

A Practical and Hierarchical Yarn-based Shading Model for Cloth

J. Zhu¹, Z. Montazeri², J. Aubry³, L. Yan¹ and A. Weidlich⁴

¹University of California, Santa Barbara

²University of Manchester

³Wētā Digital / Unity

⁴NVIDIA




Figure 1: In this scene, we compare our hierarchical yarn-based model to a geometric reference on a knitted glove (left). Our method achieves a high level of accuracy in reproducing the reference appearance in both near- and far-field. The model can be decomposed into four sub-components, namely R and D , simulating single scattering, and F^{yarn} and B^{yarn} which are responsible for the physically-based simulation of the multiple scattering. Our method significantly outperforms the reference in terms of computational efficiency. For the same quality, our method is $7\times$ faster than the reference, while only requiring 3% of the memory.

Abstract

Realistic cloth rendering is a longstanding challenge in computer graphics due to the intricate geometry and hierarchical structure of cloth: Fibers form plies which in turn are combined into yarns which then are woven or knitted into fabrics. Previous fiber-based models have achieved high-quality close-up rendering, but they suffer from high computational cost, which limits their practicality. In this paper, we propose a novel hierarchical model that analytically aggregates light simulation on the fiber level by building on dual-scattering theory. Based on this, we can perform an efficient simulation of ply and yarn shading. Compared to previous methods, our approach is faster and uses less memory while preserving a similar accuracy. We demonstrate both through comparison with existing fiber-based shading models. Our yarn shading model can be applied to curves or surfaces, making it highly versatile for cloth shading. This duality paired with its simplicity and flexibility makes the model particularly useful for film and games production.

CCS Concepts

• *Computing methodologies* → *Reflectance modeling*;

1. Introduction

Fabrics are one of the most common materials that appear in virtual scenes yet one of the most challenging topics in the field of computer graphics for many years. This is due to its complex structure – a piece of cloth is a complex aggregation of fibers, twisted into plies, then twisted into yarns, then woven or knitted together. Each level of aggregation forms a distinct type of geometric structure, therefore resulting in a unique kind of appearance for different types of cloth.

Different methods have been proposed for reproducing the appearance of cloth. These might be roughly categorized as curve-based and surface-based shading models. Curve-based methods explicitly model individual fibers/plies/yarns as curves and model their scattering using the Bidirectional Curve Scattering Distribution Function (BCSDF) similar to hair rendering [ZJMB11; SKZ11]. While they represent cloth with high precision, they are slow to render, require heavy storage, are difficult to edit, and cannot be efficiently filtered in render time. On the other hand, surface-based methods model cloth as an explicit surface and encode its appearance mostly as reflectance models or scattering at a given point using the Bidirectional Scattering Distribution Function (BSDF) [IM12; SBDJ13]. These methods are significantly more efficient to represent and render, but they result in a flat and opaque appearance, fail at representing high-frequency details such as highlights from individual fibers, and usually lack physical accuracy. Moreover, efficient surface-based models rely on assumptions that are specific to woven cloth, thus handling knitted cloth remains to be challenged.

In this work, we propose a yarn-based cloth shading model that utilize the hierarchical structure of yarns to simulate light scattering more efficiently and practically. Specifically, we demonstrate a yarn-level shading model with ply-level geometry details and fiber-level highlight. Note that the term "fiber-level highlight" in this context refers to the accurate macro-scale highlight shape, as if the actual fibers with correct tangent directions have been generated, rather than the detailed highlight on each individual fiber. Our proposed shading model can uniformly render either surface-based geometries with implicitly defined yarn curves given weave or knitted patterns, or similarly, curve-based input models with explicit yarn curves to be passed as input geometry. We model the shading of a single fiber using a BCSDF with three lobes (R , TT and D). Then we derive the shading model of yarns from the property of a single ply composed of fibers and their geometry statistics such as the number of fibers in a ply, the amount of twist, and similarly the number of plies in a yarn and their twist amount. The proposed shading model simulates single and multiple scattering between individual fibers and encapsulates them to formulate an aggregated model for the ply and further into the yarn representation. As illustrated in Fig. 1, our model successfully captures ply geometries on-the-fly given the yarn geometries, and the fiber-level details are modeled statistically to match the reference highlights. Our hierarchical technique provides an accurate yet practical approach and addresses the best of both worlds.

The core of our method involves two key components: an implicit solution for the ray intersection and a yarn cylinder, and an analytical multi-level dual scattering shading model. Geometri-

cally, we approach the problem of solving the intersection of a ray with a yarn cylinder in order to obtain implicit information about the alignment of ply geometry within the yarn cross-section. To model the appearance, we approximate the yarn shading model by incorporating the number of fibers that the ray passes through and aggregates light scattering through the bundle. This allows us to analytically derive a complete shading model that incorporates both single and multiple scattering effects.

Concretely, our main contributions include:

- A new hierarchical geometric model to compute the procedural ply geometries and the fiber details on-the-fly to preserve high fidelity with minimum overhead on memory and performance, as elaborated in Sec. 3.
- A new analytical single and multiple light scattering model for yarns as an aggregated formulation of individual fibers using dual-scattering theory. Our proposed method is explained in Sec. 4.

To demonstrate the practical usefulness of our technique, we have integrated it into a physically-based production renderer.

2. Related Work

Render cloth as surface. Traditionally, cloth has been rendered using surface-based models defined as 2D thin sheets utilizing lightweight Bidirectional Reflection Distribution Function (BRDF) [AMF03; IM12; SBDJ13; JWH*22]. While these models are popular because of their fast performance, they lack geometric sub-yarn details and do not offer realism in close-up views. Furthermore, they approximate multiple scattering using a simple diffuse lobe yield inaccurate results. Besides, data-driven approaches such as Bidirectional Texture Functions (BTF) [DvGNK99] offer a more faithful appearance, but they are data intensive and not efficient for practical usage and editability.

Render cloth as curves. More recently, as an alternative to surface-based approaches, a family of appearance models are introduced describing the cloth as explicit fibers (e.g. [KSZ*15]), plies ([MGZJ20]), or yarns (e.g. [YKJM12]) providing high-fidelity renderings in close-up views. Despite their rich detailed appearance, they are inefficient for rendering or editing due to the high complexity of the models. To address this issue, precomputation-based methods [ZHRB13; KWN*17; MXF*21] were proposed and improved the performance while preserving the realism. However, better efficiency remains to be challenged due to their intensive computation. While GPU-based methods [WY17] address the performance issue, they lose physical accuracy due to the estimation of inner fibers.

Render cloth as volume. Alternatively, the micro-appearance models that define the cloth down to their fiber details can be rendered using their volumetric representation following the original study by Kajiya and Kay [KK89] and further improved later [SKZ11; ZJMB11; LN18]. However, they share the same disadvantages with fiber-based models due to their cumbersome computation and challenging manipulation.

Render using aggregation. A series of aggregation-based methods were introduced to the community utilizing the generalization

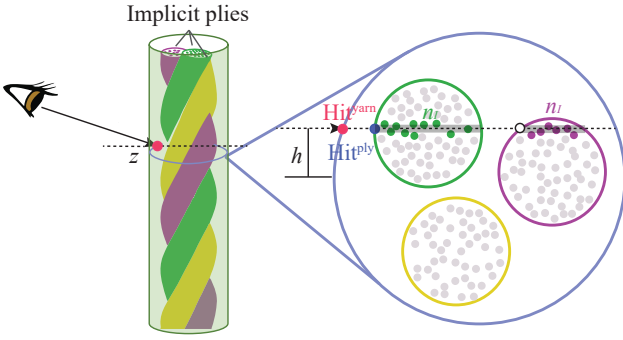


Figure 2: Geometric representation and light intersections

Table 1: Summarization of the notations used for geometry.

Notation	Definition
h	azimuthal offset of the yarn cylinder
z	longitudinal offset of the yarn cylinder
n	number of fibers in a ply
θ_{fiber}	fiber twist angle for ply
d	fiber density in a ply
r_f	radius of fiber
r_p	radius of ply
n_i	number of intersected fibers
n_{plies}	number of plies in a yarn

of sub-components of the light scattering model to enable more efficient rendering. Dual scattering [Yuk08] aggregate the complex light scattering of human hair into simple lobes. But the number of fibers in a yarn is notably less than in human hair and the light scattering of fibers is not as specular as in human hair. Thus, direct usage of a dual scattering model would not be applicable to cloth rendering.

Montazeri et al. [MGZJ20] proposed an aggregated technique that encapsulates the light scatterings of individual fibers to approximate the overall appearance in ply level for woven fabrics, similarly for knitted cloth [MGJZ21]. Their model is observation driven and uses texture mapping while ours is an analytical solution. Also, our aggregated yarn-based model outperforms their ply-based representation and further accelerates the render process. Likewise, a bundle of fur fibers can be also aggregated together for a more efficient rendering ([ZZW*22], [YSJR17]). In the same fashion, our proposed method aggregates the single and multiple scattering of light but unlike this line of work, we can immediately obtain the matching results to the reference using the same sets of parameters with no need for parameter tweaking [MGZJ20] or cumbersome neural networks [ZZW*22].

3. Our Geometric Model

3.1. Overview

The structure of cloth is complex at many scales. A garment is composed of numerous yarns, via manufacturing techniques such as

weaving and knitting. Most of the studies involving rendering cloth and textiles have been aimed toward woven fabrics due to their simpler structure. While woven fabrics consist of two perpendicular yarns called warps and wefts, knits have usually made of one long yarn for the entire piece of the fabric. Many real-world yarns contain multiple sub-strands, or *plies*, each of which consists of hundreds of *fibers* twisted around a common center that can be procedurally generated [SZK15] and fit with real samples [ZJMB11].

To describe the geometry of the cloth, we need the key properties for each level of the hierarchical structure. To define the fibers, their quantity and distribution in a bundle of fiber, twisting amount, and irregularities (known as migration) need to be specified. Similarly, at the ply level, their quantity and the distance from the center are needed. Lastly, the alignment of the yarns forming repeating woven or knitted patterns identifies the geometry of the fabrics and can be passed for the rendering stage in a few ways.

3.2. Our Implicit Ray-tracing

We only need to explicitly define the yarn curves, then we introduce an implicit technique to get the intersection information of the plies at render time and finally, the fiber details are added statistically.

First, the ray intersection with the surface of the yarn is computed, then we further intersect the ray to compute the ply geometry. We assume the ray incident is perpendicular to the ply axis and the ply intersections are computed within a circular cross-section of the yarn, as illustrated in Fig. 2.

Let's define h as the azimuthal offset similar to the hair literature, and z as the position along the yarn from one end to the other. The plies are twisted around the yarn centerline as z increases given the rotation speed ϕ_{ply} . The total rotation of the cross-section at z with an initial phase ϕ_{init} is computed as follows:

$$\phi_z = \phi_{init} + z\phi_{ply} + \phi_{view}. \quad (1)$$

where ϕ_{view} is the updated angle from where the camera aligns with ϕ_{init} to make the yarn rotation to be view-independent. When the camera moves, the updated phase ϕ_{view} can correctly represent the different ply geometry with changes in viewing angle which can ensure temporal coherence.

Once we have obtained ϕ_z and h , we can know all the intersections against all plies, which are then used for shading (as shown in Fig. 2). If there is no intersection, we report transparency. In essence, if the ray passes through the cross-section without hitting the plies, it is marked as transparent to preserve the correct silhouette of ply geometry. This assumes the plies are not further away from the center which causes gaps.

For the intersection, we also calculate the corresponding fiber details by considering their normals. To do so, we rotate the ply normals along their axis to follow the fiber frames. Later, we build the azimuthal-longitudinal frame for the BCSDf during rendering.

Next, we calculate the number of fibers an incident ray passes through using a 2D cross-sectional ray intersection assuming the ray incident is perpendicular to the ply axis. As shown in Fig. 2, the distribution of fibers within the ply circular cross-section defines the density d . The simple ratio of the occupied area by the

Table 2: Summarization of the notations used for shading.

Notation	Definition
R	reflection lobe
TT	transmission lobe
D	diffuse lobe defined
F^{ply}	forward scattering lobe for ply
B^{ply}	backward scattering lobe for ply
F^{yarn}	forward scattering lobe for yarn
B^{yarn}	backward scattering lobe for yarn

fibers over the entire cross-section area is a mathematical approximation and cannot represent the geometric details at the fiber level. Once the fibers are twisted along the ply centerline, their cross-sections are stretched from circular to elliptical, so we introduce the fiber twist θ_{fiber} , together with a user-specified constant c to compensate for the enlarged elliptical area and compute the overall unitless density

$$d := \frac{cn\pi r_f^2}{\pi r_p^2 \cos \theta_{fiber}}, \quad (2)$$

where r_f and r_p are radii of fibers and the ply, respectively.

With the fiber density parameter in place, we can compute the number of fibers an incident ray passes through using a 2D cross-sectional ray intersection assuming the ray incident is perpendicular to the ply axis. Given the length of the intersected ray inside the ply cross-section ($l \leq 2$, because $l = 2h$ and $h \leq 1$). The number of intersected fibers n_l can be computed as $dl\sqrt{n/\pi}$, where $\sqrt{n/\pi}$ is an estimate of the number of fibers intersected per unit distance.

By using of our implicit tracing method, we can obtain intersection information between rays and yarns. This includes determining if the intersection point is transparent, as well as implicitly obtaining the normal (ply level), tangent (fiber level), the intersected plies information (ray-intersected order, the number of fibers). The geometric representation is used in the next section to be shaded on-the-fly with no precomputation needed and computed procedurally.

4. Our Aggregated Shading Model

Our shading technique leverages the dual scattering model [Yuk08], hence we first provide an overview in Sec. 4.1. Further, we describe how light is scattered by the defined geometry explained in Sec. 3 and present our novel yarn shading model (Sec. 4.4) and building upon fiber single scattering and aggregated formulation in the ply level as elaborated in Sec. 4.2 and Sec. 4.3, respectively. In order to facilitate the understanding of the effects of our method, in Fig. 5 and Fig. 1, we show the rendering of different shading components.

4.1. Overview on Dual Scattering

Dual scattering (DS) [Yuk08] approximates global illumination effects within the hair volume at a specific shading point as a combination of two components: global and local scattering. The former

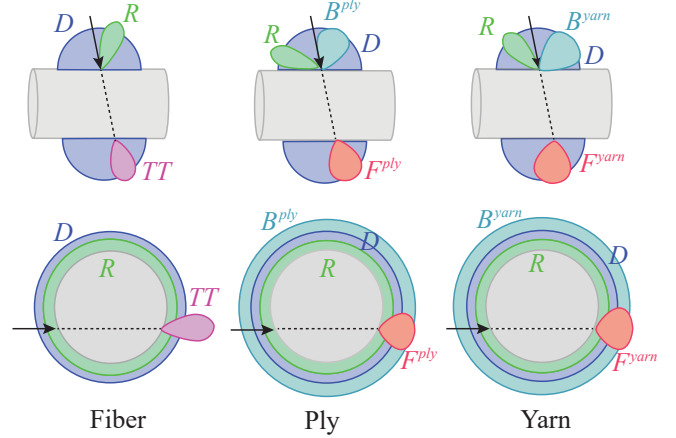


Figure 3: Illustration of our single and multiple scattering model per fiber, ply, and yarn for longitudinal (top) and azimuthal (bottom) components. Note that the R , TT and D lobes are defined under the fiber frames, and B^{ply} , F^{ply} , B^{yarn} and F^{yarn} lobes are defined under the ply frames.

approximates how much light arrives at x after passing through n_l fibers along the light path and the latter captures the portion of the light that scatters inside the volume, together forming forward and backward lobes. The global transmission in dual scattering (DS) is considered as the local shading in our model as an approximation for the attenuation.

Averaged forward/backward lobes: following the same notation in [Yuk08] as illustrated in Fig. 2, given a specific longitudinal incident angle with unit radiance, the total radiance to the front and back hemisphere can be computed. The back hemisphere refers to $\phi_r - \phi_i \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, and the rest is captured by the front hemisphere. To compute the averaged attenuation of light, we simply enumerate all ϕ_i and sum the corresponding outgoing radiance toward the front or back hemisphere, and compute the average using a 3D integral as follows:

$$a_F(\theta_i) = \frac{1}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \int_{\Omega_f} f(\omega_i, \omega_r) \cos \theta_r \, d\omega_r \, d\phi_i. \quad (3)$$

Similarly, the backward attenuation is:

$$a_B(\theta_i) = \frac{1}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \int_{\Omega_b} f(\omega_i, \omega_r) \cos \theta_r \, d\omega_r \, d\phi_i, \quad (4)$$

where $f(\omega_i, \omega_r)$ is the single fiber's scattering model described in §4.2.

4.2. Fiber Shading Model

As first introduced by Marschner et al [MJC*03], fibers are glass-like cylinders and fiber scattering models generalize the light interactions with a single fiber into reflection and transmission lobes (R and TT respectively) [KSZ*15]. To approximate the multiple scatterings, the fiber-level shading model considers an additional diffuse lobe (D), as suggested in [YTJR15]. The further introduced D




Figure 4: Contribution of the four lobes of our model to the final result (reflection, diffuse, forward scattering and backward scattering), rendered using an environment light.

lobe enhances the realism and flexibility of the fiber shading model, enabling it to represent a wider range of cloth. The reason is that fibers are not purely glass-like tubes with no internal scatterings as traditionally assumed and a diffuse component addresses this limitation similar to hair or fur fibers [YTJR15].

To formulate the fiber shading, the incident direction ω_i and outgoing direction ω_r are first parameterized into θ_i and θ_r in the longitudinal plane, and ϕ_i and ϕ_r in the azimuthal plane. Shading on these two planes is computed separately as M and N , and then multiplied together along with the attenuation a . The fiber BCSDF is:

$$f(\theta_i, \theta_r, \phi_i, \phi_r) = \sum_{p \in \{R, TT, D\}} a_p \cdot M_p(\theta_h) \cdot N_p(\phi) / \cos^2 \theta_i, \quad (5)$$

where p represents different types of lobes and the cosine term follows the conventional hair model [MJC*03]. θ_h is the half vector $\theta_h = (\theta_i + \theta_r)/2$ and ϕ is the difference of ϕ_i and ϕ_r and can be written as $\phi = (\phi_i - \phi_r)$.

To simplify the fiber model, we use a classic Gaussian distributions, following the literature [MJC*03; KSZ*15; YTJR15], to approximate M_R , M_{TT} and N_{TT} , while other lobes are modeled as uniform distributions, as follows:

$$\begin{aligned} M_R(\theta_h) &= g(\beta_R^M; \theta_h), & N_R(\phi) &= \frac{1}{2\pi} \\ M_{TT}(\theta_h) &= g(\beta_{TT}^M; \theta_h), & N_{TT}(\phi) &= g(\beta^N; \phi) \\ M_D &= \frac{1}{\pi}, & N_D &= \frac{1}{2\pi}, \end{aligned} \quad (6)$$

where β^M and β^N are the longitudinal and azimuthal roughness, respectively. In our model, we define $g(\beta, \theta) = \frac{1}{\sqrt{2\pi} \cdot \beta} \cdot e^{-\theta^2 / (2\beta^2)}$. Here we use the traditional Gaussian instead of the spherical Gaussian because β is small enough, that g is negligible outside of $[-\pi/2, \pi/2]$.

4.3. Ply Shading Model

We present an aggregated model to approximate light scattering within a ply constructed of n fibers using the BCSDF from Sec. §4.2. Given the twist angle θ_{fiber} , the fiber tangents are computed by deviating the tangent direction of the plies around the axis to accurately model the fiber-level highlights. As already mentioned, the fiber geometries are not explicitly computed and their contribution is only statistically considered. The interactions between these virtual fibers within the ply cross-section can be modeled using a multiple-scattering component in addition to the reflection and transmission lobes, R and TT , respectively. This is illustrated in Fig. 3.

To simplify multiple scattering, we assume the scattering events happen only along the main path, which is the segment intersected by the incident ray and the ply cross-section. The scattered energy leaves the ply either by passing forward through the yarn, or backward toward the incident direction. Therefore, the multiple scattering component is captured using two lobes: the forward scattered lobe (F^{ply}) and the backward scattered lobe (B^{ply}). The weight to split the energy among them is dependent on the ply physical properties, mainly the fiber density.

Enhanced dual scattering. To approximate the multiple scattering components (F^{ply} and B^{ply}), we utilize the classical dual scattering method [Yuk08] with two main modifications to enhance the performance specifically for fabrics. Firstly, the conventional DS theory is built upon only R and TT lobes to approximate multiple scattering. Both these components exhibit high level of specularity. Consequently, regardless of the number of scattering events, the resulting F^{ply} or B^{ply} lobe remains relatively specular. However, in the case of cloth DS, an additional diffuse component (D) is considered to better capture of the scatterings with more controllability. Secondly, the traditional DS model presumes an infinite number of similar hairs while the quantity of the fibers in yarn bundles is finite and predefined.

Formulating the scattering components using dual-scattering introduces two main properties to be computed, following the original theory: attenuation a , and distributions M and N referring to longitudinal and azimuthal, respectively. Accordingly, the forward and backward scattering lobes can be written as:

$$\begin{aligned} F^{ply} &= a_F^{ply} \cdot M_F^{ply} \cdot N_F^{ply}, \\ B^{ply} &= a_B^{ply} \cdot M_B^{ply} \cdot N_B^{ply}, \end{aligned} \quad (7)$$

where attenuation A and distributions M and N are defined in what follows.

Attenuation. The aggregated attenuation of a ply that consists of several fibers, is the product of the attenuation per intersected fibers along the scattering path. To compute the forward component, the light can have an even number of backward bounces (B) multiple times with the rest being forward bounces (F). For instance, FBBFF is a forward light path for $n_f = 5$. We observed, in the case of forward attenuation, the energy is mostly derived from fully forward light paths (F^+) and the pair of backward bounces are negligible.

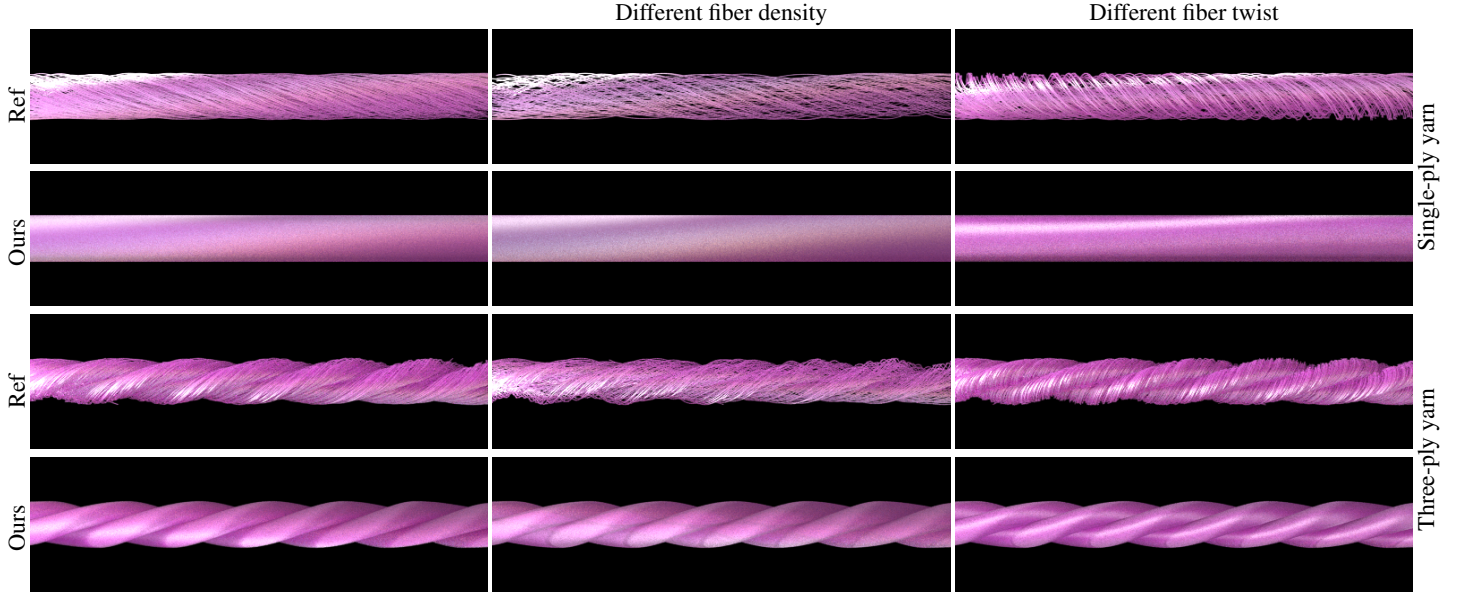


Figure 5: Single-yarn comparisons with the reference (fiber-based path-tracer). Changing the fiber properties (i.e. twist and density) yield to different highlight shape and overall color and model can successfully match the reference. Please see the supplementary video for more parameter tweaking.

Therefore, using Eqn. 3 for every intersected fiber, a_F^{ply} is:

$$a_F^{\text{ply}}(\theta_d) = \prod_{k=1}^{n_l} a_F(\theta_d). \quad (8)$$

Similarly, to compute the total backward attenuation of the ply, an odd number of backward bounces per fiber can be considered while the rest of the energy comes from the forward bounces. Following traditional DS, we also assume that only single and triple backward bounces mainly contribute to the total backward attenuation.

To compute the average backward scattering with a single B, all the possible scattering paths including one B is F^+BF^+ —where + notation means at least one occurrence. Consequently, the possible backward scattering paths with three B bounces can be written as $F^+BF^+BF^+BF^+$. The five and more bounces contributions are negligible, as elaborated in original DS [Yuk08]. Similar to their derivations, the average attenuation for single and triple backward scattering are as follows:

$$\begin{aligned} a_{1B}^{\text{ply}} &= a_B(\theta_d) \sum_{i=0}^{n_l-1} a_F^{2i}(\theta_d), \\ a_{3B}^{\text{ply}} &= a_B^3(\theta_d) \sum_{i=1}^{n_l-1} \sum_{j=0}^{i-1} \sum_{k=j+1}^{n_l-1} a_F^{2(i-j-1+k)}(\theta_d). \end{aligned} \quad (9)$$

Therefore, the total attenuation for backward scattering is:

$$a_B^{\text{ply}} = a_{1B}^{\text{ply}} + a_{3B}^{\text{ply}}. \quad (10)$$

Distributions. To compute forward scattering distributions for the ply, we model Gaussian functions with wider variance (v) compared to individual fibers by accumulating all variances along the possible light paths F^+ . β_F is the weighted average of roughness

for fiber shading lobes (R , TT , D) using their corresponding attenuation, the ply variance for longitudinal and azimuthal are as follows, respectively, following by their distribution:

$$\begin{aligned} v_F^{\text{M,ply}} &= \sum_{k=1}^{n_l} \beta_F^{M,k}(\theta_d), & M_F^{\text{ply}}(\theta_h) &= g(v_F^{\text{M,ply}}, \theta_h), \\ v_F^{\text{N,ply}} &= \sum_{k=1}^{n_l} \beta_F^{\text{N}}(\theta_h), & N_F^{\text{ply}}(\phi) &= g(v_F^{\text{M,ply}}, \phi). \end{aligned} \quad (11)$$

Similarly, the backward scattering distributions for longitudinal and azimuthal components are:

$$\begin{aligned} M_B^{\text{ply}} &= g(v_B^{\text{M,ply}}, \theta_h), \\ N_B^{\text{ply}} &= \frac{1}{2\pi}, \end{aligned} \quad (12)$$

where $v_B^{\text{M,ply}}$ is the ply variance for backward scatterings that is computed using the weighted averages of variance from all possible light paths similar to the formulation of aggregated attenuation in 10. Please see the Appendix for the derivation.

4.4. Yarn Shading Model

So far, we explained how to aggregate the ply shading model using the fiber scattering model. In this section, we follow the same approach by aggregating the scattering components of the plies (F^{ply} and B^{ply}) further to compute the yarn scattering lobes (F^{yarn} , B^{yarn}). The F^{yarn} and B^{yarn} are also BCSDFs similar to Eqn. 7 and can be

defined as:

$$\begin{aligned} F^{\text{yarn}} &= a_F^{\text{yarn}} \cdot M_F^{\text{yarn}} \cdot N_F^{\text{yarn}} \\ B^{\text{yarn}} &= a_B^{\text{yarn}} \cdot M_B^{\text{yarn}} \cdot N_B^{\text{yarn}} \end{aligned} \quad (13)$$

In a similar fashion to §4.3, we compute the attenuation and the distribution of the lobes, this time for the yarn. Contrary to the ply where we were only used a statistical model, we now examine all plies explicitly and compute their aggregation. This has two reasons. First, the number of plies in a yarn is significantly smaller (between 3 to 12), and adding them explicitly is not costly. Second, the plies may offer distinct properties. Hence, we accumulate the contribution of each of the plies explicitly, instead of numerically averaging them, such as in the fiber scenario—that assumes fibers are identical.

Attenuation. The yarn forward attenuation, follows Eqn. 8 and can be computed as the product of the underlying ply attenuation. The attenuation for backward scattering, however, is not as straightforward. Given the arrangement of the plies in a yarn, a forward-directed light can intersect the plies at most twice before it leaves the yarn. Hence, only two light paths can occur in this scenario (B and FBF). The aggregated ply attenuations are as follows:

$$\begin{aligned} a_F^{\text{yarn}}(x, \omega_d) &= \prod_{k=1}^{n_{\text{plies}}} a_F^{\text{ply},k}(\theta_d) \\ a_B^{\text{yarn}}(x, \omega_d) &= \sum_{k=1}^{n_{\text{plies}}} a_B^{\text{ply},k}(\theta_d) \left(1 + a_F^{\text{ply},k}(\theta_d)^2\right) \end{aligned} \quad (14)$$

Distribution. In a similar approach to aggregated ply model, the distribution functions for forward scattering can be assumed to be Gaussian. The yarn variance can be computed by accumulating the individual ply variance to form a wider and more scattered lobe. The forward variance for longitudinal and azimuthal components is computed here, following their distributions.

$$\begin{aligned} v_F^{M,\text{yarn}} &= \sum_{k=1}^{n_{\text{plies}}} v_B^{M,\text{ply},k}(\theta_d) \cdot a_F^{\text{ply},k}(\theta_d), & M_F^{\text{yarn}} &= g(v_F^{M,\text{yarn}}, \theta_h), \\ v_F^{N,\text{yarn}} &= \sum_{k=1}^{n_{\text{plies}}} v_F^{N,\text{ply},k}(\theta_d) \cdot a_F^{\text{ply},k}(\theta_d), & N_F^{\text{yarn}} &= g(v_F^{N,\text{yarn}}, \phi). \end{aligned} \quad (15)$$

Consequently, the backward scattering component is approximated as follows, similar to Eqn. 12:

$$\begin{aligned} N_B^{\text{yarn}} &= g(v_B^{N,\text{yarn}}, \phi) \\ N_B^{\text{yarn}} &= \frac{1}{2\pi}, \end{aligned} \quad (16)$$

where $v_B^{N,\text{yarn}}$ is the yarn-level aggregated variance that is derived using the ply-level variances.

$$\begin{aligned} M_F^{\text{yarn}} &= g(v_F^{M,\text{yarn}}, \theta_h) \\ v_F^{M,\text{yarn}} &= \sum_{k=1}^{n_{\text{plies}}} v_B^{M,\text{ply}}(\theta_d)^k \cdot a_F^{\text{ply}}(\theta_d)^k \\ N_F^{\text{ply}} &= g(v_F^{N,\text{ply}}, \phi) \\ v_F^{N,\text{yarn}} &= \sum_{k=1}^{n_{\text{plies}}} v_F^{M,\text{ply}}(\theta_d)^k \\ M_B^{\text{ply}} &= g(v_B^{M,\text{ply}}, \theta_h) \\ v_B^{M,\text{yarn}} &= \sum_{k=1}^{n_{\text{plies}}} \left(\prod_{i=1}^k a_F^{\text{ply},i} \right)^2 \cdot v_B^{M,\text{ply}}(\theta_d)^k \\ N_B^{\text{yarn}} &= \frac{1}{2\pi} \end{aligned} \quad (17)$$

5. Rendering

Our aggregated shading model explained in §4 is integrated into both Mitsuba [Jac20] and a physically-based production renderer. The results shown in this paper are rendered using the Mitsuba implementation that supports on-the-fly geometry generation to efficiently create our geometry model explained in §3. During render time, we use a custom shader to both evaluate our shading model as described in §4. In §5.1, we provide more details on our proposed importance sampling technique and lastly, we discuss the performance of our technique in comparison with previous works in §5.2.

5.1. Importance Sampling.

Our importance sampling is similar to [dMH13] and can be summarized in the following three stages:

- *Sample lobe p .* The lobes are weighted according to their corresponding energy. Since the longitudinal term M_p and the azimuthal distribution term D_p are normalized, the energy that lobe p carries depends directly on its attenuation term a_p . The probability for sampling lobe p is determined by a_p over the total amount of attenuation that is summed over all lobes.
- *Sample outgoing direction.* Similar to previous works, we explicitly sample the outgoing direction from the selected lobe p in longitudinal and azimuthal following their distributions. We sample from Gaussian or uniform cosine sampling function based on the distribution.
- *Calculate Probability Density Function (PDF).* The PDF is the product of the selected lobe p 's azimuthal and longitudinal PDF followed by a conversion from (θ, ϕ) space to solid angle space. The final sampling weight is the selected p 's BCSDf value divided by its chosen probability.

5.2. Performance

In what follows, we analyze the performance of our model, especially in comparison with previous works. In Table 3, we listed the rendering time, memory consumption and bounce count of knit and woven samples shown in Fig. 6. As the number of fibers in a yarn increases, the performance and memory requirements of our model

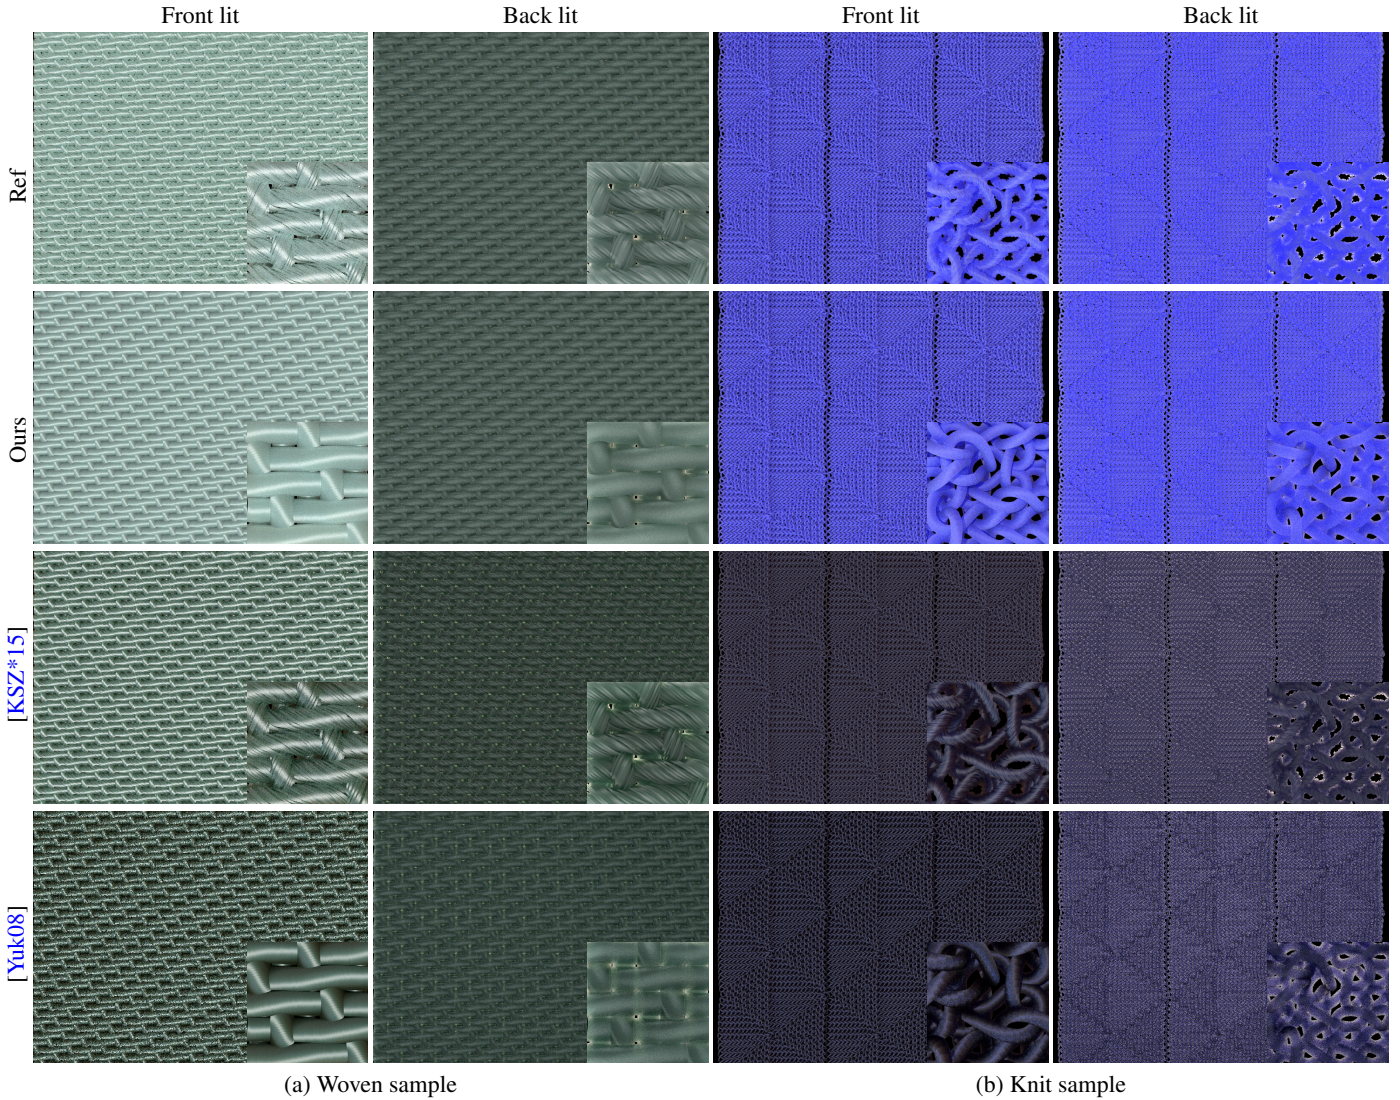


Figure 6: Large-scale comparisons with the reference (fiber-based path-tracer). Results use environment light containing only one light source in the front and back, referred as front and back lit, respectively. In addition to our reference, in the last two rows, we compare our method with [KSZ*15] that does not use diffuse lobe for fiber BCSDf, as well as [Yuk08] as classical dual scattering instead of our enhanced version.

stay mainly intact. However, the reference model — that explicitly models the fibers — performs slower as geometric complexity increases. In turn, memory usage increases as a function of the number of fibers, while our method remains independent of these geometric modifications. As shown in the 3, under the same configuration, our method is 4-10 times faster than the reference model, besides the memory usage is only 0.5%-4% of the reference model. All scenes shown in the paper are rendered using path tracing on an AMD Ryzen 9 7950X 16-Core Processor 4.50 GHz machine.

6. Results

In this section, we present the rendering results of our method and compare them with the reference, which was rendered using ex-

plicit fibers (R , TT and D). The cloth geometric parameters in all scenes refer to real measured data [ZJMB12]. Please also check out our accompanying video, where we provide more validation on animations.

Knitted glove. In Fig. 1, we present a comparison between our method and the reference in a knitted glove scenario. This scene is illuminated by both smooth ambient light and a sharp point light source from above. There are 69 yarns and 2417 fibers in the scene, with a twist angle of 25° for the fibers. We compare the rendering results of our method and the reference at the same noise level. Our method only requires the shading parameters of the fibers and the geometric parameters of the yarns to accurately simulate the single and multiple scattering of a yarn. Our method is $7\times$ faster than

the reference, and requires only 3% of the memory. In Fig. 7, we also show the rendering with different parameters of our method to show how easily our method can control the rendering through simple parameter adjustments. The glove yarn geometry is taken from the yarn dataset by Yuksel et al. [Yuk] and the fiber curves (reference) are generated procedurally [ZLB16].

Single Yarn. In Fig. 5, we compare our method to the reference on a single yarn, and we compare the rendering of a yarn containing one ply and three plies. Our method matches the reference in terms of both rendering results and implicitly geometric modeling at the ply level. From the rendering results, we compare the results of different densities and twists on single and multiple plies, and our method can accurately simulate the changes in specular highlights and multiple scattering. From the geometric modeling, our method can accurately model the ply-level geometry ply level on a single yarn.

Woven and knitted sample. Our method can be applied to both woven and knitted cloth. In Fig. 6, we demonstrate two sample cloth models under the front and back lighting conditions and compare the results with reference images. We also compare the rendering results with [KSZ*15] and the traditional dual scattering method [Yuk08]. For the woven sample, we use the actual measured geometry data of silk (refer to [ZLB16], silk 1) and for the fiber shading model, we use $a_r = \{0.2, 0.2, 0.2\}$, $a_{tt} = \{0.1, 0.4, 0.25\}$ and $a_d = \{0.03, 0.35, 0.1\}$. It can be seen that the traditional dual scattering model has lower saturation in the rendering results and more shadowing masking in the front lighting condition due to the two assumptions mentioned in §4.1. Our method matches the reference images while on average, it takes 17% render time of reference, 0.08% memory, and 14% number of bounces. The resolution of images is 768×576 and the performance statistics are listed in Table 3. We also use our model to modify the fiber parameters of the glove model shown in Fig. 1 and show them in Fig. 7.

7. Conclusion and Discussion

Limitations. In the oblique viewing direction, our geometric representation of the ply is not accurate enough, particularly at grazing angles. This limitation is due to the assumption that the rays enter the yarn perpendicularly and implicitly traced in the yarn cross section. Note that, this limitation is only for the geometric representation, our shading model does not suffer from this deficiency and remains accurate even at grazing angles.

Table 3: Performance statistics. All rendering times are counted at equal quality (EQ).

	Time (s)		Memory (MB)		#Bounce	
	Ref	Ours	Ref	Ours	Ref	Ours
Fig. 6.a	187	29	210	1.8	12.4	1.6
W/ twice fibers	265	29	420	1.8	17.4	1.6
Fig. 6.b	343	72	452	4.0	16.5	3.2
W/ twice fibers	547	72	905	4.0	21.3	3.2
Fig. 1	519	73	135	4.0	26.4	4.2




Figure 7: Rendered results using our method (a) with measured fiber parameters of Polyester material [ZLB16], (b) increased number of fibers (higher density), (c) different fiber twist, and (d) three-ply yarn with different albedo for each ply. In this figure, we show a close view at the grazing angle to emphasize that our implicit representation of yarn geometry does not result in noticeable artifacts.

Additionally, some of the limitations of traditional DS still apply to our extended model. For example, the assumption of uniform fiber distribution behind the shading point is still valid, which results in an inaccurate angular distribution of the forward and backward lobes for yarn with twisted fibers. However, since the width (variance) of the forward and backward lobes are relatively broad, the impact on the final rendered results is negligible.

Our method does not consider the contribution of fly-away fibers because the orientation of fly-away fibers is completely different from the regular fibers. However, it is possible to represent and render fly-away fibers by modeling them separately, a similar approach has already been used in [MGZJ20].

Future works. In addition to addressing the limitations, a straightforward future work avenue is to implement our BCSDf model for real-time rasterization-based applications. We also would like to extend our method with a multi-resolution manner to have an efficient level of detail, especially for far-field views.

Conclusion. We introduce an efficient and accurate model that can represent the analytical solution for single and multiple scattering of light at the yarn level, using our hierarchical geometric structure. Our proposed yarn-based model takes into account the physical properties of the yarn, ply and fibers and accurately computes the yarn shading, derived from the fiber-level BCSDf model. By incorporating the hierarchical geometry of yarns, our yarn-based shading model can produce highly realistic results with minimal computational cost. Furthermore, unlike previous works, our model is general and can be applied to render woven and knitted fabrics.

Acknowledgements

We thank Marc Droske for his valuable input. Ling-Qi Yan is

supported by gift funds from Adobe, Intel, Meta and XVerse. This project is supported by Wētā Digital.

References

- [AMF03] ADABALA, NEEHARIKA, MAGNENAT-THALMANN, NADIA, and FEI, GUANGZHENG. “Real-Time Rendering of Woven Clothes”. VRST ’03. Osaka, Japan: Association for Computing Machinery, 2003, 41–47. ISBN: 1581135696. DOI: [10.1145/1008653.1008663](https://doi.org/10.1145/1008653.1008663). URL: <https://doi.org/10.1145/1008653.1008663>.
- [dMH13] D’EON, EUGENE, MARSCHNER, STEVE, and HANIKA, JOHANNES. “Importance sampling for physically-based hair fiber models”. *SIGGRAPH Asia 2013 Technical Briefs*. 2013, 1–4 7.
- [DvGNK99] DANA, KRISTIN J., van GINNEKEN, BRAM, NAYAR, SHREE K., and KOENDERINK, JAN J. “Reflectance and Texture of Real-World Surfaces”. *ACM Trans. Graph.* 18.1 (Jan. 1999), 1–34. ISSN: 0730-0301. DOI: [10.1145/300776.300778](https://doi.org/10.1145/300776.300778). URL: <https://doi.org/10.1145/300776.300778>.
- [IM12] IRAWAN, PITI and MARSCHNER, STEVE. “Specular reflection from woven cloth”. *ACM Transactions on Graphics (TOG)* 31.1 (2012), 1–20 2.
- [Jac20] JACOB, WENZEL. *Mitsuba*. <https://www.mitsuba-renderer.org>. 2020 7.
- [JWH*22] JIN, WENHUA, WANG, BEIBEI, HASAN, MILOS, et al. “Woven Fabric Capture from a Single Photo”. *SIGGRAPH Asia 2022 Conference Papers*. 2022, 1–8 2.
- [KK89] KAJIYA, J. T. and KAY, T. L. “Rendering Fur with Three Dimensional Textures”. *Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’89. New York, NY, USA: Association for Computing Machinery, 1989, 271–280. ISBN: 0897913124. DOI: [10.1145/74333.74361](https://doi.org/10.1145/74333.74361). URL: <https://doi.org/10.1145/74333.74361>.
- [KSZ*15] KHUNGURN, PRAMOOK, SCHROEDER, DANIEL, ZHAO, SHUANG, et al. “Matching Real Fabrics with Micro-Appearance Models”. *ACM Trans. Graph.* 35.1 (2015), 1–1 2, 4, 5, 8, 9.
- [KWN*17] KHUNGURN, PRAMOOK, WU, RUNDONG, NOECKEL, JAMES, et al. “Fast rendering of fabric micro-appearance models under directional and spherical gaussian lights”. *ACM Transactions on Graphics (TOG)* 36.6 (2017), 1–15 2.
- [LN18] LOUBET, GUILLAUME and NEYRET, FABRICE. “A new microflake model with microscopic self-shadowing for accurate volume downsampling”. *Computer Graphics Forum* 37.2 (May 2018), 111–121. DOI: [10.1111/cgf.13346](https://hal.science/hal-01702000). URL: <https://hal.science/hal-01702000>.
- [MGJZ21] MONTAZERI, ZAHRA, GAMMELMARK, SØREN, JENSEN, HENRIK WANN, and ZHAO, SHUANG. “Practical Ply-Based Appearance Modeling for Knitted Fabrics”. *Eurographics Symposium on Rendering - DL-only Track*. Ed. by BOUSSEAU, ADRIEN and MCGUIRE, MORGAN. The Eurographics Association, 2021. ISBN: 978-3-03868-157-1. DOI: [10.2312/sr.20211297](https://doi.org/10.2312/sr.20211297).
- [MGJZ20] MONTAZERI, ZAHRA, GAMMELMARK, SØREN B, ZHAO, SHUANG, and JENSEN, HENRIK WANN. “A practical ply-based appearance model of woven fabrics”. *ACM Transactions on Graphics (TOG)* 39.6 (2020), 1–13 2, 3, 9.
- [MJC*03] MARSCHNER, STEPHEN R., JENSEN, HENRIK WANN, CAMMARANO, MIKE, et al. “Light Scattering from Human Hair Fibers”. *ACM Trans. Graph.* 22.3 (July 2003), 780–791. ISSN: 0730-0301. DOI: [10.1145/882262.882345](https://doi.org/10.1145/882262.882345). URL: <https://doi.org/10.1145/882262.882345>.
- [MXF*21] MONTAZERI, ZAHRA, XIAO, CHANG, FEI, YUN, et al. “Mechanics-Aware Modeling of Cloth Appearance”. English. *IEEE Transactions on Visualization and Computer Graphics* (Jan. 2021), 137–150. ISSN: 1077-2626. DOI: [10.1109/TVCG.2019.2937301](https://doi.org/10.1109/TVCG.2019.2937301).
- [SBDJ13] SADEGHI, IMAN, BISKER, OLEG, DE DEKEN, JOACHIM, and JENSEN, HENRIK WANN. “A practical microcylinder appearance model for cloth rendering”. *ACM Transactions on Graphics (TOG)* 32.2 (2013), 1–12 2.
- [SKZ11] SCHRODER, K, KLEIN, REINHARD, and ZINKE, ARNO. “A volumetric approach to predictive rendering of fabrics”. *Computer Graphics Forum*. Vol. 30. 4. Wiley Online Library, 2011, 1277–1286 2.
- [SKZ15] SCHRÖDER, KAI, ZINKE, ARNO, and KLEIN, REINHARD. “Image-Based Reverse Engineering and Visual Prototyping of Woven Cloth”. *IEEE Trans. Vis. Comp. Graph.* 21.2 (2015) 3.
- [WY17] WU, KUI and YUKSEL, CEM. “Real-time Fiber-level Cloth Rendering”. *Proceedings of I3D*. 2017. ISBN: 978-1-4503-4886-7/17/03. DOI: [10.1145/3023368.3023372](https://doi.org/10.1145/3023368.3023372). URL: <http://doi.acm.org/10.1145/3023368.3023372>.
- [YKJM12] YUKSEL, CEM, KALDOR, JONATHAN M., JAMES, DOUG L., and MARSCHNER, STEVE. “Stitch Meshes for Modeling Knitted Clothing with Yarn-Level Detail”. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2012)* 31.3 (2012), 37:1–37:12. DOI: [10.1145/2185520.2185533](https://doi.org/10.1145/2185520.2185533). URL: <http://doi.acm.org/10.1145/2185520.2185533>.
- [YSJR17] YAN, LING-QI, SUN, WEILUN, JENSEN, HENRIK WANN, and RAMAMOORTHY, RAVI. “A BSSRDF Model for Efficient Rendering of Fur with Global Illumination”. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2017)* 36.6 (2017) 3.
- [YTJR15] YAN, LING-QI, TSENG, CHI-WEI, JENSEN, HENRIK WANN, and RAMAMOORTHY, RAVI. “Physically-accurate fur reflectance: Modeling, measurement and rendering”. *ACM Trans. Graph.* 34.6 (2015), 1–13 4, 5.
- [Yuk] YUKSEL, CEM. *Yarn-level Cloth Models*. <http://www.cemyuksel.com/research/yarnmodels9>.
- [Yuk08] YUKSEL, CEM. “Dual Scattering for Real-Time Multiple Scattering in Hair”. *ACM SIGGRAPH 2008 Computer Animation Festival*. Los Angeles, California: ACM, 2008, 39. ISBN: 978-1-60558-345-7. DOI: [10.1145/1400468.1400494](https://doi.org/10.1145/1400468.1400494). URL: <http://doi.acm.org/10.1145/1400468.1400494>.
- [ZHRB13] ZHAO, SHUANG, HAŠAN, MILOŠ, RAMAMOORTHY, RAVI, and BALA, KAVITA. “Modular flux transfer: efficient rendering of high-resolution volumes with repeated structures”. *ACM Transactions on Graphics (TOG)* 32.4 (2013), 1–12 2.
- [ZJMB11] ZHAO, SHUANG, JAKOB, WENZEL, MARSCHNER, STEVE, and BALA, KAVITA. “Building volumetric appearance models of fabric using micro CT imaging”. *ACM Transactions on Graphics (TOG)* 30.4 (2011), 1–10 2, 3.
- [ZJMB12] ZHAO, SHUANG, JAKOB, WENZEL, MARSCHNER, STEVE, and BALA, KAVITA. “Structure-aware synthesis for predictive woven fabric appearance”. *ACM Transactions on Graphics (TOG)* 31.4 (2012), 1–10 8.
- [ZLB16] ZHAO, SHUANG, LUAN, FUJUN, and BALA, KAVITA. “Fitting procedural yarn models for realistic cloth rendering”. *ACM Transactions on Graphics (TOG)* 35.4 (2016), 1–11 9.
- [ZZW*22] ZHU, JUNQIU, ZHAO, SIZHE, WANG, LU, et al. “Practical Level-of-Detail Aggregation of Fur Appearance”. *ACM Trans. Graph.* 41.4 (July 2022). ISSN: 0730-0301. DOI: [10.1145/3528223.3530105](https://doi.org/10.1145/3528223.3530105). URL: <https://doi.org/10.1145/3528223.3530105>.

Appendix A: Derivation of a_B^{ply} .

In Eqn. 10, we accumulate the attenuation of each possible F^+BF^+ and $F^+BF^+BF^+BF^+$ path to define the a_B^{ply} . Here we derive the

analytical solution to solve the a_{1B}^{ply} and a_{3B}^{ply} .

$$\begin{aligned} a_{1B}^{\text{ply}}(\theta_d) &= a_B(\theta_d) \sum_{i=0}^{n_I-1} a_F^{2i}(\theta_d) \\ &= a_B \left(\frac{1 - a_F^{2n_I}}{1 - a_F^2} \right) \end{aligned}$$

$$\begin{aligned} a_{3B}^{\text{ply}}(\theta_d) &= a_B^3(\theta_d) \sum_{i=1}^{n_I-1} \sum_{j=0}^{i-1} \sum_{k=j+1}^{n_I-1} a_F^{2(i-j-1+k)}(\theta_d), \\ &= a_B^3 \frac{(1 - 2n_I)a_F^{2n_I+2} + a_F^{4n_I} + (2n_I - 1)a_F^{2n_I} - a_F^2}{(a_F^2 - 1)^3} \end{aligned}$$

$$a_B^{\text{ply}} = a_B \left(\frac{1 - a_F^{2n_I}}{1 - a_F^2} \right) + a_B^3 \frac{(1 - 2n_I)a_F^{2n_I+2} + a_F^{4n_I} + (2n_I - 1)a_F^{2n_I} - a_F^2}{(a_F^2 - 1)^3}$$

Then a_B^{ply} can be easily integrated in the implementation.