain Lefebvre, Rob Cook, Tony DeRose, George Drettakis, David S Ebert, John P Lewis, Ken Perlin, and Matthias Zwicker. 2010. A survey of procedural noise functions. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, USA, 2579–2600.
- Carmen LaTorre and Bharat Bhushan. 2005. Nanotribological characterization of human hair and skin using atomic force microscopy. *Ultramicroscopy* 105, 1-4 (2005), 155–175.
- Tomas Linder. 2014. *Light scattering in fiber-based materials: a foundation for characterization of structural properties*. Ph. D. Dissertation. Luleå tekniska universitet.
- Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light Scattering from Human Hair Fibers. *ACM Trans. Graph.* 22, 3 (July 2003), 780–791. <https://doi.org/10.1145/882262.882345>
- Rashmi Nanjundaswamy. 2016. Scanning electron microscope image of a human hair (black and white). <https://www.nisenet.org/catalog/scientific-image-sem-human-hair>.
- Mark Olano, John C Hart, Wolfgang Heidrich, Bill Mark, and Ken Perlin. 2002. Real-time shading languages.
- Ken Perlin. 1985. An image synthesizer. *ACM Siggraph Computer Graphics* 19, 3 (1985), 287–296.
- Ken Perlin. 2002. Improving noise. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. 681–682.
- Ken Perlin and Fabrice Neyret. 2001. Flow noise. In *28th International Conference on Computer Graphics and Interactive Techniques (Technical Sketches and Applications)*. Siggraph, 187.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically based rendering: From theory to implementation*. Morgan Kaufmann.
- Andrew J Poggio and Edmund K Miller. 1970. *Integral equation solutions of three-dimensional scattering problems*. MB Assoc., USA.
- Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. 2010. An artist friendly hair shading system. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 56.
- Jukka Sarvas. 2003. Performing interpolation and antepolation entirely by fast Fourier transform in the 3-D multilevel fast multipole algorithm. *SIAM J. Numer. Anal.* 41, 6 (2003), 2180–2196.
- Jiming Song, Cai-Cheng Lu, and Weng Cho Chew. 1997. Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects. *IEEE Transactions on Antennas and Propagation* 45, 10 (1997), 1488–1493.
- Jos Stam. 1999. Diffraction shaders. *Siggraph* 99 (1999), 101–110.
- Shlomi Steinberg and Ling-Qi Yan. 2022. Rendering of Subjective Speckle Formed by Rough Statistical Surfaces. *ACM Transactions on Graphics (TOG)* 41, 1 (2022), 1–23.
- Allen Taflov, Susan C Hagness, and Melinda Picket-May. 2005. Computational electromagnetics: the finite-difference time-domain method. *The Electrical Engineering Handbook* 3 (2005), 629–670.
- H. C. van de Hulst. 1957. Light scattering by small particles.
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B Hullin. 2017. Scratch iridescence: Wave-optical rendering of diffractive surface structure. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 1–14.
- Te-Kao Wu and L Tsai. 1977. Scattering by arbitrarily cross-sectioned layered, lossy dielectric cylinders. *IEEE Transactions on Antennas and Propagation* 25, 4 (1977), 518–524.
- Mengqi Xia, Bruce Walter, Eric Michielssen, David Bindel, and Steve Marschner. 2020. A wave optics based fiber scattering model. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering Specular Microgeometry with Wave Optics. *ACM Trans. Graph.* 37, 4, Article 75 (July 2018), 10 pages. <https://doi.org/10.1145/3197517.3201351>
- Ling-Qi Yan, Henrik Wann Jensen, and Ravi Ramamoorthi. 2017. An Efficient and Practical near and Far Field Fur Reflectance Model. *ACM Trans. Graph.* 36, 4, Article 67 (July 2017), 13 pages. <https://doi.org/10.1145/3072959.3073600>
- Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. Physically-Accurate Fur Reflectance: Modeling, Measurement and Rendering. *ACM Trans. Graph.* 34, 6, Article 185 (Oct. 2015), 13 pages. <https://doi.org/10.1145/2816795.2818080>
- Arno Zinke, Martin Rump, Tomás Lay, Andreas Weber, Anton Andriyenko, and Reinhard Klein. 2009. A Practical Approach for Photometric Acquisition of Hair Color. In *ACM SIGGRAPH Asia 2009 Papers* (Yokohama, Japan) (SIGGRAPH Asia '09). Association for Computing Machinery, New York, NY, USA, Article 165, 9 pages. <https://doi.org/10.1145/1661412.1618511>