

Realistic Cloth Rendering with a Ray-Wave Hybrid Shading Model

YUNCHEN YU, Cornell University & NVIDIA, USA
BRUCE WALTER, Cornell University, USA
STEVE MARSCHNER, Cornell University & NVIDIA, USA
ANDREA WEIDLICH, NVIDIA, Canada



Fig. 1. Our wave-ray hybrid ply model can simulate a wide range of woven fabrics, from translucent to completely opaque (left). By augmenting the average scattering response with per-instance variations, we can simulate imperfections on fabrics, resulting in a more realistic and visually rich appearance (right).

Realistic fabric rendering is still a significant challenge due to their complex structures and varying fiber properties. We present a new fabric shading technique, which models both reflection and transmission using a hybrid of ray and wave optics methods, grounded in simulation data.

We target fabrics woven from yarns, each formed by twisting together one or more plies, which further contain twisted fibers. Our model is based on simulations that predict the scattering of a narrow Gaussian beam by a single ply. Comparing results from full-wave simulations and path tracing, we found that ray optics can accurately simulate the average far field scattering from an ensemble of plies, but not the variation among individual ply instances, and ray tracing overlooks important diffraction effects. Following these observations, our model is built from ray simulations performed for many ply instances, with simulation data fitted by Gaussian mixtures to be used during rendering. Wave simulations are used to calibrate noise functions that account for instance-to-instance variation, and an aperture diffraction model is used to handle light passing between plies and yarns.

The result is a hybrid model capable of producing realistic appearance and highlight structure in fabrics, while capturing spatial break-ups and irregularities and simulating the subtle color shifts and blurriness that occur in transmission. We validate our results by comparing rendered images with

photographs, demonstrating the effectiveness of our approach in achieving realistic cloth rendering.

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

Additional Key Words and Phrases: wave optics, material appearance, scattering, fabrics, yarn

ACM Reference Format:

Yunchen Yu, Bruce Walter, Steve Marschner, and Andrea Weidlich. 2025. Realistic Cloth Rendering with a Ray-Wave Hybrid Shading Model. *ACM Trans. Graph.* 44, 6, Article 217 (December 2025), 16 pages. <https://doi.org/10.1145/3763359>


1 INTRODUCTION

In recent years, significant progress has been made in fabric rendering, particularly through ply-based models [Zhu et al. 2024, 2023a]. These models have substantially improved realism by relying on physical principles and measurable properties, rather than on empirical methods [Estevez and Kulla 2017; Zeltner et al. 2022] that match appearance purely through observation.

Despite these advancements, faithfully capturing the appearance of fabrics remains challenging. Modeling both reflection and transmission accurately is difficult, and characteristic irregularities of real materials are often inadequately addressed. Additionally, many current models rely on idealized assumptions—for instance, fibers are assumed to be perfectly circular and lack natural variation, and wave optics effects are ignored—which do not reflect the complexity of real-world fabrics.

We propose a ply model derived from first principles, combining the realism of wave optics with the speed of ray optics. Our ply

Authors' addresses: Yunchen Yu, yy735@cornell.edu, Cornell University & NVIDIA, USA; Bruce Walter, bruce.walter@cornell.edu, Cornell University, USA; Steve Marschner, srm@cs.cornell.edu, Cornell University & NVIDIA, USA; Andrea Weidlich, aweidlich@nvidia.com, NVIDIA, Canada.

Please use nonacm option or ACM Engage class to enable CC licenses 
This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2025 Copyright held by the owner/author(s).
ACM 0730-0301/2025/12-ART217
<https://doi.org/10.1145/3763359>

scattering model is based on multi-fiber full-wave simulations and path tracing, where both simulation methods account for the true, generalized shapes of fibers. Comparisons between wave and ray simulations show that ray optics can accurately compute the average far field scattering from a large collection of plies, but not the variation among individual ply instances. Thus, we model average ply scattering by fitting ray simulation results to a small set of Gaussian lobes, achieving a lightweight representation. Moreover, since the spatially varying and sometimes glinty appearance of fabrics arises from instance-to-instance differences in ply structures, we derive noise data from a limited set of wave simulations and use them to augment our average ply scattering model.

Moreover, our model handles multi-bounce interactions between multiple plies and in double-yarn configurations, resulting in a new, more diffuse transmission model that captures complex scattering and absorption effects across ply layers. Additionally, we introduce wave optics based diffraction lobes to more realistically model transmission through gaps and around yarn edges. Together, these components enable a robust and accurate approach to fabric rendering, contributing to the advancement of appearance modeling.

2 RELATED WORK

In this section, we review prior works related to fabric rendering, with a focus on far field reflectance models, ply-based yarn models, and recent applications of wave optics in rendering.

2.1 Far Field Fabric Reflectance Models

Rendering fabric has a long tradition. Early models were mostly empirical, focusing on phenomenological features [Ashikhmin and Premoze 2007], or used microfacets to fit measured data [Wang et al. 2008]. Later works enhanced realism using geometric features such as implicit micro-cylinders [Adabala et al. 2003; Irawan and Marschner 2012; Sadeghi et al. 2013] or microflake models that split surface reflectance into single and multiple scattering components [Jin et al. 2022; Tang et al. 2024; Wang et al. 2022]. [Garces et al. 2023] measured fabrics at two scales and used a neural network to propagate microscale SVBRDF parameters to the mesoscale.

Alternatively, volumetric models simulate fabrics using fiber-like participating media with microflakes [Jakob et al. 2010] or as heterogeneous volumes [Schroder et al. 2011]. While effective in many scenarios and sometimes efficient enough for real-time use, these models lack detailed fiber-level representation.

Data-driven methods, such as bidirectional texture functions (BTFs) [Sattler et al. 2003] or volumetric scans from techniques like micro-CT [Zhao et al. 2011, 2012], mitigate this by directly capturing rich structural details. However, they come with high memory demands, complex acquisition setups, and limited support for editing. An extensive survey on recent methods in fabric rendering is available in [Castillo et al. 2019].

2.2 Ply-Based Yarn Models

In recent years, fiber-based and ply-based models have gained attention for their ability to incorporate measurable physical parameters, leading to more realistic appearance. Notably, works such as [Aliaga




Fig. 2. Photos of various types of fabrics demonstrating subtle wave optics effects, in both reflection and transmission.

et al. 2017; Khungurn et al. 2015; Zhao et al. 2016] model yarns explicitly at the fiber level, achieving high detail at the cost of memory and performance inefficiencies. Consequently, several approaches simulate at the ply level, avoiding explicit modeling of individual fibers in favor of cluster-level scattering approximations or precomputation techniques [Khungurn et al. 2017; Montazeri et al. 2021, 2020; Wu and Yuksel 2017; Zhu et al. 2023b].

Ply-based models can be applied to both curves and surfaces; however, when used on surfaces, additional components are often required. For instance, [Zhu et al. 2023a] introduced a Dirac delta transmission term to simulate transmission through weave gaps, and [Zhu et al. 2024] considered parallax effects. More recently, [Khattar et al. 2025] extended their model to support a second yarn layer resulting from weaving. Our model also belongs to the ply-based category: similar to data-driven methods, it uses simulation data as input but approximates reflectance using analytical functions.

Neural techniques have also been employed to render fabrics [Chen et al. 2024; Soh and Montazeri 2023], while we consider these works orthogonal to ours.

2.3 Wave Optics in Appearance Modeling

Wave optics remains a niche area in graphics, but various phenomena have been explored, including interference from layered structures [Dias 1991; Gondek et al. 1994; Hirayama et al. 2001; Icart and Arquès 1999, 2000], dispersion [Sun et al. 1998], and diffraction from small periodic structures [Egholm and Christensen 2006; Okada et al. 2013; Stam 1999; Sun et al. 2000]. Kirchhoff and Fraunhofer theories have been used to model diffraction around cylinders [Benamira and Pattanaik 2021] and polygonal solids [Steinberg et al. 2024b]. In this work, we use Kirchhoff diffraction computations to model light passing through micro-scale fabric weave gaps.

Textile fibers typically range from 10–50 μm in diameter [Houck 2009], though some synthetic fibers can reach as small as 1 μm [Chang et al. 2020]. Yarn spacings rarely exceed 1 mm and is often much smaller depending on weave tightness. These scales suggest that diffraction may influence fabric appearance; indeed, subtle color effects can be observed in photographs (see Fig. 2). Nevertheless, to our knowledge, no prior work has incorporated wave optics into fabric appearance models, and we address this gap by embedding wave-based effects into our ply and fabric model.

The closest works to ours are [Xia et al. 2020], which used the boundary element method (BEM) to derive a BCSDf for hair, and [Xia et al. 2023], which introduced a noise-based model to simulate imperfections. A similar technique was also adopted by [Yu et al. 2024] for modeling the appearance of iridescent feathers.




Fig. 3. Cross sections of different fibers. While wool fibers are near-elliptical, cotton fibers are more kidney-shaped with a hollow space in the center, and silk fibers are triangular or stone-shaped. Images from [Markova 2019].

3 BACKGROUND

In this section, we describe the structural organization of fabrics and our surface-based representation of woven cloth. We then introduce the role of wave optics in our framework and summarize the core ideas behind our ply-level scattering model.

3.1 Fabric Structure

Rendering realistic fabrics is challenging due to their geometric complexity. Fabrics are composed of yarns that are either knitted or woven, and these yarns consist of twisted plies. Each ply is made up of numerous fine fibers further twisted together. Fibers fall into two broad categories: natural and manufactured. Natural fibers include materials such as wool, cotton, and silk, while manufactured fibers—produced by processing natural or synthetic polymers—include examples like polyester and viscose.

The cross section shape of a fiber depends on its material. While most previous works assume circular cross sections, real fibers exhibit a range of shapes, resulting in distinct light scattering characteristics. Fig. 3 shows examples of various fiber shapes.

3.2 Surface-Based Representation

Our work focuses on woven fabrics, which we represent using surface-based models. Each fabric is modeled as a surface mesh, which is easy to generate from cloth simulators and requires less memory than representing yarns as curves. Surface-based rendering is also generally less expensive than curve-based shading, despite known limitations such as inter-shadowing and parallax effects.

Surface-based models of woven cloth typically consist of two main components: the weave pattern and the geometry at the yarn or ply level. As in many previous works, we represent relevant information using textures. The weave pattern (e.g. plain weave, satin, twill, etc.), is textured using a weaving grammar; moreover, each point on the fabric surface maps to a position on a yarn, and local attributes such as yarn direction and yarn width are stored in texture maps. At each shading point, we reconstruct local yarn geometry by evaluating these textures using UV coordinates.

Each yarn contains one or more plies, and in multi-ply yarns, we model the plies as circular helices twisted together. For fabric rendering, we are interested in ply-level information: for a ray that intersects our fabric surface, we use our textured information to determine whether this ray hits a yarn, and further, which ply in




Fig. 4. Beam-ply interaction events. We distinguish between three different events, a near-center hit where the beam fully hits the ply, an off-center hit where the beam center intersects the ply but a fraction of the beam leaks from the ply edge, and a “miss” case where the beam center does not hit the ply but part of the beam may still scatter from the ply edge.

this yarn and where on the ply it hits, using an implicit ray-helix intersection calculation as in [Zhu et al. 2023b]. If an intersection is found, we shade this point on the fabric with our ply model.

3.3 Wave Optics Extension

Importantly, our use of wave optics sets us apart from all previous methods: an ideal delta ray does not exist in wave optics [Steinberg et al. 2024a], and each ray in our framework actually represents a *beam* with finite spatial extent. These beams are still traced as rays, with each ray corresponding to the center of a beam.

Unlike in pure ray optics, beam-ply interactions are not binary. In addition to full hits and complete misses, there exists an extra case: partial hits, where our incident beam partially overlaps with a ply. More specifically, we consider three types of beam-ply interaction events in our context, as shown in Fig. 4 (a)–(c):

- If a beam hits near the center of a ply, it fully interacts with the ply and scatters from fibers.
- If it hits near the edge of a ply, part of its energy will not interact with fibers and instead undergoes edge diffraction.
- If the beam center does not intersect any ply, the beam diffracts through a weave gap, and small portions of it may scatter from fibers in surrounding plies.

Identifying different beam-ply interactions is straightforward with our surface representation. Given a beam, we can determine the beam center’s intersection point on a ply, and when no intersection is detected, we evaluate distances between the beam center and nearby plies to identify relevant edge-scattering or diffraction events. This is possible because the full fabric structure—including information on neighboring shading points on the fabric surface—is readily accessible through textures. Moreover, our three types of beam-ply interactions relate to different components of our ply shading model, as summarized in the following section.

3.4 Method Overview

Our physically based, wave-augmented ply scattering model addresses different types of light-ply interactions and integrates well with our surface-based fabric representation. The model is fundamentally based on accurate simulations of multi-fiber scattering. As presented in Section 4, our work begins with experimenting on simulation methods for computing the collective scattering of fibers

in a ply. Specifically, wave simulation results serve as our ground truth data, validating the accuracy of more practical path tracing simulations—what we primarily use to build our model.

In Section 5.1 and 5.2, we describe our compact analytical representation of single-ply scattering, obtained by carefully fitting Gaussian lobes to simulated scattering distributions. Alongside this multi-lobe ply scattering model, we also introduce a ply-edge diffraction lobe to account for partial beam-ply overlap. Furthermore, as presented in Section 5.3 and 5.4, we extend our single-ply model to handle multi-ply scattering and account for light interactions with double yarn layers resulting from weaving.

Furthermore, in Section 6.1, we present our diffraction model based on the Kirchhoff diffraction integral [Heurtley 1973], which describes light propagation around plies or through weave gaps. In Section 6.2, we introduce our wave optics based noise functions for simulating spatially varying appearance along plies and across fabric surfaces. As we will see in our results, our wave-augmented ply shading model leads to highly realistic fabric rendering.

4 ACCURATE SIMULATIONS OF PLY SCATTERING

In this section, we present our geometric representation of a ply and discuss a few 2D and 3D simulation methods for computing ply scattering, which form the basis of our single-ply shading model.

Table 1. Important notations in Section 4.

R	Radius of a ply
p_{fiber}	Pitch describing fiber twists in plies
α_{fiber}	Fiber roughness
σ_k	Fiber absorption coefficients
ω_i, ω_o	Incident and outgoing directions
θ_i	Longitudinal incident angle
h	Incident position on the ply
λ	Wavelength of the light
w_0	Waist of the incident Gaussian beam
d_{cut}	Gaussian beam distance threshold
f^{miss}	Distribution of paths that miss all the fibers
f^{refl}	Outgoing distribution of reflection-only paths
f^{high}	Outgoing distribution of all other paths

4.1 Ply Geometric Modeling

We first introduce our geometric modeling of plies in both 2D and 3D simulations. Each ply is modeled as a bundle of moderately twisted, locally aligned fibers, with each fiber represented as a generalized cylinder having an arbitrary, procedurally generated cross section. The fibers consist of dielectric material characterized by a refractive index and an absorption coefficient. As noted in Section 3.1, the shape and size of the cross section depend on the fiber type.

In 2D simulations, each ply is represented with an integrated cross section model, procedurally created by generating and clustering a number of fiber cross sections, as shown in Fig. 5 (a). Measurements (see supplemental document) show the number of fibers per ply in fabric yarns typically ranges from 20 to 50. As 2D models only capture plies' cross sections, the fibers are assumed to be perfectly parallel, and thus 2D simulations target plies with no fiber twist.

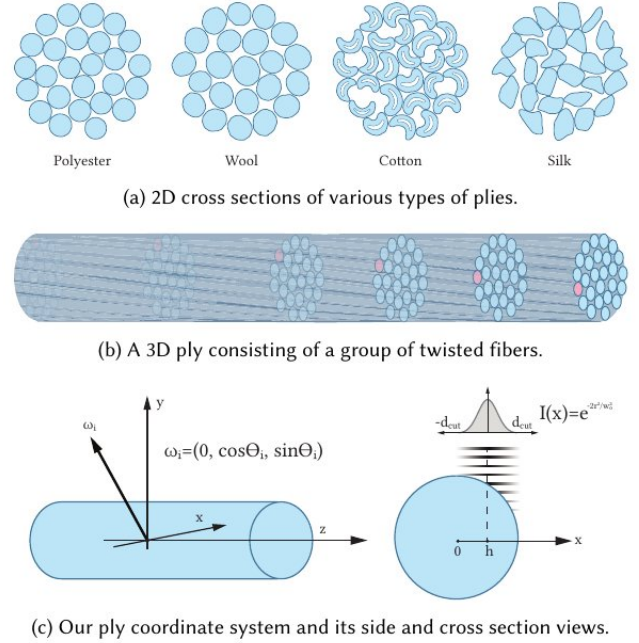


Fig. 5. Ply geometric models and coordinate system used in our simulations.

For 3D simulations on plies with twisted fibers, we extend our 2D cross section representation into 3D by sweeping each fiber's cross section along a helical path. Each fiber follows a uniform helical twist, with its cross section rotating smoothly along the fiber axis. The amount of twist is governed by a parameter p_{fiber} , which defines the pitch of the helix. This procedure generates a twisted 3D geometry for each fiber, which we discretize into a mesh. Assembling the 3D meshes of all fibers produces our complete ply geometry, as illustrated in Fig. 5 (b).

It is worth noting that each 2D or 3D simulation discussed below targets a single ply instance with a specific geometric realization. However, our final ply model is based on the average scattering behavior over an ensemble of such instances, requiring us to generate many ply geometries and perform simulations for each.

4.2 Ply Coordinate System and Simulation Setup

Having introduced our geometric representation of a ply, we now formulate our single ply simulation problem. Each simulation takes a ply geometry and an incident beam as inputs. The beam is modeled with a Gaussian intensity profile, inspired by Gaussian beams in wave optics [Dickson 1970], and is characterized by a beam waist w_0 (we use $w_0 = 10\mu\text{m}$). To specify the beam's direction and position, we first introduce our ply coordinate system.

As shown in Fig. 5 (c), the z -axis in ply space aligns with the ply axis. The incident beam is characterized by an incident angle $\theta_i \in (-\pi/2, \pi/2)$, measured with respect to z . The y -axis is chosen such that the beam always enters from the $+y$ direction in this local frame. Mathematically, the incident direction in ply space is

$$\omega_i = \langle 0, \cos \theta_i, \sin \theta_i \rangle \quad (1)$$

where, following BSDF convention, ω_i points away from the ply. The x -axis is determined from y and z , allowing us to specify the beam's position with the parameter h —the x -coordinate of the beam center, as illustrated in Fig. 5 (c). Given the beam's Gaussian intensity profile, its intensity becomes negligible beyond a distance d_{cut} from the beam center (we use $d_{\text{cut}} = 2w_0$, assuming an intensity profile $I(r) = e^{-2r^2/w_0^2}$). Thus, letting R denote the ply radius, the three beam–ply interaction cases shown in Fig. 4 (a)–(c) correspond to different ranges for the h parameter:

$$0 \leq |h| < R - d_{\text{cut}}; \quad R - d_{\text{cut}} \leq |h| < R; \quad R \leq |h| < R + d_{\text{cut}} \quad (2)$$

We characterize each incident beam using the parameters (θ_i, h) . Given such an incident condition, our simulation computes the scattering distribution from a ply as a function of the outgoing direction, resulting in a model of the form $f(\theta_o, h; \omega_o)$, or equivalently $f(\theta_o, h; \theta_o, \phi_o)$, where $\omega_o \sim (\theta_o, \phi_o)$ is an outgoing direction in the ply coordinate system:

$$\omega_o = \langle \cos \theta_o \cos \phi_o, \cos \theta_o \sin \phi_o, \sin \theta_o \rangle \quad (3)$$

Note that while our model distinguishes between incident positions, it stores scattering as a function of outgoing direction in the far field limit. Also, the model does not depend on an incident azimuthal angle, as it is based on ensemble-averaged simulation results from many ply instances; we assume this averaging makes the far field scattering invariant to rotation around the ply axis.

We next present our 2D and 3D simulations for computing scattering distributions from plies, including our wave-based ground truth and more practical ray-based methods.

4.3 Full-Wave Simulations

The constituent fibers in plies and yarns are very small, with diameters typically between $10\mu\text{m}$ and $20\mu\text{m}$. Thus, we need to consider the wave nature of light when simulating scattering from plies. Since it is important to compute ground truth distributions from multi-fiber scattering, accounting for wave effects as well as scattering of all orders, we built a 2D full-wave simulator based on the Boundary Element Method (BEM) for multiple objects [Gibson 2021].

BEM is an accurate method for computing object scattering, with no approximation other than numerical discretization. The method has been thoroughly discussed in recent works [Xia et al. 2020; Yu et al. 2023], where it was used for computing scattering from single objects. Our multi-fiber BEM solver generalizes previous simulators to multiple scattering objects, while our simulation costs increase quickly with the number and sizes of fibers—despite GPU and numerical approximation acceleration [Bleszynski et al. 1996]. For scalability, we keep our wave simulations in 2D, where all computations are performed on cross section models. These simulations compute scattering from plies with parallel, non-twisted fibers.

Each wave simulation starts with a 2D cross section model of a ply as well as a Gaussian beam incident field with wavelength λ , described by the incident parameters (θ_i, h) . Note that these simulations are capable of computing ply scattering under out-of-plane incidence, i.e. $\theta_i \neq 0$. The result is a 1D scattering distribution with respect to the outgoing azimuthal angle ϕ_o , as the translational symmetry in the ply constrains its scattering distribution in a specular cone where $\theta_o = -\theta_i$, as in [Xia et al. 2020] and [Yu et al. 2024].




Fig. 6. For path tracing, we simulate a Gaussian beam incidence on a ply using a ray bundle, where the spatial and angular distributions of the rays match those in the Gaussian beam.

Since these 2D simulations cannot generate full scattering distributions of plies containing twisted fibers and are also expensive, we do not use them directly for computing ply scattering models. Instead, we use this simulator as a reference and diagnostic tool that we compare with more practical methods.

4.4 Path Tracing Simulations

A natural alternative to wave simulations is standard path tracing, which is significantly more tractable and can be easily implemented in 2D (tracing fiber cross sections) or 3D (tracing full meshes).

When using path tracing to compute scattering distributions as a function of the incident direction, we are computing samples of a function, and to avoid aliasing, these samples should be filtered over some range of incident directions. We provide this filtering by illuminating the ply with a bundle of rays with Gaussian distributions in their incident positions and incident angles. This also allows us to simulate a Gaussian beam incidence on our ply, by matching the spatial and angular distributions of the rays with those in the Gaussian beam. Our path tracing uses the same incident beam for all wavelengths, and since a fixed-waist Gaussian beam's angular divergence is wavelength dependent, the angular distribution underlying our ray bundle exactly matches the Gaussian beam for one wavelength and is slightly mismatched for others.

This way, our path tracing simulation starts with incident rays distributed over the spatial and angular extent of an incident Gaussian beam, as shown in Fig. 6. The spatial distribution of rays is determined by the beam waist w_0 and incident position h , and the angular distribution matches the angular divergence of a Gaussian beam with waist w_0 at our center wavelength, $\lambda = 550\text{nm}$. These rays are traced through our 2D or 3D ply geometry, as we randomly sample a new ray direction from a rough dielectric BSDF at each intersection with a fiber surface, controlled by a fiber roughness parameter α_{fiber} , until the path escapes the ply.

The results are binned according to the exiting direction, as well as the number and types of scattering events in each path, producing three separate far field distributions, f^{miss} , f^{refl} , and f^{high} . Among them, f^{miss} corresponds to incident rays that never intersect any fiber and directly propagate to the far field, f^{refl} describes paths that only contain reflection events—paths that never travels to the interior of fibers. Lastly, f^{high} records all other paths, describing multiple scattering of all orders. We independently record f^{miss} in order to better study the scattering from fibers, and we distinguish




Fig. 7. Comparison between wave and ray simulations. While scattering distributions from a single ply predicted by wave and ray methods look very different, the average scattering distributions over more (50) ply instances show much stronger agreement. The only major difference between ensemble-averaged wave and ray results occurs from off-center illumination, where ray simulations predict much stronger forward scattering, due to not addressing edge diffraction.

between f^{refl} and f^{high} because it makes it easier to model plies with different amount of absorption. With the same ply geometry, f^{refl} is independent from the fibers' absorption coefficient while f^{high} varies with it. In practice, in each simulation we compute this high-order scattering distribution for a collection of absorption coefficients σ_k ($k = 0, 1, \dots, K - 1$), resulting in $f_0^{\text{high}}, f_1^{\text{high}}, \dots$, and f_{K-1}^{high} . This is achieved by tracking the total path length L in the interior of fibers, and weight its contribution to f_k^{high} by $e^{-\sigma_k L}$.

Our 2D path tracing, just like our wave simulations, is performed on fiber cross sections but supports out-of-plane incident angles, and the resulting distributions are defined in the specular cone (recall Section 4.3), as 1D functions of the outgoing azimuthal angle. Thus, our 2D path tracing results are directly comparable to 2D wave simulation results, as we will present in the next section. Moreover, our 3D path tracing on general, twisted fibers are what we rely on for building our ply model, and we dive into these 3D simulation results as well as our modeling process in Section 5.

4.5 Wave-Ray Comparison

As mentioned earlier, our 2D simulations serve as experiments for ply scattering computation in the special case of plies with untwisted fibers. With wave simulation data as ground truth, we compare our wave and ray results to study when path tracing provides sufficient accuracy and to identify key differences when it does not.

We first compare scattering distributions from a single ply instance. For path tracing, we use the full distribution given by $f^{\text{miss}} + f^{\text{refl}} + f^{\text{high}}$ from Section 3.4, computed with $\alpha_{\text{fiber}} = 0$ —i.e. assuming specular reflection or refraction at each surface interaction—since wave simulations use exact fiber geometry and do not model implicit fiber roughness. As shown in Fig. 7(a), with results from path tracing alongside wave simulations at $\lambda = 400\text{nm}$ and $\lambda = 694\text{nm}$, the two methods yield notably different predictions. Path tracing exhibits several sharp peaks absent in wave-based results, and more importantly, wave simulations reveal strong wavelength dependence in the scattering distribution, whereas path tracing lacks any notion of wavelength and produces a single, wavelength-independent result.

Nevertheless, once averaged over many ply instances, the scattering distributions from wave and ray simulations become more

similar. As shown in Fig. 7(b), when the incident beam strikes near the center of the plies, the 50-instance average scattering distributions from path tracing and wave simulations at $\lambda = 400\text{nm}$ and $\lambda = 694\text{nm}$ yield much stronger agreement. Notably, even the wave-based results exhibit minimal wavelength dependence in the averaged scattering. In Fig. 7(c), where the beam hits near the ply edge, ray and wave predictions remain closely aligned. The key difference is that path tracing allows rays that miss the ply to travel directly to the far field, producing a sharp forward scattering peak. In contrast, wave simulations diffract the unoccluded portion of the beam at the ply edge, yielding a broader angular spread around the forward direction—one that varies with wavelength.

Our 2D wave-ray comparison therefore indicates that, for ensemble averages of ply scattering distributions, path tracing is generally an accurate approach for modeling scattering from fibers. However, our ply model also requires a ply-edge diffraction component under off-center illumination, which must be computed wave optically.

Lastly, we believe the demonstrated accuracy of 2D ray simulations in computing ensemble-averaged multi-fiber scattering generalizes well to 3D. Therefore, we use 3D path tracing to generate data underlying our single-ply scattering model. These simulations are tractable using efficient, off-the-shelf ray tracing libraries [Wald et al. 2014], and we further optimized runtime through path reuse.

5 MULTI-LEVEL PLY SHADING MODEL

As concluded in Section 4, we use 3D path tracing as the foundation for building our ply scattering model. Each 3D simulation produces scattering data for a single ply, a single incident beam, and a full sphere of outgoing directions. To support ply shading under arbitrary illumination and within a fabric composed of many plies and yarns, we develop a multi-level model as follows:

- We compute ensemble-averaged single-ply scattering distributions over a grid of incident parameters (θ_i, h) (5.1)
- For each incident condition, we fit the resulting outgoing distribution to an interpolated Gaussian mixture model (5.2)
- We extend our single-ply model to multi-ply configurations through numerical convolutions (5.3)
- We further convolve our single-ply or multi-ply model to account for the two yarn layers in woven fabrics (5.4)

Table 2. Important notations in Section 5.

G	Anisotropic spherical Gaussian (ASG) lobe
$\theta_z, \phi_z, \phi_{xy}$	ASG lobe axis parameters
μ_1, μ_2	ASG lobe bandwidth parameters
c	ASG lobe amplitude parameters
\mathbf{q}_l	Umbrella notation for all ASG parameters
f^{single}	Single-ply scattering distribution
f^{multi}	Multi-ply scattering distribution

5.1 Tabulated Single-Ply Scattering Distributions

The ensemble-averaged single-ply scattering data from our 3D simulations can be viewed as a function of five variables: the incident beam, described by (θ_i, h) ; the outgoing direction, defined by (θ_o, ϕ_o) ; and for high-order scattering within fibers, the fiber absorption coefficient σ . Following discussions in Section 4.4, we denote these tabulated, high dimensional scattering distributions as $f^{\text{miss}}(\theta_i, h; \theta_o, \phi_o)$, $f^{\text{refl}}(\theta_i, h; \theta_o, \phi_o)$, and $f_k^{\text{high}}(\theta_i, h; \theta_o, \phi_o)$, where each f_k^{high} corresponds to an absorption value σ_k we simulate.

f^{miss} captures rays that bypass the fibers entirely. This component vanishes when the beam is fully occluded by the ply—i.e. when $|h| \leq R - d_{\text{cut}}$ (recall Section 4.2). For larger $|h|$, f^{miss} contains a sharp forward scattering peak unique to path tracing, where rays that miss the ply continue directly to the far field without undergoing diffraction. Because this component lacks diffraction and produces an overly sharp forward scattering peak, we do not retain the full 4D distribution $f^{\text{miss}}(\theta_i, h; \theta_o, \phi_o)$. Instead, we integrate it over all outgoing directions to obtain a scalar energy fraction $r_d(\theta_i, h)$ for each incident condition, which represents the total portion of the beam that escapes without interacting with any fiber. In particular, this value is zero for small $|h|$. The scalar $r_d(\theta_i, h)$ represents the energy that should be redirected into an edge-diffraction term f^{edge} in our model, whose construction is deferred to Section 6.1.

Moreover, f^{refl} describes the reflection-only scattering events from the fibers, and f_k^{high} corresponds to other, higher-order scattering within the ply. Each of these is a 4D function sampled over a grid of (θ_i, h) and several hundreds of values for θ_o and ϕ_o . As storing these 4D tables is impractical, we fit analytical approximations to each of these tabulated functions. These fitted models provide a compact and accurate representation of the scattering behavior, while supporting smooth interpolation across all the 5 dimensions.

5.2 Anisotropic Spherical Gaussian Fitting

Our analytical representations for the reflection-only and higher-order scattering distributions are constructed using the anisotropic spherical Gaussian (ASG) function [Xu et al. 2013]. For each incident configuration (θ_i, h) , we fit the corresponding 2D slice of the 4D scattering table—i.e. $f^{\text{refl}}(\theta_i, h; \omega_o)$ or $f_k^{\text{high}}(\theta_i, h; \omega_o)$ —to a single ASG lobe or a mixture of lobes. These local fits provide smooth analytical approximations over outgoing directions $\omega_o \sim (\theta_o, \phi_o)$, enabling fast evaluation. Moreover, to obtain a continuous scattering function across the incident domain, we need to interpolate between ASG fits at different values of (θ_i, h) , and for high-order scattering, interpolate across the absorption coefficient σ_k .

First invented in [Xu et al. 2013], an ASG function is defined with respect to unit vectors \mathbf{v} on the unit sphere:

$$G(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mu_1, \mu_2, c; \mathbf{v}) = c \cdot \max(\mathbf{v} \cdot \mathbf{z}, 0) \cdot e^{-\mu_1(\mathbf{v} \cdot \mathbf{x})^2 - \mu_2(\mathbf{v} \cdot \mathbf{y})^2} \quad (4)$$

where \mathbf{z} is the lobe center, and \mathbf{x}, \mathbf{y} are the tangent and bi-tangent axes of the ASG lobe. μ_1, μ_2 are positive values representing the lobe’s bandwidths for x and y axes, and c is the lobe amplitude. In our implementation, we parameterize $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$ using three scalars $(\theta_z, \phi_z, \phi_{xy})$ (see supplemental document).

Our reflection-only distribution f^{refl} exhibits a clean, single-lobe structure at each incident configuration (θ_i, h) , as it primarily represents single surface reflection from the illuminated region of the ply. Therefore, we can fit each 2D slice in f^{refl} to one single ASG lobe, such that for each incident parameter sample (θ_i^m, h^n) , we have:

$$f^{\text{refl}}(\theta_i^m, h^n; \omega_o) \approx G(\theta_z^{mn}, \phi_z^{mn}, \phi_{xy}^{mn}, \mu_1^{mn}, \mu_2^{mn}, c^{mn}; \omega_o) \quad (5)$$

where the six lobe parameters are optimized using the L-BFGS-B algorithm [Zhu et al. 1997]. Note that m and n are superscripts rather than exponents. The fitted lobe parameters vary smoothly with the incident samples (θ_i^m, h^n) . Thus, for an arbitrary (θ_i^*, h^*) , we can perform bicubic interpolation on the precomputed parameters $(\theta_z^{mn}, \phi_z^{mn}, \phi_{xy}^{mn}, \mu_1^{mn}, \mu_2^{mn}, c^{mn})$, obtaining parameters that define the reflection lobe corresponding to (θ_i^*, h^*) .

In contrast, fitting the high-order scattering data f_k^{high} requires a mixture of ASG lobes, as f_k^{high} incorporates all scattering events where light travels into fibers, resulting in complex structures in the scattering distribution. Also, this ASG mixture fitting can no longer be performed independently for each (θ_i^m, h^n) , because lobes fitted to neighboring incident conditions may represent entirely different scattering features, making it difficult to interpolate the fitted parameters across the incident domain in a meaningful way.

To address this challenge, we propose a joint fitting technique that uses the notion of “fitting blocks.” Specifically, we partition our discretized incident domain into overlapping 3×3 neighborhoods, i.e. our fitting blocks. Each block contains nine adjacent incident configurations, and within each block, we fit a mixture of L ASG lobes to each of the nine data slices, but only optimize the lobe parameters at the four corners. The five remaining slices are assigned parameters by averaging the fitted parameters for the corner slices, ensuring that each fitted lobe refers to a consistent scattering feature across the block. This fitting technique is first applied to f_0^{high} , representing high-order scattering in non-absorptive fibers ($\sigma_0 = 0$).

The idea of fitting blocks is illustrated in Fig. 8 (a), and mathematically, each fitting block, indexed (\tilde{m}, \tilde{n}) , includes nine 2D slices:

$$f_0^{\text{high}}(\theta_i^{2\tilde{m}+a}, h^{2\tilde{n}+b}; \omega_o) \quad \text{for } a, b \in \{0, 1, 2\} \quad (6)$$

On these data slices, we perform a large-scale ASG fitting so:

$$f_0^{\text{high}}(\theta_i^{2\tilde{m}+a}, h^{2\tilde{n}+b}; \omega_o) \approx \sum_{l=0}^{L-1} G_l(\mathbf{q}_{l,a,b}^{\tilde{m}\tilde{n}}; \omega_o) \quad \text{for } a, b = 0, 1, 2 \quad (7)$$

where \mathbf{q} is an umbrella notation representing the six parameters $(\theta_z, \phi_z, \phi_{xy}, \mu_1, \mu_2, c)$. Importantly, only $\mathbf{q}_{l,0,0}^{\tilde{m}\tilde{n}}$, $\mathbf{q}_{l,0,2}^{\tilde{m}\tilde{n}}$, $\mathbf{q}_{l,2,0}^{\tilde{m}\tilde{n}}$, and $\mathbf{q}_{l,2,2}^{\tilde{m}\tilde{n}}$ are our free parameters to be optimized, whereas the other five sets of parameters are forced to be averages (linear combinations

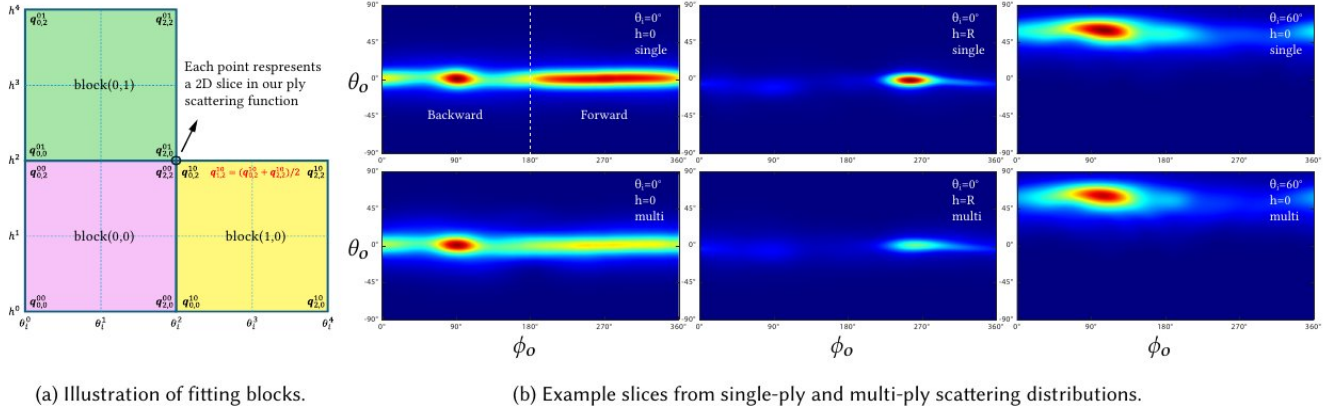


Fig. 8. Illustration of our ply scattering distributions and ASG lobe fitting. (a) Fitting blocks in Section 5.2, where each point in the block structure refers to a 2D slice in our ply scattering distribution. (b) Slices from the scattering distributions of a polyester ply containing circular fibers, featured for three incident conditions. The top row demonstrates single-ply scattering, and the bottom row presents corresponding multi-ply scattering distributions for a 3-ply yarn.

with weights 1/2 or 1/4, see supplemental document) of these corner parameters. By design, this guarantees that for incident parameters (θ_i^*, h^*) such that $\theta_i^* \in [\theta_i^{2\tilde{m}}, \theta_i^{2\tilde{m}+2}]$ and $h^* \in [h^{2\tilde{n}}, h^{2\tilde{n}+2}]$, the 2D slice $f_0^{\text{high}}(\theta_i^*, h^*; \omega_o)$ can be constructed as an ASG mixture with bilinearly interpolated parameters. Globally, this leads to a piecewise-interpolatable analytical approximation of f_0^{high} .

Lastly, to handle absorption, we reuse the ASG lobes from f_0^{high} and apply scalar weights per lobe to model attenuation in f_k^{high} for $k \geq 1$. These weights are optimized using least squares and can be interpolated linearly across σ . Full details and equations are provided in the supplemental document. One limitation of our piecewise fitting is that the resulting representation of f_k^{high} is not globally continuous. In particular, 2D slices located at block corners can belong to multiple overlapping blocks and therefore have multiple valid ASG mixture fits. However, using $L = 8$ lobes, we find that each of these fits is individually accurate, and we have never observed visual artifacts due to this discontinuity in any of our renderings.

In summary, our single-ply scattering model, temporarily excluding ply-edge diffraction, can be written as:

$$f^{\text{single}}(\theta_i, h, \sigma; \omega_o) = f^{\text{refl}}(\theta_i, h; \omega_o) + f^{\text{high}}(\theta_i, h, \sigma; \omega_o) \quad (8)$$

where the reflection term f^{refl} is modeled by a single ASG lobe, and the high-order term f^{high} is represented as a mixture of ASG lobes with absorption-dependent amplitudes. All ASG parameters can be interpolated over (θ_i, h) , and f^{high} also supports lobe amplitude interpolation across σ . Some visualizations of our single-ply scattering distributions can be found in Fig. 8 (b) (top row).

5.3 Multi-Ply Scattering Extension

Our model f^{single} can be directly used for rendering fabrics with single-ply yarns; however, when our fabric consists of multi-ply yarns, we need to extend our single-ply scattering to account for multiple scattering among plies. Prior work [Zhu et al. 2023b] modeled scattering from a multi-ply yarn through performing aggregation with their ply scattering function, formulating yarn-level forward

and backward scattering lobes from ply-level scattering function parameters. In contrast, our single-ply model—represented using numerous ASG lobes—is more expressive but contains less clean separation between forward and backward scattering, and deriving multi-ply scattering requires a more simulation-driven approach.

Fig. 9 (a) illustrates a multi-ply yarn. When one ply is illuminated by an incident beam described with (θ_i, h) , this *primary* ply scatters light according to $f^{\text{single}}(\theta_i, h, \sigma; \omega_o)$, while some of the scattered light may hit another *secondary* ply before escaping to the far field, and thus be redirected or absorbed. To model relevant multiple scattering events, we split our multi-ply scattering into:

- Direct, unoccluded scattering from the primary ply
- Scattering that includes *one* additional interaction with a secondary ply before reaching the far field
- Remaining energy from more inter-ply interactions

Mathematically, for a yarn with P plies, we can write:

$$f^{\text{multi}}(\theta_i, h, \sigma; \omega_o) = f^{\text{direct}}(\theta_i, h, \sigma; \omega_o) + \sum_{p=1}^{P-1} f_p^{\text{conv}}(\theta_i, h, \sigma; \omega_o) + f^{\text{resi}}(\theta_i, h, \sigma; \omega_o) \quad (9)$$

The single-to-multi extension needs to be performed for different (θ_i, h) with respect to the primary ply and for different absorption coefficients, and each resulting 2D slices in f^{multi} are temporarily tabulated with respect to ω_o . To compute each component in Eq. 9, we first simulate shadowing between plies in 2D, as illustrated in Fig. 9 (a). For each azimuthal angle ϕ_o , we trace out a set of rays from the primary ply (indexed 0), and measure what fraction $r_0(\phi_o)$ of these rays escapes and what fraction $r_p(\phi_o)$ hits each of the other plies p . Following our shadowing tests, we compute f^{direct} as:

$$f^{\text{direct}}(\theta_i, h, \sigma; \omega_o) = r_0(\phi_o) \cdot f^{\text{single}}(\theta_i, h, \sigma; \omega_o) \quad (10)$$

The secondary scattering of the outgoing light from the primary ply 0 by another ply p is computed through “numerically convolving” two single-ply scattering functions describing ply 0 and ply p .

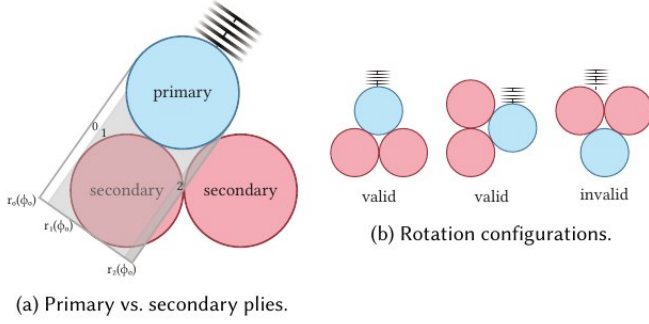


Fig. 9. Multi-ply yarns. (a) The ply illuminated by an incident beam is the primary ply, and other plies, which may rescatter the outgoing light from the primary ply, are secondary plies. (b) A rotation configuration of the plies is valid if the primary ply can be directly seen by the incident beam, which illuminates it from the upright direction in the ply space. In other words, a primary ply cannot be “buried” under other plies in a valid configuration.

On the high level, the outgoing energy from ply 0 along each direction ω_o is weighted by $r_p(\phi_o)$ and treated as incident energy on p , with incident direction and position relative to p determined from coordinate transformations. This energy is rescattered according to our single-ply model—evaluated on ply p , and accumulating the rescattered energy corresponding to all outgoing directions from ply 0 gives us a convolved scattering function. This convolved scattering function is further weighted down after another 2D shadowing test from ply p , retaining only the rescattered energy that is not again occluded. This rescattering term is denoted f_p^{conv} in Eq. 9, and more details on computing f_p^{conv} are in the supplemental document.

Since higher-order effects are too costly to simulate, we instead keep track of the fraction of light scattered and absorbed in previously considered ply interactions that involve at most one secondary ply, and then use a geometric series to estimate the total energy that should be piped into a more diffuse residual term f^{resi} , which incorporates all further inter-ply scattering (see supplemental).

Notably, special care is needed when the incident position h is close to the ply edge—when a portion of the incident beam misses the primary ply, it either misses the yarn entirely or lands upon another ply in the yarn. The former case will be addressed by our ply-edge diffraction modeling (see Section 6.1), and for the latter case, where there are essentially two plies illuminated by the incident beam, we perform our direct and secondary scattering analysis twice and combine together the two resulting distributions.

Furthermore, we repeat our computation on a set of representative “valid rotation configurations” and average the results, as explained in Fig. 9 (b). This improves our modeling of the expected response from the whole yarn, when one ply, whose relative position to other plies varies due to ply twisting, is illuminated. The final multi-ply model averaged across various ply configurations—still in its tabulated form—can then be refitted to ASG lobes as in Section 5.2, resulting in an identical form as the single-ply model¹. At render time, f^{multi} is used in place of f^{single} on fabrics with multi-ply yarns.

¹We maintained separation of reflection-only and higher-order scattering throughout our multi-ply computations. Moreover, the energy that should be redirected to the ply-edge diffraction term in the multi-ply model is tracked, and the diffracted energy is

Fig. 8 (b) presents some example slices in our single-ply and multi-ply model, under a few incident conditions. For each (θ_i, h) , our multi-ply scattering function resembles its single-ply counterpart along backward scattering directions but appears weaker and blurrier along for forward scattering—due to increased inter-ply interactions associated with forward scattering of the primary ply.

5.4 Double-Layer Scattering

Our previous sections provide a directionally detailed model for single-ply and multi-ply scattering, allowing us to compute scattering from a single yarn. However, to characterize the overall color, especially transmitted color, of an entire fabric, it is also essential to account for scattering between the two yarn layers in a woven cloth. Since this double-layer scattering involves interactions between yarns that are approximately perpendicular, we expect directionality to be significantly reduced in resulting distributions. Thus, we perform analysis in terms of scalar reflection and transmission factors for the two yarn layers, which we compute through our single-ply or multi-ply model, averaged across incident positions.

We provide a full description of our double-layer scattering model in the supplemental document. In brief, we simulate two steps of scattering through a top and bottom yarn layer, estimate their reflection and transmission factors, and derive total forward and backward scattering by modeling inter-reflection events with a geometric series. These scalar corrections allow us to account for additional absorption and scattering in double yarn layers and result in a more diffuse model for fabric transmission.

6 WAVE OPTICS AUGMENTATION

In this section, we augment our ply scattering model with wave optics. Specifically, we introduce models for diffraction around and between plies, and for instance-to-instance variations in ply scattering. Both are crucial for improving the realism of rendered fabrics.

Table 3. Important notations in Section 6.

A, x_w, x_e, y_s, y_n	Quantities describing a diffraction aperture
ξ_i, ξ_o	Incident and outgoing directions
K, f_x, f_y	Kirchhoff diffraction integral related quantities
f^{gap}	Weave-gap diffraction term
f^{edge}	Ply-edge diffraction term
r_d	Ply-edge diffraction energy fraction
$N_{2D,r}^{(n)}, N_{2D,h}^{(n)}$	2D wave optics noise function instances
$f_{\text{refl}}^{(n)}, f_{\text{high}}^{(n)}$	Noised instances of ply scattering distributions

6.1 Diffraction Modeling

Diffraction comes into play in two cases. First, when the incident beam center misses all the plies and passes through a gap between yarns. Second, when the beam partially overlaps with a ply, i.e. an off-center hit. In ray optics, both cases are modeled with delta impulses, but wave optics predicts diffraction, producing blurrier

strictly reduced in the multi-ply case, since an incident beam that misses one ply can be captured by another ply, therefore not considered diffracted.

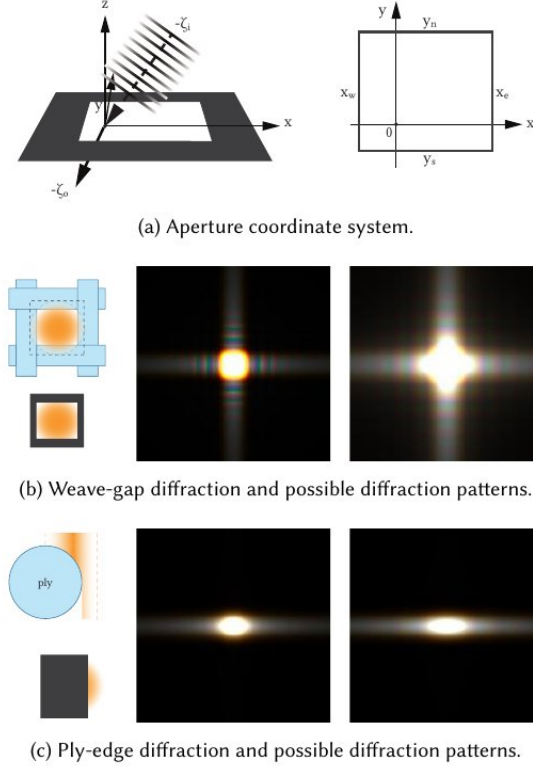


Fig. 10. Illustration of our aperture coordinate system and some diffraction patterns corresponding to weave-gap diffraction and ply-edge diffraction.

transmission profiles. We address these with two related models—*weave-gap diffraction* and *ply-edge diffraction*—based on the far field diffraction of a Gaussian beam through a rectangular aperture.

We formulate our problem in a coordinate system aligned to the beam and the aperture, as in Fig. 10 (a). The aperture is the rectangle $[x_w, x_e] \times [y_s, y_n]$ in the xy -plane, the beam's propagation direction is given by $-\xi_i$, where (as before) the unit vector $\xi_i = \langle \xi_i^x, \xi_i^y, \xi_i^z \rangle$ points away from the aperture, and the origin is located where the beam center intersects the aperture plane.

With our incident field modeled as a Gaussian-windowed plane wave, we compute the forward diffracted field intensity along each direction $\xi_o = \langle \xi_o^x, \xi_o^y, \xi_o^z \rangle$ ($\xi_o^z < 0$) in the far field limit, using the Kirchhoff diffraction integral [Heurtley 1973]. As fully derived in our supplemental material, this diffracted intensity is given by

$$I(A, \lambda, \xi_i; \xi_o) \approx \frac{1}{4\lambda^2} (|\xi_i^z| + |\xi_o^z|)^2 K^2(A, f_x, f_y) \quad (11)$$

where $A = (x_w, x_e, y_s, y_n)$ represents the aperture, while $f_x = (\xi_i^x + \xi_o^x)/\lambda$ and $f_y = (\xi_i^y + \xi_o^y)/\lambda$. Moreover, we have:

$$K(A, f_x, f_y) = \int_{x_w}^{x_e} \int_{y_s}^{y_n} e^{-\frac{x^2+y^2}{w_0^2}} \cdot e^{-2\pi i(f_x \cdot x + f_y \cdot y)} dy dx \quad (12)$$

Using Eq. 11 and Eq. 12, we can readily compute the far field diffraction pattern for any aperture $A = (x_w, x_e, y_s, y_n)$, wavelength λ , and incident direction ξ_i , but the computation is too expensive to

perform during rendering. Instead, we create tabulations separately for weave-gap diffraction and ply-edge diffraction.

For weave-gap diffraction, we precompute and tabulate K as functions of (f_x, f_y) for different aperture configurations A . Since ξ_i, ξ_o are unit vectors, we have $|\xi_i^x + \xi_o^x|, |\xi_i^y + \xi_o^y| \leq 2$, and our wavelengths are between 400nm and 700nm, bounding f_x, f_y in a relatively small range. Moreover, we only need to evaluate weave-gap diffraction if A satisfies:

$$-d_{\text{cut}} \leq x_w < 0 < x_e \leq d_{\text{cut}}; \quad -d_{\text{cut}} \leq y_s < 0 < y_n \leq d_{\text{cut}} \quad (13)$$

recalling from Section 4.2 that at points farther away than d_{cut} to the beam center, the beam intensity is negligible. In such cases, further enlarging the aperture has no effect on the diffraction pattern, so we clamp our aperture edge distances to d_{cut} .

Also, we find that apertures of the same size give rise to similar diffraction patterns, and the relative position of the beam center to aperture edges merely changes the pattern's color saturation (see the supplemental material for an example). Following these observations, we only tabulate K for a collection of apertures satisfying

$$x_w = -x_e; \quad y_s = -y_n; \quad 0 < x_e, y_n \leq d_{\text{cut}} \quad (14)$$

and apply a color saturation correction heuristic when rendering.

With K precomputed, the weave-gap diffraction model f^{gap} is:

$$f^{\text{gap}}(A, \lambda, \xi_i; \xi_o) = \frac{1}{4\Phi_i(A) \cdot |\xi_i^z|} \left[\frac{|\xi_i^z| + |\xi_o^z|}{\lambda} \cdot K(A, f_x, f_y) \right]^2 \quad (15)$$

where the normalization factor $\Phi_i(A) \cdot |\xi_i^z|$ is the beam's total incident power on the aperture, with

$$\Phi_i(A) = \int_{x_w}^{x_e} \int_{y_s}^{y_n} e^{-(x^2+y^2)/w_0^2} dy dx \quad (16)$$

precomputed for different apertures and $|\xi_i^z|$ representing the cosine falloff in non-normal incidence.

Moreover, ply-edge diffraction, as a component of the ply scattering model, is evaluated in the ply coordinate system and required when part of the incident beam extends beyond the edge of the ply, i.e., $R - d_{\text{cut}} \leq h < R$ (recall Section 4.2). A simple sketch in our supplemental document reveals that for a beam with incident position h on the ply, the relevant aperture A is given by

$$A = (R - h, d_{\text{cut}}, -d_{\text{cut}}, d_{\text{cut}}) \quad (17)$$

for $h > 0$, and $h < 0$ can be handled symmetrically. Moreover, comparing Fig. 10 (b) and (c) shows that ply-edge diffraction patterns are much simpler and smoother, and edge diffraction primarily serves to broaden the angular divergence of the forward propagating beam along the ply space's x -axis.

Therefore, for each wavelength λ and incident position h , we compute the ply-edge diffraction pattern using Eq. 11, Eq. 12, and Eq. 17 and approximate the resulting pattern with one ASG lobe. As this diffraction lobe is centered in the forward direction and aligned with the ply space's x -axis, we only need to store the ASG bandwidths $\mu_1(\lambda, h), \mu_2(\lambda, h)$. Our ply-edge diffraction term is then

$$f^{\text{edge}}(\lambda, \theta_i, h; \omega_o) = r_d(\theta_i, h) \cdot G(-\theta_i, \frac{3\pi}{2}, 0, \mu_1, \mu_2, \frac{\mu_1\mu_2}{\pi}; \omega_o) \quad (18)$$

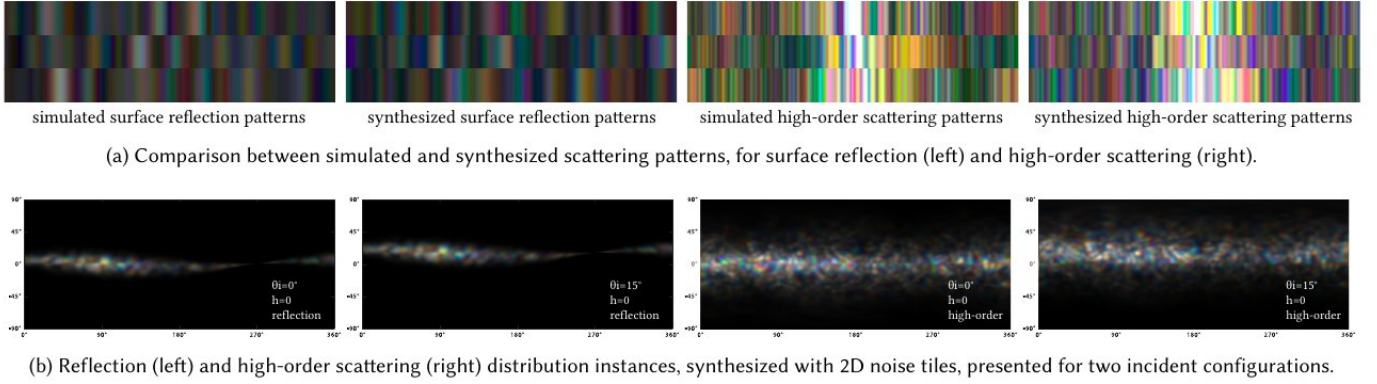


Fig. 11. Illustrations of our 1D and 2D noise functions and synthesized scattering distributions. Our 1D noise functions can synthesize fiber scattering patterns very similar to those computed from wave simulations. We can generalize these noise functions to 2D and use them to augment our ply scattering model defined on the entire unit sphere. In our noise application, we try to preserve the speckle pattern details in surface reflection patterns as they shift with the incident light direction, but do not need to model this memory effect in higher-order scattering.

where $r_d(\theta_i, h)$ is the total energy in this edge-diffraction term, first introduced in Section 5.1². The ASG lobe parameters follow the definition from Section 5.2, and $\mu_1 = \mu_1(\lambda, h)$, $\mu_2 = \mu_2(\lambda, h)$ are pre-computed and interpolated over λ and h . $\mu_1\mu_2/\pi$ is an approximated normalization factor of the ASG lobe such that G integrates to 1 [Xu et al. 2013]. Thus, f^{edge} , which corresponds to the portion of an off-center beam that misses the ply, completes our ply model.

6.2 Wave Optics Noise Functions

The ply scattering model in Section 5 captures average behavior across similar plies, but without per-instance variation, it will produce overly uniform appearance if used in rendering. To address this, we introduce a wave optics informed noise model that augments scattering distributions based on single-instance simulations.

Because the structures of our idealized yarns and plies are different from those in real fabrics, where fibers are more irregularly arranged, we focus on modeling fiber-to-fiber variation in directional scattering from a specific fiber type, providing a distribution that can be used at an appropriate spatial scale to capture the spatio-angular variation of fabrics made from that type of fiber. For each given fiber type (characterized by size and shape), we perform 2D wave simulations on 200 single fiber instances, where in each simulation the fiber scatters light from a wide Gaussian beam. We repeat each simulation for wavelengths between 400nm and 700nm, and for both black (highly absorbing) and white (non-absorbing) fibers. The black fibers produce mostly surface reflection, and subtracting their scattering distributions from those of their white counterparts yields the high-order scattering component.

To summarize the variation across instances and wavelengths, we design a 1D noise function that matches the statistical properties of fiber scattering distributions, following Yu et al. [2024]. Grounded in speckle theory [Goodman 2007; Steinberg and Yan 2022; Xia et al.

2023], our noise function takes the form

$$N_{1D}^{(n)}(\lambda, \phi_o) = \frac{1}{2} [r_{1D}^{(n)}(\lambda, \phi_o)]^2 + \frac{1}{2} [i_{1D}^{(n)}(\lambda, \phi_o)]^2 \quad (19)$$

where $r_{1D}^{(n)}$ and $i_{1D}^{(n)}$ are independent Gaussian processes with

$$r_{1D}^{(n)}(\lambda, \phi_o), i_{1D}^{(n)}(\lambda, \phi_o) \sim \mathcal{N}(0, 1) \quad (20)$$

This guarantees unit expectation at each (λ, ϕ_o) :

$$\mathbb{E}[N_{1D}^{(n)}(\lambda, \phi_o)] = 1 \quad (21)$$

A synthesized scattering distribution is then obtained by multiplying the mean distribution $\bar{I}(\lambda, \phi_o)$ with this noise:

$$I^{(n)}(\lambda, \phi_o) = \bar{I}(\lambda, \phi_o) \cdot N_{1D}^{(n)}(\lambda, \phi_o) \quad (22)$$

where $I^{(n)}$ is a synthesized scattering distribution, and \bar{I} is computed from averaging all simulated instances.

Moreover, as in [Yu et al. 2024], $r_{1D}^{(n)}$ and $i_{1D}^{(n)}$ are generated using a 1D Gaussian process with dual-point statistics governed by a power spectral density (PSD) function, derived from the angular correlations in our simulated scattering data. To further maintain correlation across wavelengths, we generate noise only at $\lambda = 400\text{nm}$ and produce longer-wavelength variants by rescaling and resampling, following the same prior approach.

In Fig. 11 (a), we compare 1D angular scattering patterns synthesized from our noise model against real simulation results. For both single reflection and higher-order scattering, the synthesized and simulated results show visually similar intensities and chromatic banding, validating the statistical accuracy of our approach. Building on this, we generalize the noise model to 2D angular domains—over outgoing directions on the unit sphere—so it can be applied to our full ply scattering model. Details of this 2D construction are provided in the supplemental document.

For rendering, we generate two sets of 2D noise tiles $N_{2D,r}^{(n)}(\lambda, \omega_o)$, $N_{2D,h}^{(n)}(\lambda, \omega_o)$ for modeling reflection and higher-order scattering. At each shading point, we select noise instances based on a hash

²This energy portion is decreased in our multi-ply model, as briefly discussed in our previous footnote. With multiple plies in a yarn, light that leaks from one ply may be captured and scattered by another ply and not considered diffracted.

function of the ply ID and fiber position. These noise instances are queried to modulate the average ply scattering functions via:

$$\begin{aligned} f_{\text{refl}}^{(n)}(\lambda, \theta_i, h; \omega_o) &= f^{\text{refl}}(\theta_i, h; \omega_o) \cdot N_{2D,r}^{(n)}(\lambda, \omega'_o) \\ f_{\text{high}}^{(n)}(\lambda, \theta_i, h, \sigma; \omega_o) &= f^{\text{high}}(\theta_i, h, \sigma; \omega_o) \cdot N_{2D,h}^{(n)}(\lambda, \omega_o) \end{aligned} \quad (23)$$

where for single reflection, we preserve the memory effect [Xia et al. 2023]—a phenomenon where the speckle pattern shifts predictably with the illumination angle—by computing ω'_o from ω_i and ω_o to reflect this shift (see supplemental document). High-order scattering does not exhibit this effect, and we find that querying $N_{2D,h}^{(n)}$ directly with ω_o (no ω_i dependence) yields reasonable results. Fig. 11 (b) demonstrates some noised scattering distributions from our method, featuring two incident conditions, $(\theta_i = 0^\circ, h = 0)$ and $(\theta_i = 15^\circ, h = 0)$. For the same ply instance, the single reflection patterns (left) are almost identical up to a shift, thanks to the memory effect.

Lastly, we mention that due to the high dimensions of our ply scattering function and the complex correlations in scattering across different dimensions, designing an entirely physically based noise function is extremely challenging. Our noise function design and application scheme is physically informed while containing approximations, but it results in realistic appearance in practice and is easily implemented in rendering—as a multiplication operation.

7 IMPLEMENTATION AND RENDERING

In this section, we describe our fabric rendering implementation, demonstrate how our ply-based model captures diverse material appearance under front-lit and back-lit conditions, and showcase the visual effects enabled by wave optics.

7.1 Fabric Shading Implementation

We implemented our ply-based shading model as well as a surface-based fabric rendering pipeline in Mitsuba3. Each fabric is represented with two tangent textures for yarn directions, two position textures containing each shading point’s position along and across a yarn, a yarn ID texture encoding the weave pattern, and some additional parameters describing yarn widths and ply twists. Each ply scattering model is specified by a collection of parameter tables and tabulated patterns: fitted ASG lobe parameters, double-layer scattering correction factors, edge-diffraction bandwidths, weave-gap diffraction patterns, and 2D wave optics noise tiles.

To use the shading model, we follow our beam-based pipeline described in Section 4–6. At a fabric shading point, we reconstruct the local ply geometry and set up the ply coordinate frame to determine the incident beam parameters (θ_i, h) . Based on the value of h , we select one of the three cases:

- (a) $|h| < R - d_{\text{cut}}$: the beam hits a ply near its center. We evaluate f^{single} or f^{multi} depending on the number of plies (5.2 & 5.3).
- (b) $R - d_{\text{cut}} \leq |h| < R$: the beam hits near a ply edge. We blend f^{single} or f^{multi} with the edge-diffraction term f^{edge} (6.1).
- (c) No ply intersection. We estimate distances to nearby plies to define an aperture, query the corresponding weave-gap diffraction pattern f^{gap} (6.1). If a nearby ply is close enough such that $|h| < R + d_{\text{cut}}$, we also weight in its scattering using f^{single} or f^{multi} .




Fig. 12. Comparison of photographs of four different fabric samples with renderings produced by our model and the model from [Zhu et al. 2023b]. While both models capture the overall reflection well, only our model can simulate the glinty or irregular appearance on these samples.

If both yarn layers are intersected in (a) or (b), we apply our double-layer energy adjustment (5.4). We also apply wave optics noise to modulate the our ply scattering functions (6.2).

Our shading model operates in a spectral renderer and requires fabric color to be specified via an absorption spectrum rather than an albedo. The procedure we use to convert colors to absorption coefficients is detailed in the supplemental document.

Lastly, our scattering model supports importance sampling by selecting ASG lobes proportionally to their weights and drawing samples from their Gaussian distributions—an approach that is especially effective when lobes are sharp.

7.2 Rendering Fabrics with Various Materials

We first demonstrate our fabric model using four samples with distinct weaves and fiber materials, and include three addition samples




Fig. 13. Transmission through the fabric samples from Fig. 12. Delta transmission leads to unrealistically sharp forward scattering and reveals the light source behind the sample. Our modeling of diffraction through weave gaps produces significantly closer match with the reference images.

in the supplemental document. For each sample, we estimated ply geometry from measuring detached yarns and performed scattering simulations to derive ply models (see supplemental). Notably, measured fiber counts per ply are much lower than those reported in [Zhao et al. 2016], which used statistics for sewing threads.

Each fabric was photographed under calibrated front-lit and back-lit setups. We reproduced these conditions in scenes and rendered with our ply models. For comparison, we rendered the same scenes using the model in [Zhu et al. 2023b]; this method is the closest to ours, as we share a similar single-ply to multi-ply structure. Parameters such as fiber count and refractive index were shared between models, and less observable parameters such as the transparency in [Zhu et al. 2023b] were fitted to the best of our abilities. We converted their ply model to a fabric model by including a delta transmission component [Zhu et al. 2023a] to simulate gaps, and added a double-layer scattering feature as in [Khattar et al. 2025].

As shown in Fig. 12, both models match front-lit references reasonably well. The main difference lies in the lack of ply-to-ply variation in the prior method, which leads to flat or overly smooth appearance, especially in the orange polyester. Our model, by contrast, captures the glinty, spatial variation seen in photos.

Back-lit comparisons in Fig. 13 reveal stronger differences. While the translucent organza remains bright under fixed exposure, the three opaque fabrics appear nearly black in photos. Prior methods produce overly transparent renderings for these, even with double-layer scattering addressed. Reducing transparency further required

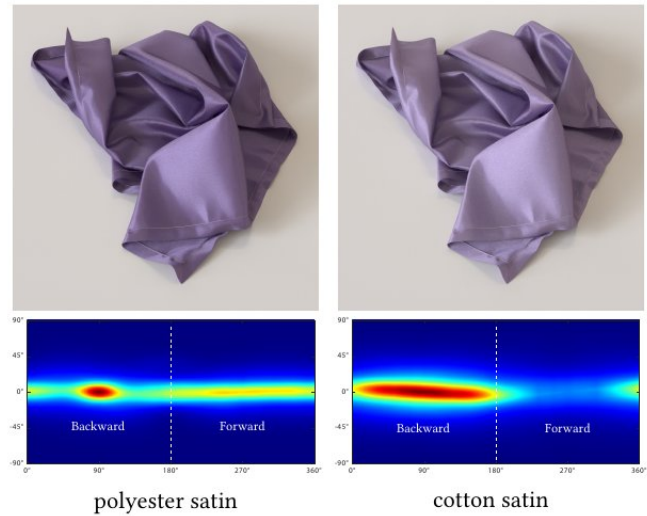


Fig. 14. Renderings of a polyester satin and a cotton satin, using ply models derived from nearly identical fiber scattering simulations, with the only difference being the simulated fiber shapes. Compared to polyester, the cotton satin features stronger backward scattering and less absorption, and also exhibits slightly softer highlights.

unrealistically high fiber counts (10000+). In contrast, while our model also overestimates the brightness in transmission images, it produces convincingly more opaque results for thick fabrics and recovers the distinctive cross-shaped highlight in the loose organza. Replacing our beam-gap handling with delta transmission also results in unrealistically sharp area light outlines (3rd column), while our combination of weave-gap diffraction and edge scattering from nearby plies significantly softens forward scattering. Our double-layer modeling further makes the transmission more diffuse, yielding smooth appearance (4th column) that matches reference photos.

Furthermore, we analyze the visual impact of fiber shapes by comparing polyester and cotton satin fabrics rendered with ply models that differ only in fiber shapes. The same fiber absorption spectrum was used for rendering both fabrics, and as shown in Fig. 14, our polyester and cotton fabrics exhibit different shades of purple. Specifically, the cotton satin appears lighter and features more backward scattering, most likely owing to its hollow-core fibers. Slight differences in highlight softness are also visible. However, the overall similarity between the two renderings suggests that other factors like fiber roughness and twist likely have an even larger influence on fabric appearance.

7.3 More Wave-Related Visual Effects

We now highlight additional visual effects enabled by wave optics, beyond those already demonstrated in prior sections. Fig. 15 shows fabrics rendered with and without our wave optics noise functions. Without noise, rendered fabrics appear smooth and less organic, resembling a far field appearance despite using near field ply geometry, while our noise-based irregularities improve realism with negligible cost. More results rendered with and without our wave




Fig. 15. Several fabrics rendered without and with our wave optics based structural irregularities. While the overall appearance is similar, the visual break-ups introduced with our noise functions significantly improve realism.

optics noise are in the supplemental material, where we also compare the effects of using our sophisticated, wave-driven noise model and adopting less controlled, pseudo-random noise.

We also explore subtle color phenomena sometimes observed in real fabrics under back-lighting. While our organza sample in Fig. 13 does not show color fringing under a large area light, we were able to capture pink and green tints (Fig. 2, last two) on some back-lit organza fabrics in real life, when using a very narrow light source. These colors are unlikely caused by the weave-gap diffraction modeled in our framework, since the wide gaps in organza fabrics should yield near-delta transmission that lacks chromatic features. Instead, we think that these colors more likely arise from diffraction through small gaps between fibers or between twisted plies, which our current model does not explicitly capture.

Therefore, although not formally part of our model, we implemented a bonus feature to reproduce the subtle pink and green color patterns that sometimes appear in loose fabrics back-lit with a very small light source. We modeled a rectangular aperture that is narrow in one dimension ($< 5\mu m$) and unbounded in the other, mimicking thin gaps between locally aligned fibers. Using the method from Section 6.1, we computed the resulting diffraction pattern, which exhibits a linear shape and strong chromatic effects, as seen in Fig. 16. We blended this pattern with our regular scattering model, assuming roughly 25% of the forward scattering arises from diffraction through such tiny inter-fiber or inter-ply gaps (slightly abusing our ply-level shading pipeline). This 25% / 75% mix produced plausible color effects in the back-lit organza rendering in Fig. 16. While not meant to match any specific photo, our rendered color patterns resemble those in Fig. 2 (last two). A formal integration of this bonus feature is left for future work.

8 DISCUSSION

In this section, we briefly discuss limitations, design decisions, and generalizations of our current model.

8.1 Ideal Structure Assumptions

Our ply-based fabric shading model makes several simplifying assumptions about yarn structures. In particular, while real plies can




Fig. 16. A back-lit silk organza sample, rendered by combining a strongly colored diffraction distribution with our regular ply model. These color effects are likely due to the small gaps between fibers and twisted plies.

have irregular cross sections and twist patterns, we assume circular ply and ideal helical twisting for both plies and fibers. These assumptions greatly simplify our modeling—for example, our ply scattering functions depend only on the difference between azimuthal angles ($\phi_i - \phi_o$). Nonetheless, our simulation framework generalizes easily to less ideal settings. As a light example, we compute and visualize 2D slices of average scattering functions from plies with three different cross section shapes. As shown in Fig. 17, all three examples contain roughly the same number of fibers, and all scattering plots correspond to an incident angle of $\theta_i = 0^\circ$. For each ply shape, we sample five incident positions across the ply's horizontal extent. The resulting scattering functions show visible differences across ply types and incident positions—especially for off-center hits—demonstrating how cross-sectional geometry affects angular scattering. A full investigation of ply scattering using non-ideal geometry is left for future work.

8.2 Modeling Choice and Flexibility

As our method is partially data-driven, it offers less flexibility for editing compared to fully analytical approaches. For instance, our model does not allow on-the-fly changes to some properties, including roughness, IOR, and fiber twist, without re-running the simulations and re-fitting the ASGs. However, we have designed the system so that key appearance parameters, such as absorption, remain outside the simulations and can even be textured as needed. For fiber twist, it may be possible—albeit with some loss of accuracy—to simulate an untwisted fiber and apply analytical coordinate frame rotations during rendering, as in [Zhu et al. 2023b].

Our ASG mixture model is an analytical representation that can be easily implemented in any renderer. Alternatively, one could construct a semi-analytical neural model based on simulation results. Such a model may be trained on sparse simulation data and depend on a larger number of geometric and structural parameters, avoiding the need to interpolate between fitted distributions and potentially allowing more parameters to be varied on the fly. Still, our work is orthogonal to neural techniques, since we are primarily interested

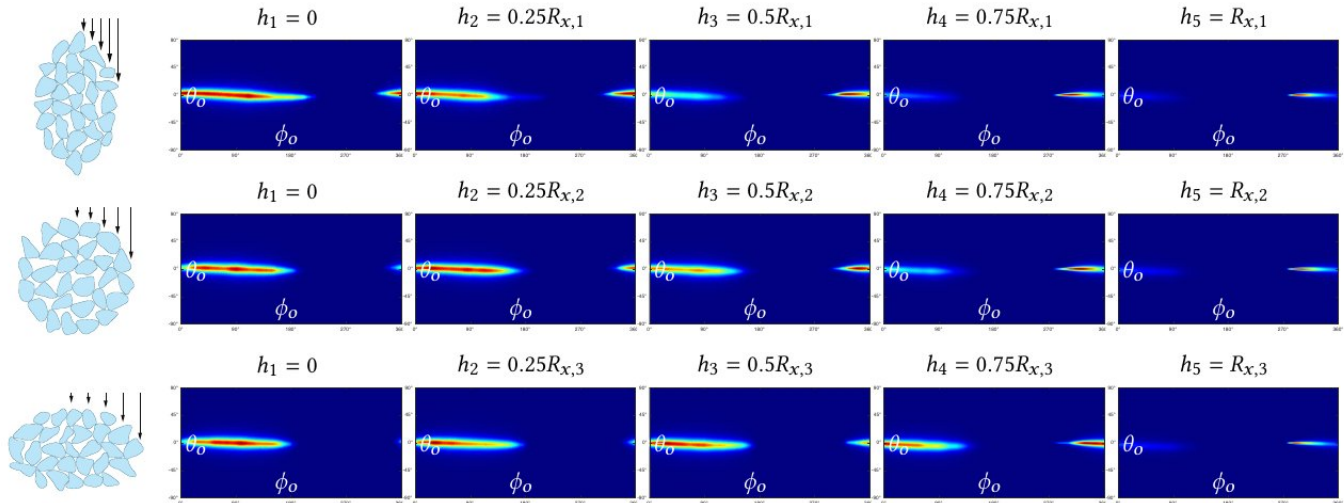


Fig. 17. Scattering from plies with elliptical cross sections. Each row shows scattering distributions (at $\theta_i = 0^\circ$) from a different ply shape, all containing a similar number of fibers. We sample five incident beam positions h_1 – h_5 across each ply’s horizontal extent. The resulting angular distributions vary across shapes and incident positions, illustrating the sensitivity of scattering to ply cross section geometry.

in fabric appearance and the influence of wave optics, rather than data representations.

8.3 Transmission

Transmission through opaque fabrics presents a particular challenge for simulation-based methods. The small amount of light that passes through typically results from high-order scattering involving multiple yarn interactions or occasional gaps due to slight variations in the weave. While our results are closer to the reference photographs than those produced by the previous method, we still tend to overestimate transmission in dense fabrics. One explanation is that real fabric structures differ from our helical fiber and ply model, especially in interior regions where yarns are tightly compressed. Such deviations could increase the optical thickness of the yarns and reduce the size of the weave gaps.

9 CONCLUSION

In this paper, we presented the first fabric appearance model that incorporates wave optics. Since full-wave simulations on large fiber assemblies remain prohibitively expensive even with our accelerations, we based our model on path tracing simulations. While ray and wave simulations yield different results for individual ply instances, their ensemble-averaged scattering distributions are remarkably similar—aside from the absence of diffraction in ray-based results. We addressed this by augmenting our ply scattering model with diffraction components grounded in wave optics, and introduced wave-based noise functions to model spatial variation. Comparisons with photographs show that our model closely reproduces reference images, with fabric irregularities and imperfections playing a key role in enhancing realism.

We therefore believe our model marks an important step toward more realistic fabric rendering. Looking ahead, we aim to overcome the current limitations posed by the simulation-driven nature of

our method, enabling applications such as inverse rendering and interactive material editing.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under grant IIS-2212084. We give special thanks to Longxiulin Deng for her help on our supplemental video and several result images.

REFERENCES

- Neeharika Adabala, Nadia Magnenat-Thalmann, and Guangzheng Fei. 2003. Real-Time Rendering of Woven Clothes. In *Proceedings of the ACM symposium on Virtual reality software and technology* (Osaka, Japan) (VRST’03). Association for Computing Machinery, New York, NY, USA, 41–47. <https://doi.org/10.1145/1008653.1008663>
- Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A. Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. 2017. An Appearance Model for Textile Fibers. *Comput. Graph. Forum* 36, 4 (July 2017), 35–45. <https://doi.org/10.1111/cgf.13222>
- Michael Ashikhmin and Simon Premoze. 2007. *Distribution-based BRDFs*. Technical Report. University of Utah. <https://api.semanticscholar.org/CorpusID:1624053>
- Alexis Benamira and Sumanta Pattanaik. 2021. A combined scattering and diffraction model for elliptical hair rendering. *Computer Graphics Forum* 40 (07 2021), 163–175.
- Elizabeth Bleszynski, Maria Bleszynski, and Thomas Jaroszewicz. 1996. AIM: Adaptive integral method for solving large-scale electromagnetic scattering and radiation problems. *Radio Science* 31, 5 (1996), 1225–1251.
- Carlos Castillo, Jorge López-Moreno, and Carlos Aliaga. 2019. Recent advances in fabric appearance reproduction. *Computers & Graphics* 84 (2019), 103–121. <https://doi.org/10.1016/j.cag.2019.07.007>
- Calvin Chang, Brian Ginn, Natalie K. Livingston, Zhicheng Yao, Benjamin Slavin, Martin W. King, Sangwon Chung, and Hai-Quan Mao. 2020. 1.4.6 - Medical Fibers and Biotextiles. In *Biomaterials Science (Fourth Edition)* (fourth edition ed.), William R. Wagner, Shelly E. Sakiyama-Elbert, Guigen Zhang, and Michael J. Yaszemski (Eds.). Academic Press, 575–600. <https://doi.org/10.1016/B978-0-12-816137-1.00038-6>
- Xiang Chen, Lu Wang, and Beibei Wang. 2024. Real-time Neural Woven Fabric Rendering. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH’24). Association for Computing Machinery, New York, NY, USA, Article 63, 10 pages. <https://doi.org/10.1145/3641519.3657496>
- Maria Lurdes Dias. 1991. Ray tracing interference color. *IEEE Computer Graphics and Applications* 11, 02 (1991), 54–60.
- Leroy D Dickson. 1970. Characteristics of a propagating Gaussian beam. *Applied Optics* 9, 8 (1970), 1854–1861.
- Janus Egholm and Niels J Christensen. 2006. Rendering compact discs and other diffractive surfaces illuminated by linear light sources. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia*

- and Southeast Asia. Association for Computing Machinery, New York, NY, USA, 329–332.
- Alejandro Conty Estevez and Christopher Kulla. 2017. Production Friendly Microfacet Sheen BRDF. In *ACM SIGGRAPH 2017 Courses*. Association for Computing Machinery, New York, NY, USA.
- Elena Garces, Victor Arellano, Carlos Rodriguez-Pardo, David Pascual-Hernandez, Sergio Suja, and Jorge Lopez-Moreno. 2023. Towards Material Digitization with a Dual-scale Optical System. *ACM Trans. Graph.* 42, 4, Article 152 (July 2023), 13 pages. <https://doi.org/10.1145/3592147>
- Walton C Gibson. 2021. *The method of moments in electromagnetics*. Chapman and Hall/CRC.
- Jay S Gondek, Gary W Meyer, and Jonathan G Newman. 1994. Wavelength dependent reflectance functions. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. Association for Computing Machinery, New York, NY, USA, 213–220. <https://doi.org/10.1145/192161.192202>
- Joseph W Goodman. 2007. *Speckle phenomena in optics: theory and applications*. Roberts and Company Publishers.
- John C Heurtley. 1973. Scalar Rayleigh–Sommerfeld and Kirchhoff diffraction integrals: a comparison of exact evaluations for axial points. *Journal of the Optical Society of America* 63, 8 (1973), 1003–1008.
- Hideki Hirayama, Kazufumi Kaneda, Hideo Yamashita, and Yoshimi Monden. 2001. An accurate illumination model for objects coated with multilayer films. *Computers & Graphics* 25, 3 (2001), 391–400.
- M.M. Houck. 2009. 2 - Ways of identifying textile fibers and materials. In *Identification of Textile Fibers*, Max M. Houck (Ed.). Woodhead Publishing, 6–26. <https://doi.org/10.1533/9781845695651.1.6>
- Isabelle Icart and Didier Arques. 1999. An illumination model for a system of isotropic substrate-isotropic thin film with identical rough boundaries. In *Rendering Techniques '99: Proceedings of the Eurographics Workshop in Granada, Spain, June 21–23, 1999*. Springer, 261–272.
- Isabelle Icart and Didier Arques. 2000. A physically-based BRDF model for multilayer systems with uncorrelated rough boundaries. In *Rendering Techniques 2000: Proceedings of the Eurographics Workshop in Brno, Czech Republic, June 26–28, 2000*. Springer, 353–364.
- Piti Irawan and Steve Marschner. 2012. Specular reflection from woven cloth. *ACM Transactions on Graphics (TOG)* 31, 1 (2012), 1–20.
- Wenzel Jakob, Jonathan T. Moon, Adam Arbree, Kavita Bala, and Steve Marschner. 2010. A Radiative Transfer Framework for Rendering Materials with Anisotropic Structure. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 29, 10 (July 2010), 53:1–53:13. <https://doi.org/10.1145/1778765.1778790>
- Wenhua Jin, Beibei Wang, Milos Hasan, Yu Guo, Steve Marschner, and Ling-Qi Yan. 2022. Woven Fabric Capture from a Single Photo. In *SIGGRAPH Asia 2022 Conference Papers* (Daegu, Republic of Korea) (SA '22). Association for Computing Machinery, New York, NY, USA, Article 33, 8 pages.
- Apoorv Khattar, Junqiu Zhu, Ling-Qi Yan, and Zahra Montazeri. 2025. A Texture-Free Practical Model for Realistic Surface-Based Rendering of Woven Fabrics. *Computer Graphics Forum* 44, 1 (2025), e15283. <https://doi.org/10.1111/cgf.15283>
- Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. *ACM Trans. Graph.* 35, 1 (2015), 1–1.
- Pramook Khungurn, Rundong Wu, James Noeckel, Steve Marschner, and Kavita Bala. 2017. Fast rendering of fabric micro-appearance models under directional and spherical gaussian lights. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 1–15.
- Ivana Markova. 2019. *Textile fiber microscopy a practical approach* (first edition. ed.). Wiley, Hoboken, NJ. <https://learning.oreilly.com/library/view/-/9781119320050/?ar>
- Zahra Montazeri, Soren Gammelmark, Henrik Wann Jensen, and Shuang Zhao. 2021. Practical Ply-Based Appearance Modeling for Knitted Fabrics. In *Eurographics Symposium on Rendering - DL-only Track*, Adrien Bousseau and Morgan McGuire (Eds.). The Eurographics Association. <https://doi.org/10.2312/sr.20211297>
- Zahra Montazeri, Soren B Gammelmark, Shuang Zhao, and Henrik Wann Jensen. 2020. A practical ply-based appearance model of woven fabrics. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–13.
- Naoki Okada, Dong Zhu, Dongsheng Cai, James B Cole, Makoto Kambe, and Shuichi Kinoshita. 2013. Rendering Morpho butterflies based on high accuracy nano-optical simulation. *Journal of Optics* 42 (2013), 25–36.
- Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Transactions on Graphics (TOG)* 32, 2 (2013), 1–12.
- Mirko Sattler, Ralf Sarlette, and Reinhard Klein. 2003. Efficient and realistic visualization of cloth. In *Rendering techniques*. Eurographics Association, Goslar, DEU, 167–178.
- K Schroeder, Reinhard Klein, and Arno Zinke. 2011. A volumetric approach to predictive rendering of fabrics. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 1277–1286.
- G. Y. Soh and Zahra Montazeri. 2023. Neural Appearance Model for Cloth Rendering. *Comput. Graph. Forum* 43 (2023), i–x. <https://api.semanticscholar.org/CorpusID:265043701>
- Jos Stam. 1999. Diffraction shaders. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. 101–110.
- Shlomi Steinberg, Ravi Ramamoorthi, Benedikt Bitterli, Eugene d'Eon, Ling-Qi Yan, and Matt Pharr. 2024a. A Generalized Ray Formulation For Wave-Optical Light Transport. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–15.
- Shlomi Steinberg, Ravi Ramamoorthi, Benedikt Bitterli, Arshiya Mollazainali, Eugene D'Eon, and Matt Pharr. 2024b. A Free-Space Diffraction BSDF. *ACM Transactions on Graphics (TOG)* 43, 4 (2024).
- 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.
- Yinlong Sun, F. David Fracchia, and Mark S. Drew. 1998. A Composite Model for Representing Spectral Functions. <https://api.semanticscholar.org/CorpusID:16718709>
- Yinlong Sun, F David Fracchia, Mark S Drew, and Thomas W Calvert. 2000. Rendering iridescent colors of optical disks. In *Rendering Techniques 2000: Proceedings of the Eurographics Workshop in Brno, Czech Republic, June 26–28, 2000*. Springer, 341–352.
- Yingjie Tang, Zixuan Li, Miloš Hašan, Jian Yang, and Beibei Wang. 2024. Woven Fabric Capture with a Reflection-Transmission Photo Pair. In *Proceedings of SIGGRAPH 2024*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3641519.3657410>
- Ingo Wald, Sven Woop, Carsten Benthin, Gregory S Johnson, and Manfred Ernst. 2014. Embree: a kernel framework for efficient CPU ray tracing. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–8.
- Beibei Wang, Wenhua Jin, Miloš Hašan, and Ling-Qi Yan. 2022. SpongeCake: A Layered Microflake Surface Appearance Model. *ACM Transactions on Graphics (TOG)* 42, 1 (2022), 1–16.
- Jiaping Wang, Shuang Zhao, Xin Tong, John Snyder, and Baining Guo. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 1–9. <https://doi.org/10.1145/1360612.1360640>
- Kui Wu and Cem Yuksel. 2017. Real-time Fiber-level Cloth Rendering. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D 2017)* (San Francisco, CA). ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3023368.3023372>
- Mengqi Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. 2023. A practical wave optics reflection model for hair and fur. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–15.
- 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.
- Kun Xu, Wei-Lun Sun, Zhao Dong, Dan-Yong Zhao, Run-Dong Wu, and Shi-Min Hu. 2013. Anisotropic spherical gaussians. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–11.
- Yunchen Yu, Andrea Weidlich, Bruce Walter, Eugene d'Eon, and Steve Marschner. 2024. Appearance Modeling of Iridescent Feathers with Diverse Nanostructures. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–18.
- Yunchen Yu, Mengqi Xia, Bruce Walter, Eric Michielssen, and Steve Marschner. 2023. A full-wave reference simulator for computing surface reflectance. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–17.
- Tizian Zeltner, Brent Burley, and Matt Jen-Yuan Chiang. 2022. Practical Multiple-Scattering Sheen Using Linearly Transformed Cosines. In *ACM SIGGRAPH 2022 Talks* (Vancouver, BC, Canada) (SIGGRAPH '22). Association for Computing Machinery, New York, NY, USA, Article 7, 2 pages. <https://doi.org/10.1145/3532836.3536240>
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building volumetric appearance models of fabric using micro CT imaging. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–10.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-aware synthesis for predictive woven fabric appearance. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.
- Shuang Zhao, Fujun Luan, and Kavita Bala. 2016. Fitting procedural yarn models for realistic cloth rendering. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–11.
- Ciyou Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. 1997. Algorithm 778: LBFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on mathematical software (TOMS)* 23, 4 (1997), 550–560.
- Junqiu Zhu, Christophe Hery, Lukas Bode, Carlos Aliaga, Adrian Jarabo, Ling-Qi Yan, and Matt Jen-Yuan Chiang. 2024. A Realistic Multi-scale Surface-based Cloth Appearance Model. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, Article 89, 10 pages. <https://doi.org/10.1145/3641519.3657426>
- Junqiu Zhu, Adrian Jarabo, Carlos Aliaga, Ling-Qi Yan, and Matt Jen-Yuan Chiang. 2023a. A Realistic Surface-based Cloth Rendering Model. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 5, 9 pages. <https://doi.org/10.1145/3588432.3591554>
- Junqiu Zhu, Zahra Montazeri, Jean-Marie Aubry, Ling-Qi Yan, and Andrea Weidlich. 2023b. A Practical and Hierarchical Yarn-based Shading Model for Cloth. *Computer Graphics Forum* (2023). <https://doi.org/10.1111/cgf.14894>