

Predicting Fabric Appearance Through Thread Scattering and Inversion Supplemental Document

Mengqi (Mandy) Xia, Zhaoyang Zhang, Sumit Chaturvedi,
Yutong Yi, Rundong Wu, Holly Rushmeier, Julie Dorsey

1 Thread Scattering Model

1.1 Fiber scattering model

Fiber scattering models represent individual fibers as cylinders and use the Bidirectional Curve Scattering Distribution Function (BCSDF) to characterize their scattering properties. Similar to the BSDF, the BCSDF describes outgoing radiance L_r as an integral of incident radiance L_i multiplied by the BCSDF S :

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

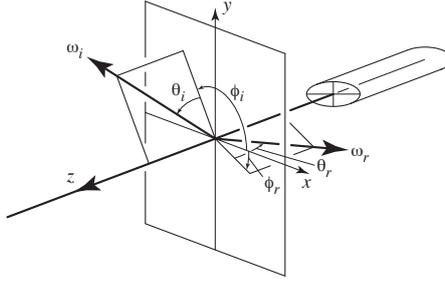


Figure 1: An illustration of the commonly adopted longitudinal-azimuthal parameterization for fiber scattering models. Each of the directions ω_i and ω_r in 3D is parameterized using the polar angle θ , defined as the angle between ω and the plane perpendicular to the cylinder axis, and the azimuthal angle ϕ , defined in that plane.

We express the BCSDF in spherical coordinates as $S(\theta_i, \theta_r, \phi_i, \phi_r)$, using the angles illustrated in Figure 1. Most existing models represent the BCSDF as a sum of reflective and transmissive modes S_p , where each mode S_p is factored into a longitudinal function M_p , an azimuthal function N_p , and an attenuation A_p :

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

$p = 0$ is the first reflection mode (R), $p = 1$ is the two transmission mode (TT), and $p = 2$ corresponds to the transmission, reflection, and transmission mode (TRT). Our model is based on [1], where the longitudinal function is from [2]

$$M_p(\theta_i, \theta_r) = \frac{1}{2v \sinh(1/v)} e^{-\frac{\sin \theta_i \sin \theta_r}{v}} I_0\left(\frac{\cos \theta_i \cos \theta_r}{v}\right), \quad (3)$$

with transformed longitudinal roughness v . The azimuthal function $N_p(\theta_i, \phi_i, \phi_r)$ evaluates a trimmed logistic function around the perfect reflection or transmission direction. The attenuation term $A_p(\theta_i)$ computes the Fresnel reflection and transmission contribution. The model relies on the Monte Carlo integration inherent in path tracing to integrate across the fiber width.

1.2 Multiple fiber scattering

We assume that azimuthal and longitudinal scattering are separable. To study azimuthal scattering behavior, we simulate ray interactions with a single circle and with multiple circles in the cross-sectional plane.

Single circle. We shoot one million rays from the left, uniformly distributed across the diameter and intersecting the circle. At each interaction, we compute the Fresnel factor and use a random number to determine whether reflection or refraction occurs. As the ray travels through the circle, we compute absorption according to Beer's law. Once the ray exits, we record its outgoing angle. Finally, we accumulate the energy of the outgoing rays into angular bins to obtain the azimuthal scattering distribution of a single fiber. The simulation naturally integrates over the width (i.e., diameter), as it does not account for the exact point of impact on the circle.

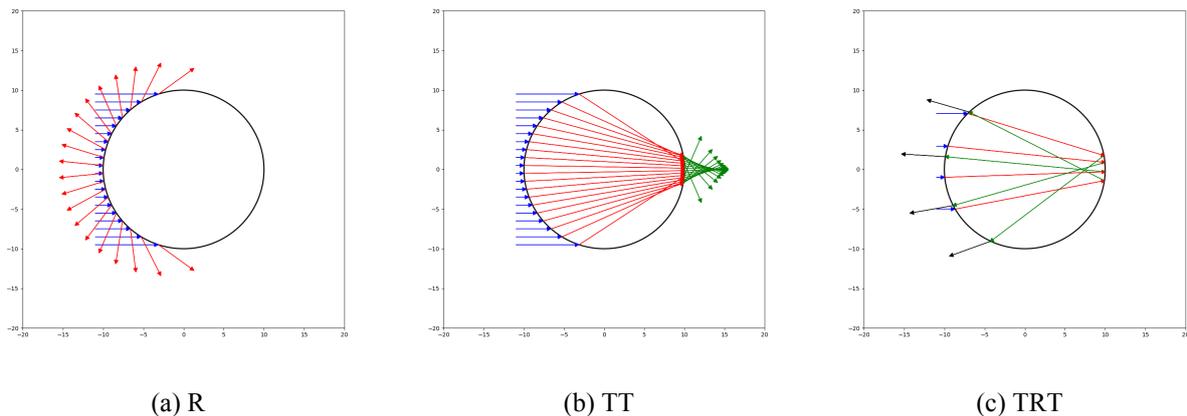


Figure 2: Illustrations of rays interacting with a circle

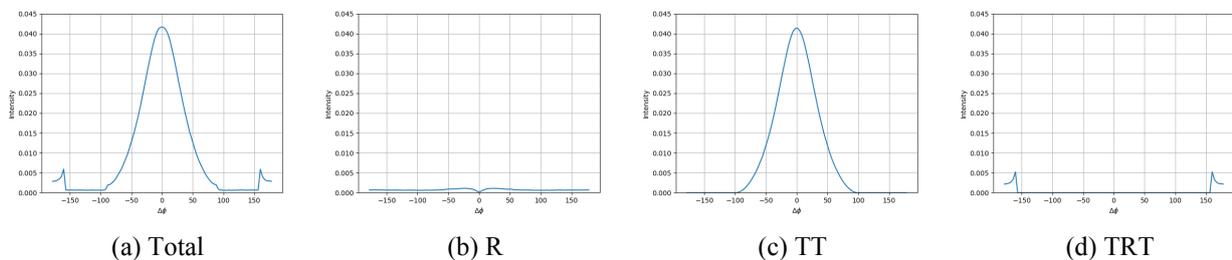


Figure 3: The angular distribution of single-circle scattering and its individual components

Multiple circles. To simulate light interacting with multiple fibers, we create small circles that represent cross sections of fibers randomly inside a circular boundary that represents the thread. We assign the radius (R) of the thread, the radius (r) of the fibers, the minimum gap (d) between the fibers and the maximum number of attempts (M) as parameters for the simulation. In each attempt, we try adding a new circle within the boundary, which will be successful if it is at least d distance away from all existing ones. This continues until all M attempts are exhausted.

For each cross-sectional configuration, we generate one million parallel rays pointing toward the thread, uniformly distributed across its diameter, to simulate specular reflection, transmission, and absorption. To account for the randomness inherent in any specific configuration, we generate 500 random circle arrangements per simulation and average the resulting angular distributions. This can be seen as approximating the average scattering behavior across different thread instances in real fabric.

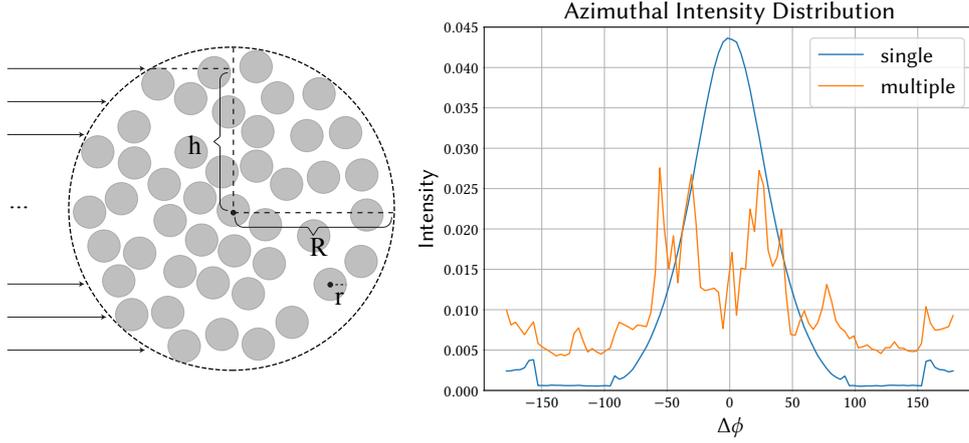


Figure 4: The left figure is an illustration of a multiple-circle configuration. compare the azimuthal scattering functions of single fiber scattering and multiple fiber scattering with different small fiber radii, averaged over random configurations.

Comparison. In Figure 5, we compare the scattering distributions of single-circle and multiple-circle cases across six different absorption levels, as shown in the subsequent plots. In each plot, the multiple-circle scattering is shown for small circle radii ranging from 0.5 to 4.0, while the large circle (thread) radius is fixed at 10.0. We observe that multiple-circle scattering exhibits stronger backscattering (corresponding to the R and TRT regions in the single-circle case) and weaker forward scattering (the TT region). As the radius of the small circles increases, the multiple-circle scattering distribution increasingly resembles that of the single-circle case. This is expected—at the limit, the configuration becomes a single circle, which corresponds exactly to the single-circle scattering case. Although the multiple-circle distribution differs from the single-circle one, it retains similar characteristics, as seen in the R, TT, and TRT lobes. Each lobe appears to be independently squashed, stretched, or scaled, and together they combine to form the overall multiple-circle scattering distribution.

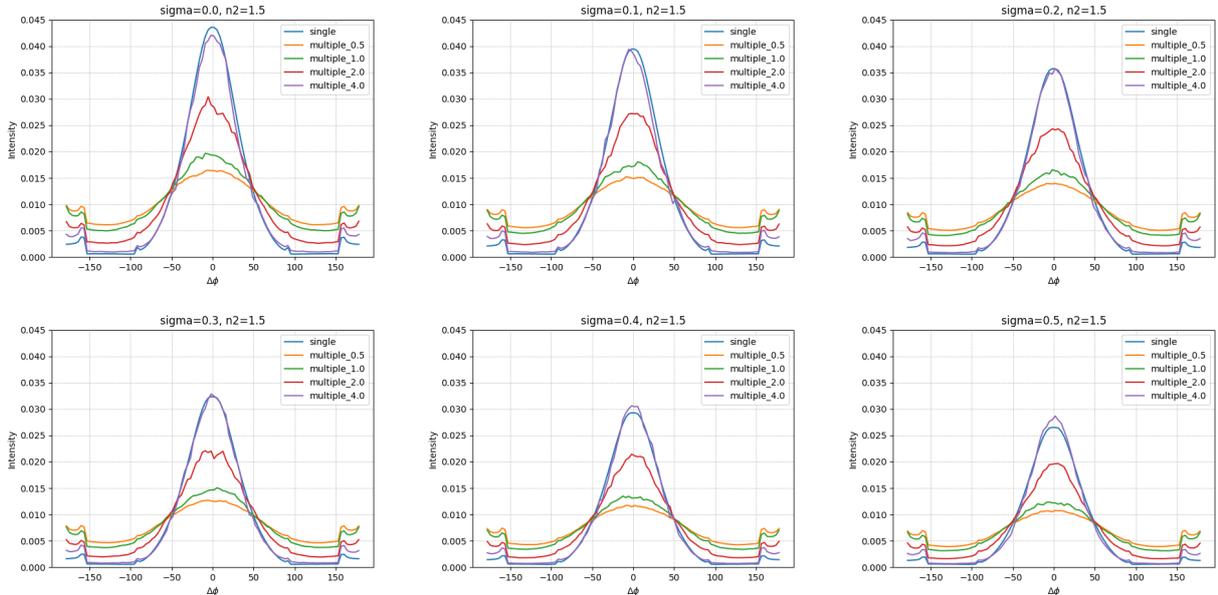


Figure 5: Azimuthal scattering function comparison

1.3 Thread BCSDf

Based on the simulation results, we propose a new model S_t for threads composed of multiple fibers, comprising three lobes: a low-order backward scattering mode adapted from the R mode in single-fiber scattering, a forward

scattering mode modifying the TT mode in single-fiber scattering, and a high-order backward scattering mode modifying the TRT mode in single-fiber scattering.

$$S_t(\theta_i, \theta_r, \phi_i, \phi_r) = \sum_{p=0}^2 M_{tp}(\alpha_p, \theta_i, \theta_r) N_{tp}(\beta_p, \theta_i, \phi_i, \phi_r) A_{tp}(\theta_i, \phi_i), \quad (4)$$

$$A_{t0} = \lambda_0(\theta_i), \quad A_{t1} = \lambda_1(\theta_i) T_r, \quad A_{t2} = \lambda_2(\theta_i) T_r$$

where α_p and β_p are the longitudinal and azimuthal roughness for each lobe p , providing independent control of the width of the distribution; λ_0 , λ_1 , and λ_2 are scale factors with longitudinal incident angle dependence. $T_r = e^{-\sigma_a l}$ is the transmittance through a single fiber, where σ_a is the absorption coefficient and l is the distance traveled within the fiber. When interacting with multiple fibers, the path length varies, but it is not analytically feasible to compute the exact average path length for each direction over random fiber configurations. Our representation effectively captures the equivalent scaling and width changes needed to account for path length variations in multiple-fiber scattering. We evaluate S_t in a non-integrated manner as proposed in [1]. However, we assume that the thread is subpixel in size and that the scattering function, integrated over the thread's width, reproduces the desired scattering behavior of multiple-fiber configurations.

To verify our model, we first perform azimuthal function fitting tests in 24 different configurations with different absorption coefficients and small circle radii (Figure 7), and for different longitudinal incident angles separately (top two rows of Figure 8). In each fitting, the target is the angular distribution obtained from the simulation. The parameters are the absorption coefficient (σ_a), the individual scale factors (λ_p), and the roughness parameters (α_p). The fitted curve is computed by evaluating 1000 different height (h) values uniformly spaced across the width. At each height, we evaluate the angular distribution of each lobe using the logistic distribution, compute the weighted sum using the scale factors, and average the distributions over h . We use the Adam optimizer to iteratively update the parameters until convergence. We observe that our proposed model fits the multiple scattering distributions very well, with only minor discrepancies near the ends of the curves.

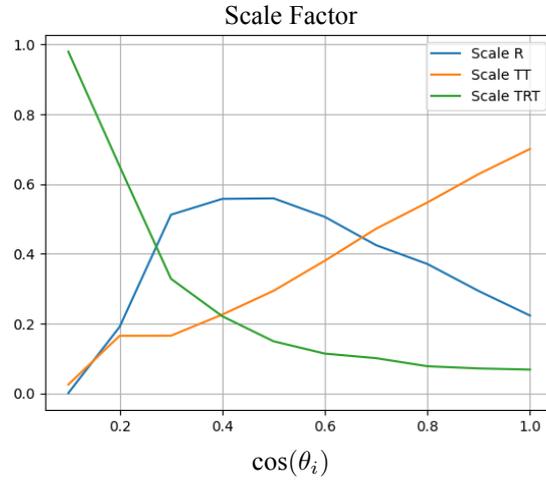


Figure 6: The scale factors as functions of $\cos(\theta_i)$.

We observed that the scale factors $\lambda_p(\theta_i)$ follow specific trends for three lobes (Figure 6), leading to the following parametrization of the scale factors

$$\begin{aligned} \lambda_0(\theta_i) &= r_{0a} \cos(\theta_i)^2 + r_{0b} \cos(\theta_i) + r_{0c} \\ \lambda_1(\theta_i) &= r_{1a} \cos(\theta_i) \\ \lambda_2(\theta_i) &= \frac{1}{r_{2a} \cos(\theta_i) + r_{2b}} + r_{2c}, \end{aligned} \quad (5)$$

With parameterized scale factors, we jointly fit different θ_i and present the results in the bottom two rows of Figure 8. Our model demonstrates significantly more accurate fitting results compared to the ply-level scattering model in [3]. As the incident angle becomes more oblique, we observe a large backscattering peak that

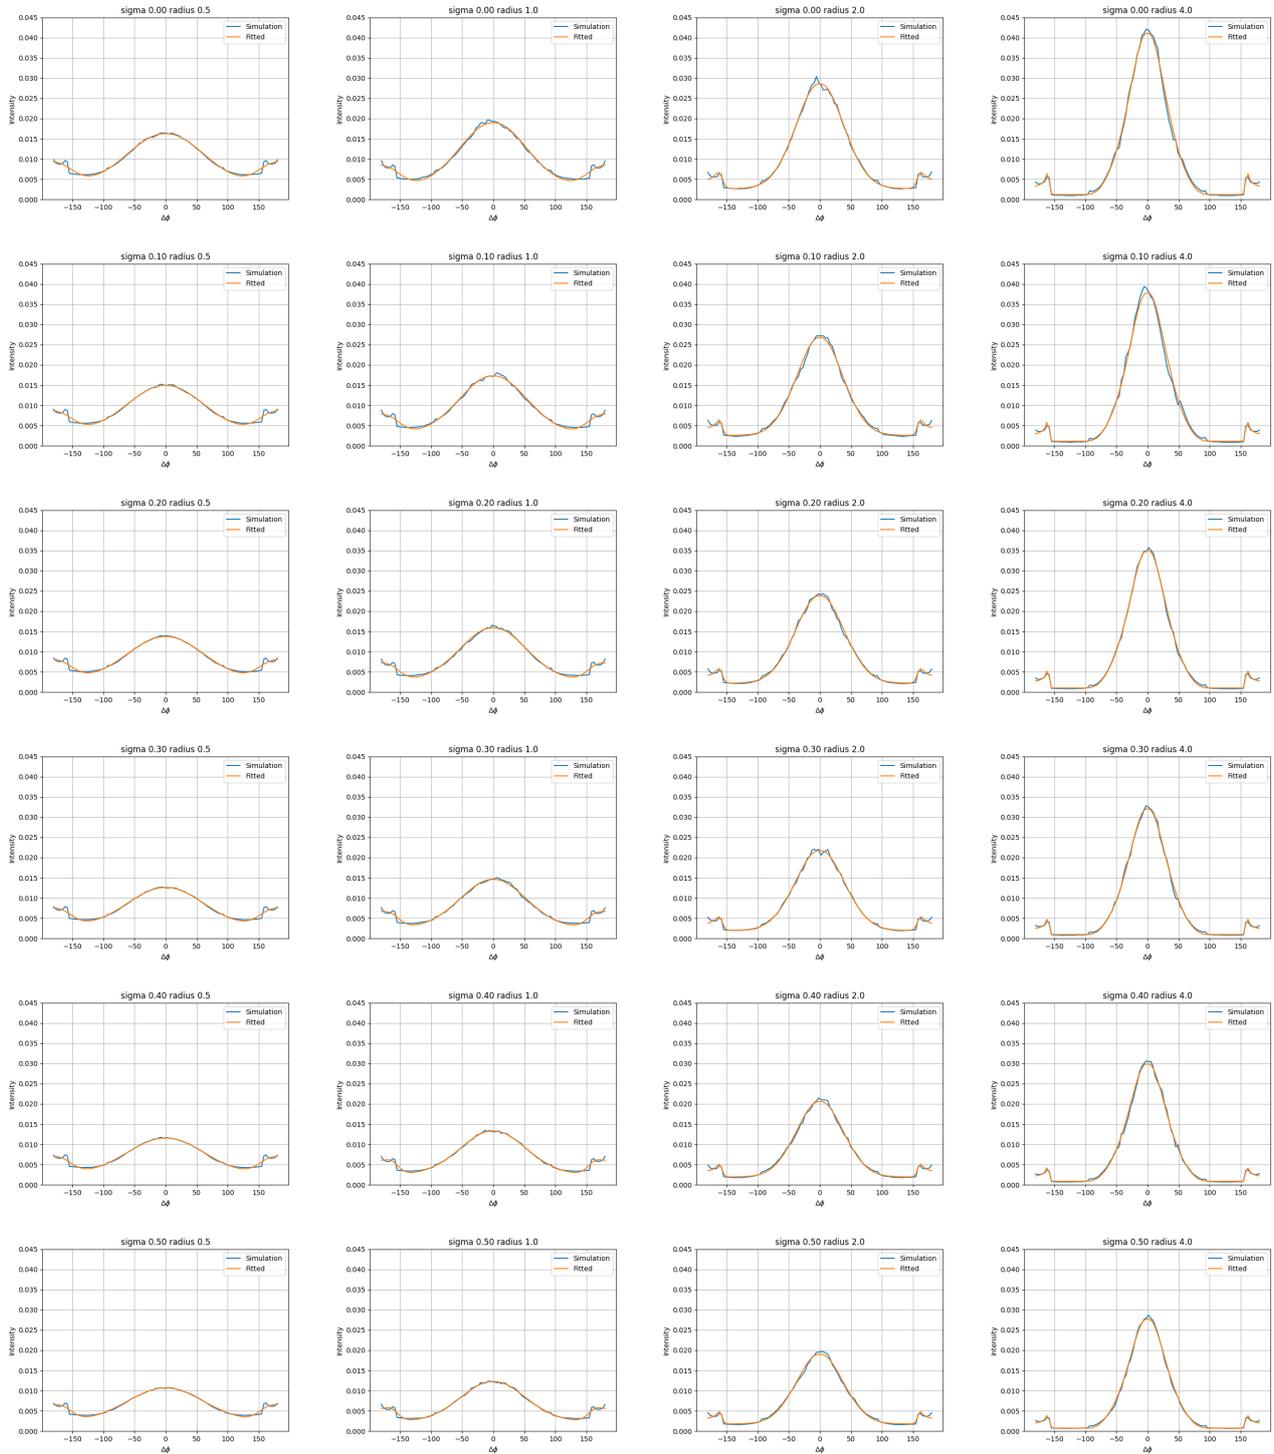


Figure 7: Fitting multiple circle scattering using independent single circle scattering lobes.

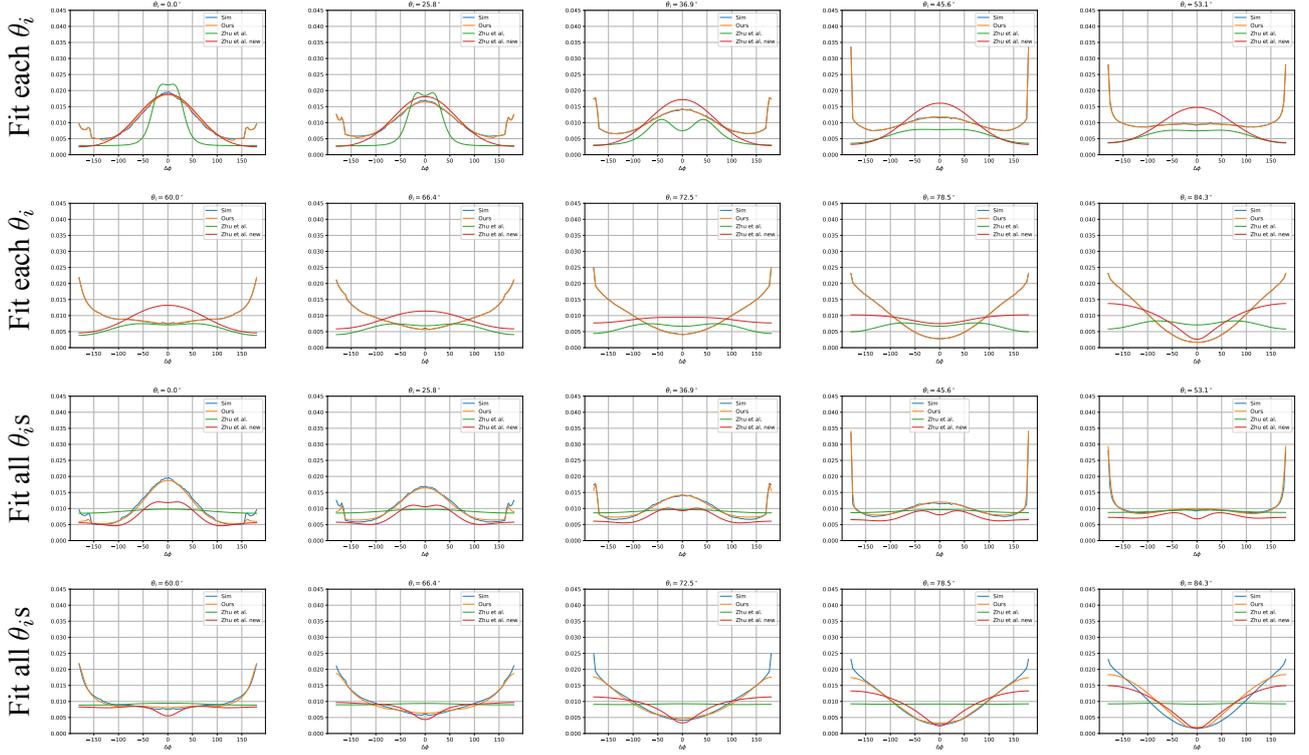


Figure 8: We compare the fitting results of our model, those from [3] (green curve) and our improvement on [3] to the simulated azimuthal functions (red curve). In the top two rows, we fit each incident longitudinal angle separately and in the bottom two rows we fit different incident angles jointly. This comparison shows that our model produces much more accurate fitting than previous work.

eventually dominates the distribution. However, in [3], the azimuthal R lobe at the fiber level and the aggregated ply-level backscattering azimuthal lobe are parameterized using a uniform distribution. This leads to a flat fitting result, compromising accuracy across different incident angles. Through ablation, we find that introducing a Gaussian parametrization for backscattering and discarding the aggregation terms for attenuation and azimuthal roughness significantly improves their model’s fit (red curves in Figure 8).

The PDF and sampling function of the thread BSDF are straightforward to implement, similar to the single-fiber scattering model but with distinct normalized attenuation PDFs.



Figure 9: Energy conservation test.

We conduct an energy conservation test by fitting our thread BSDF model to the simulated distribution of multiple-fiber scattering without absorption. We then render a head of hair geometry using the thread BSDF under a constant environment. This setup emphasizes multiple scattering, making it easier to detect any excessive energy in the model. As shown in Figure 9, our model conserves energy.

2 3D simulation

We also perform 3D simulations of multiple-fiber scattering. Thanks to the Bravais index, we can reduce ray tracing in 3D space to a 2D cross-sectional plane by computing an equivalent index of refraction. Specifically, at each intersection, we sample a new ray direction using the microfacet BSDF, following the method proposed in [4]. We track the longitudinal angle, and when the ray enters a fiber, we compute the equivalent index of refraction (η) corresponding to the current oblique angle. Additionally, we calculate the Fresnel factor and the ray's travel distance in 3D space based on η .

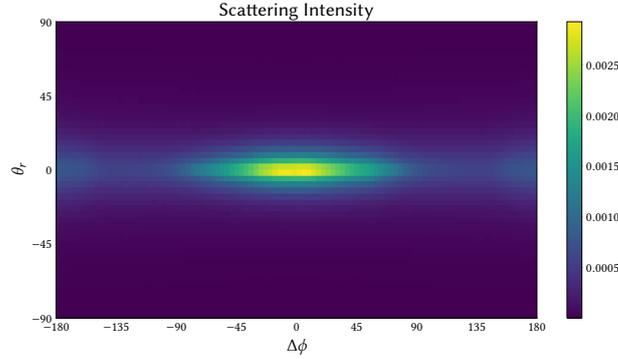


Figure 10: Scattering intensity plot with roughness value 0.1.

3 Simulation and Rendering

Yarn geometry simulation Yarn-level geometry plays a critical role in reproducing the visual appearance of woven fabrics. [5] represented yarn centerlines using parametric curves, while [6] employed a reduced simulation model where yarn deformation was restricted to the height direction. In contrast, our method utilizes a full physically-based simulation, enabling a more accurate representation of yarn geometry and its deformations.

We adopt the simulation method in [7] to generate the geometric structure of yarns in woven fabrics. For each cloth sample, a repeating patch is extracted based on the woven pattern. The initial yarn shapes are defined as parameterized curves derived from this pattern. These initial shapes are then relaxed into their deformed configurations using physically-based simulation. The simulation incorporates stretching and bending energy terms, while yarn contraction is adjusted to account for tension introduced during the weaving process. Periodic boundary conditions are applied to ensure that the patch is tileable seamlessly. Once relaxed, the patch is tiled to construct and render the complete fabric.

Twisted plies We obtain the yarn geometry through physically-based simulation, where each yarn is modeled as an elastic rod. Optionally, we model twisted plies as circularly helical curves around the yarn's center line. For a yarn centered on the Z-axis, each ply follows the form [8]:

$$\mathbf{S}(z) := (S_x(z), S_y(z), z), \quad (6)$$

with

$$\begin{aligned} S_x(z) &= R^{\text{ply}} \cos(2\pi z / \alpha^{\text{ply}} + \theta^{\text{ply}}), \\ S_y(z) &= R^{\text{ply}} \sin(2\pi z / \alpha^{\text{ply}} + \theta^{\text{ply}}). \end{aligned} \quad (7)$$

where R^{ply} is the radius of the helix, α^{ply} controls the twist rate and θ^{ply} provides an initial phase of the twist.

If we want to model twisted plies in thread inversion, we can directly apply the above formulation since the captured threads are straight. This allows us to compute the points through which the twisted plies pass. We then fit B-spline curves through these points, and the fitted B-spline control points are inputted into the rendering process for both forward and inverse rendering of the threads.

To model twisted plies around non-straight yarns from the simulation, we first process the polyline output from the simulation and determine the normal and binormal for each segment. Assume the tangents are \mathbf{t}_i , the normals are \mathbf{n}_i , the binormals are \mathbf{b}_i . Given the normal \mathbf{n}_0 and the binormal \mathbf{b}_0 of the first segment, we can compute the normals and binormals for the rest of the segments using parallel transport [9]. The parallel transport from a unit vector \mathbf{r}_1 to another unit vector \mathbf{r}_2 is the minimum rotation that aligns \mathbf{r}_1 with \mathbf{r}_2 and can be formally computed as

$$\mathbf{P}_{\mathbf{r}_1}^{\mathbf{r}_2} \equiv \mathbf{R}(\mathbf{r}_1 \times \mathbf{r}_2 / |\mathbf{r}_1 \times \mathbf{r}_2|, \angle(\mathbf{r}_1, \mathbf{r}_2)), \quad (8)$$

where $\mathbf{R}(\mathbf{r}, \psi)$ is the rotation about the unit vector \mathbf{r} by an angle ψ , and $\angle(\mathbf{r}_1, \mathbf{r}_2)$ is the angle between \mathbf{r}_1 and \mathbf{r}_2 . From segment i and $i + 1$, we use parallel transport to compute the rotation matrix $\mathbf{P}_{\mathbf{t}_i}^{\mathbf{t}_{i+1}}$ using the tangent vectors \mathbf{t}_i and \mathbf{t}_{i+1} , and rotate \mathbf{n}_i and \mathbf{b}_i to get \mathbf{n}_{i+1} and \mathbf{b}_{i+1} respectively. Assume that the yarn center line is $\mathbf{Y}(z)$ written as a function of the length parameter z . Then we can write the twisted plies as

$$\mathbf{C}(z) = \mathbf{Y}(z) + S_x(z)\mathbf{n}(z) + S_y(z)\mathbf{b}(z). \quad (9)$$

Two-scale rendering Rendering human-scale cloth is challenging, and the brute-force method of tiling the entire fabric is impractical. We leverage the two-scale rendering method proposed in [5] to efficiently render the fabric. Specifically, we implement a new integrator in Mitsuba 3 [10]. We assign a Boolean parameter to each object to denote whether it should be rendered using the two-scale rendering method and treat fabric and non-fabric objects separately. We assume the woven fabric is tiled from a patch with 6 warp yarns and 6 weft yarns, and only one copy needs to be stored, which requires just 6.6KB of memory. The patches are instanced, allowing us to perform ray intersections over an infinite domain. The two-scale rendering approach considers a macro scale for the fabric surface, which is represented using a triangle mesh in the scene, and a micro scale for the yarns. We first generate rays from the camera and compute their intersection with the macro surface. Then, we sample a light direction and transform the intersection point and the directions into the local patch space, computing the intersection in that space. Next, we start path tracing at the micro-scale. We sample the thread BSDF and simulate ray interaction with the patch, performing next-event estimation until the ray exits the micro surface. We then transform the ray back into world space and continue the macro-path tracing. This method allows us to achieve fast and accurate rendering of woven fabric. In Figure 11, we validate our implementation by comparing the rendering of a yarn plane by tracing the ground truth geometry and using the two-scale rendering with a constant environment light.

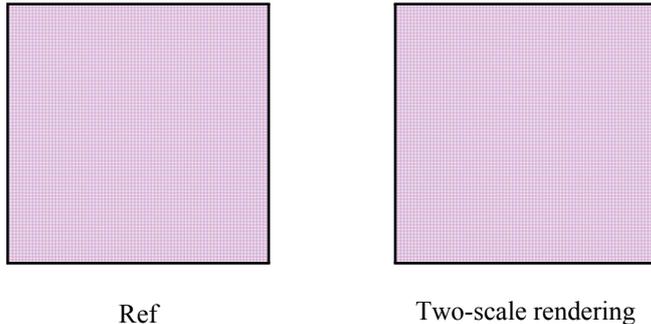


Figure 11: We validate two-scale rendering by rendering a yarn plane under constant lighting.

4 Fabric Appearance Prediction

As explained in the main paper, to validate that our method accurately predicts the appearance of the cloth, we acquire 16 fabric samples along with the corresponding threads used to make these samples. We render the fabric wrapped around a cylinder and compare it with the captured photos. We observe a highlight band with a yellow hue in the middle. As discussed in [11], the width of the highlight depends on the slope of the underlying yarns, and we achieve the desired width of the highlight by adjusting the amplitude variation as the yarn extends beyond the 2D weaving plane. In the following figure, we examine the influence of longitudinal roughness and yarn amplitude on the highlight.

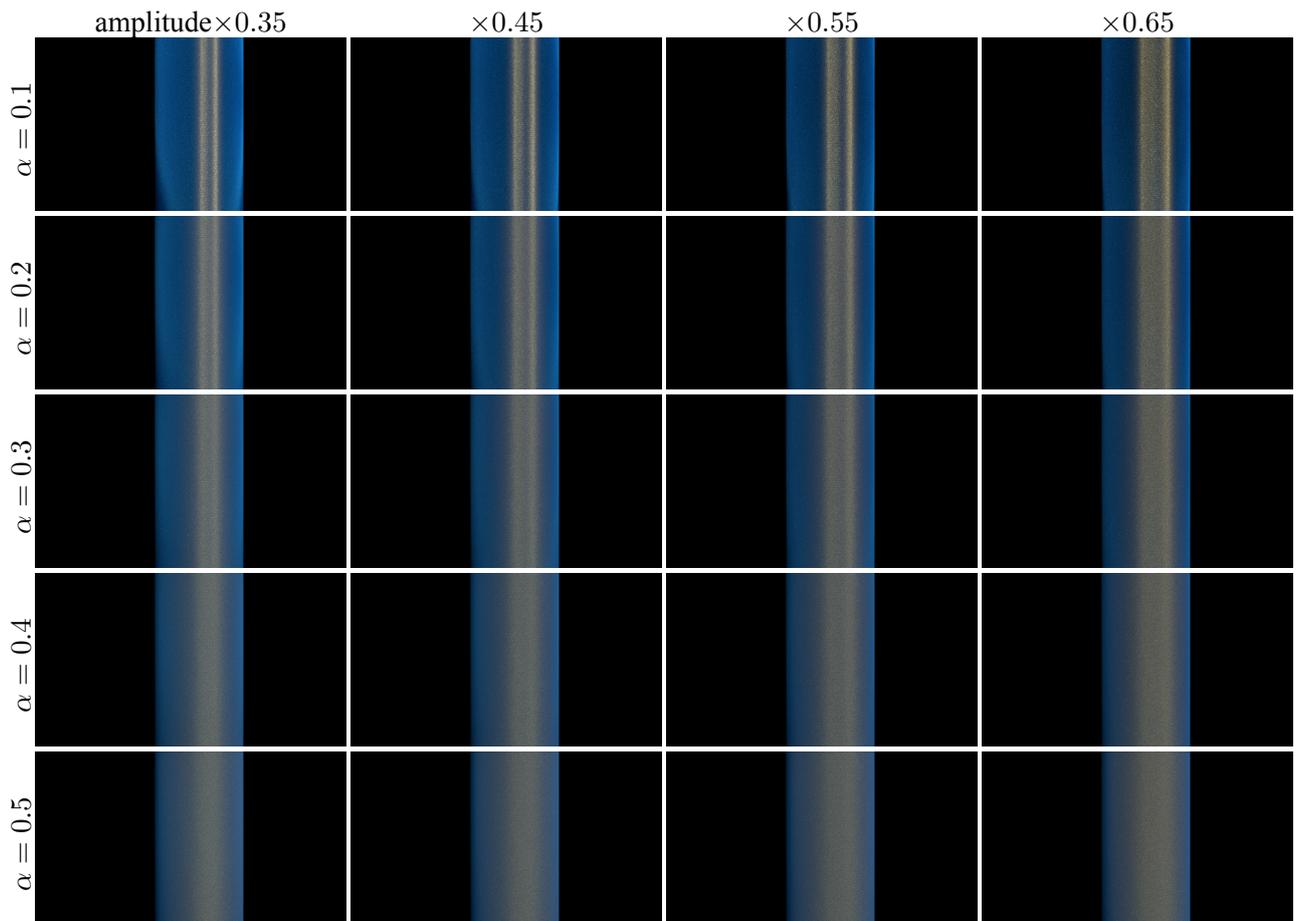


Figure 12: We render cloth wrapped around a cylinder with different longitudinal roughnesses (α) and scale the yarn amplitude using various scale factors. We find that longitudinal roughness primarily controls the sharpness of the highlight's boundary, while yarn amplitude controls the width of the highlight—a larger amplitude corresponds to a wider highlight.

In addition to the 16 silk fabric samples, we also obtain two rayon fabric samples, manually woven into a plain weave. One sample uses pink rayon threads for both the warp and weft yarns, while the other replaces the warp yarns with blue rayon threads. We also acquire the corresponding pink and blue rayon threads used to make these samples. Our pipeline accurately predicts the appearance of the fabric wrapped around the cylinder. Unlike the silk samples, these rayon fabrics exhibit no obvious highlight band, likely due to the greater height variations in the threads.



Figure 13: We compare our prediction of the fabric on the right to the captured photos on the left, demonstrating a matching appearance.

References

- [1] Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. A practical and controllable hair and fur model for production path tracing. In *ACM SIGGRAPH 2015 Talks*, pages 1–1. 2015.
- [2] Eugene d’Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. An energy-conserving hair reflectance model. In *Computer Graphics Forum*, volume 30, pages 1181–1187. Wiley Online Library, 2011.
- [3] Junqi Zhu, Zahra Montazeri, J Aubry, Lingqi Yan, and Andrea Weidlich. A practical and hierarchical yarn-based shading model for cloth. In *Computer Graphics Forum*, volume 42, page e14894. Wiley Online Library, 2023.
- [4] Jonathan Dupuy and Anis Benyoub. Sampling visible ggx normals with spherical caps. In *Computer Graphics Forum*, volume 42, page e14867. Wiley Online Library, 2023.
- [5] Zixuan Li, Pengfei Shen, Hanxiao Sun, Zibo Zhang, Yu Guo, Ligang Liu, Ling-Qi Yan, Steve Marschner, Milos Hasan, and Beibei Wang. Fiber-level woven fabric capture from a single photo, 2024.
- [6] Zahra Montazeri, Søren B Gammelmark, Shuang Zhao, and Henrik Wann Jensen. A practical ply-based appearance model of woven fabrics. *ACM Transactions on Graphics (TOG)*, 39(6):1–13, 2020.
- [7] Jonathan Leaf, Rundong Wu, Eston Schweickart, Doug L. James, and Steve Marschner. Interactive design of yarn-level cloth patterns. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2018)*, 37(6), 11 2018.

- [8] Shuang Zhao, Fujun Luan, and Kavita Bala. Fitting procedural yarn models for realistic cloth rendering. *ACM Transactions on Graphics (TOG)*, 35(4):1–11, 2016.
- [9] Miklós Bergou, Basile Audoly, Etienne Vouga, Max Wardetzky, and Eitan Grinspun. Discrete viscous threads. *ACM Transactions on graphics (TOG)*, 29(4):1–10, 2010.
- [10] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>.
- [11] Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. A practical microcylinder appearance model for cloth rendering. *ACM Transactions on Graphics (TOG)*, 32(2):1–12, 2013.