



Neural Appearance Model for Cloth Rendering

G. Y. Soh  and Z. Montazeri 

University of Manchester, United Kingdom

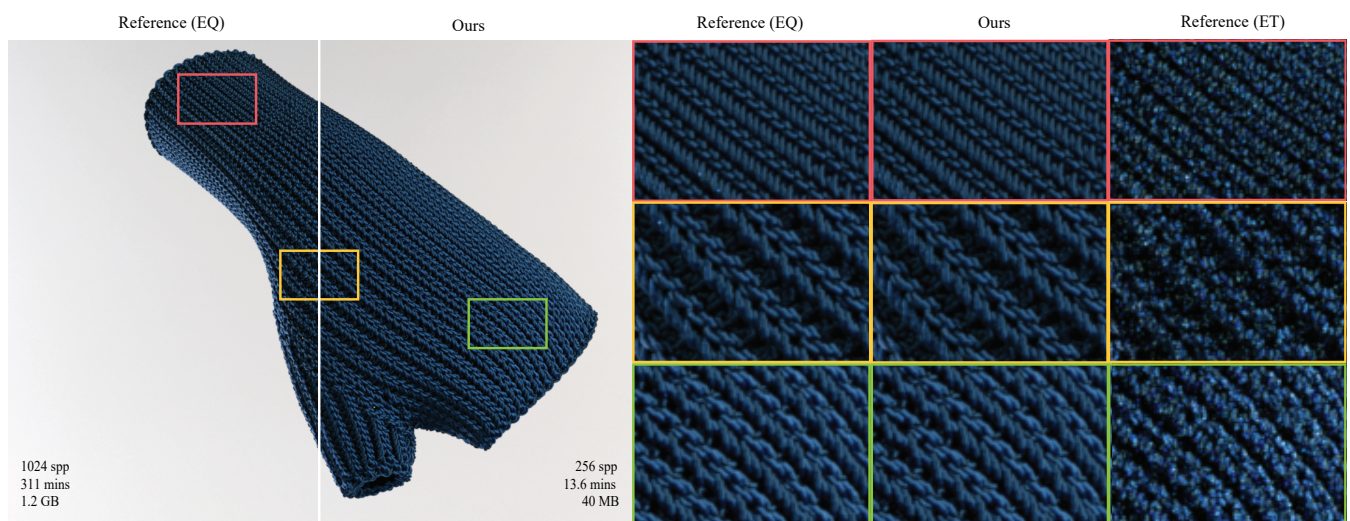


Figure 1: In this scene, we compare our neural aggregated model against the reference, which uses explicit fiber-based geometry and scattering model. The left shows the renderings side by side with equal quality (EQ) of the knitted fleece which is 23 times slower and requires 300 times more memory than ours. whereas the right shows the close-up crops to show the equal time (ET) is much noisier than ours.

Abstract

The realistic rendering of woven and knitted fabrics has posed significant challenges throughout many years. Previously, fiber-based micro-appearance models have achieved considerable success in attaining high levels of realism. However, rendering such models remains complex due to the intricate internal scatterings of hundreds of fibers within a yarn, requiring vast amounts of memory and time to render. In this paper, we introduce a new framework to capture aggregated appearance by tracing many light paths through the underlying fiber geometry. We then employ lightweight neural networks to accurately model the aggregated BSDF, which allows for the precise modeling of a diverse array of materials while offering substantial improvements in speed and reductions in memory. Furthermore, we introduce a novel importance sampling scheme to further speed up the rate of convergence. We validate the efficacy and versatility of our framework through comparisons with preceding fiber-based shading models as well as the most recent yarn-based model.

CCS Concepts

• **Computing methodologies** → **Reflectance modeling**;

1. Introduction

Fabrics are important in our everyday lives and their virtual representation has long been a key focus in computer graphics research. Fabrics, with their detailed structure of fibers, plies, and

yarns, present a unique hierarchical geometric structure at each aggregation level, offering a wide array of appearances for different cloth types.

The challenge of accurately modeling the detailed geometry and

scattering of fabrics has led to the development of different methods, mainly split into curve-based and surface-based shading models. Curve-based models, using the Bidirectional Curve Scattering Distribution Function (BCSDF), aim to explicitly represent individual elements like fibers [KSZ*15], plies [MGZJ20] and yarns [ZMA*23], similar to methods used in hair rendering. However, these models, while accurate, face challenges like long rendering times and high storage needs.

Micro-appearance models focus on representing fabrics at the microscale, detailing each fiber using high-resolution volumes or fiber meshes [ZJMB11]. These models are great at rendering with high detail but are limited in practical use due to their data-intensive nature and their challenges in manipulation and rendering.

In contrast, surface-based models, which depict fabric as a 2D sheet and use specific reflectance models for appearance, are known for being lightweight and user-friendly e.g. [SBDJ13]. These models, widely used in the computer graphics industry, can accurately reproduce the overall appearance of fabrics. However, they often fail to capture the fine details necessary for realistic close-up images.

In this paper, we aim to combine the light scattering of a twisted yarn, made up of hundreds of fibers, by simulating the paths of many light rays into the yarn and analyzing their scattering properties. From this analysis, we show that the scattering can be described as three distinct components, and we introduce a new way to model each component using various neural networks and analytical solutions. Additionally, we derive an analytical importance sampling scheme that closely matches the combined scattering distribution. We demonstrate that our model is able to run up to 23 times faster while using up to 600 times less memory when compared to previous fiber-based methods. The memory gain is directly dependent on the fiber count which is often a few hundred. In summary, our main contributions include:

- We introduce a novel neural framework for modeling the light scattering within a bundle of fibers in the yarns. By dividing the scattering into components, we can efficiently model various types of yarns across a broad range of parameters. Our proposed method runs significantly faster and uses substantially less memory.
- We further improve on existing neural network approaches by using the channel-wise PReLU activation function to increase performance. We demonstrate its effectiveness by comparing its performance against various model architectures.
- From our observations, we derive a new analytical fitting of the scattering for importance sampling. We have managed to derive a new observation-based empirical and invertible importance sampling scheme that matches the scattering distribution to further accelerate the rate of convergence.

2. Prior Work

Surface-Based Cloth Models - Cloth rendering has been a subject of extensive research, with various models being developed to achieve a balance between realism and computational efficiency. Traditional models have often depicted cloth as 2D surfaces, utilizing Bidirectional Reflectance Distribution Functions (BRDFs)

or Bidirectional Texture Functions (BTFs) to illustrate light-cloth interactions [SSK03; AMF03; IM12; SBDJ13; RJGW19; KMX*21; JWH*22; ZJA*23]. While these surface-based models are lightweight and capable of producing high-quality results at mid-to-far distances, they typically lack the fine-grained details necessary for close-up views.

Micro-appearance Cloth Models - On the contrary, micro-appearance models have emerged, focusing on the fabric's micro-geometry down to the fiber level, offering a high fidelity and detail [SKZ11; ZJMB11; KSZ*15; LN18; MXF*21; ACG*17]. However, the high complexity of these models presents a significant challenge in rendering them efficiently. Various precomputation-based rendering methods have been developed to address this, such as the techniques proposed by [ZHRB13; KWN*17; LZB17] to improve performance and GPU-based methods developed by [WY17] for procedurally generated fabrics. Nevertheless, these methods often compromise either on performance or physical accuracy, as well as being difficult to edit and render.

Aggregation Based Techniques - In recent years, aggregation-based methodologies have been introduced to the domain of cloth rendering, aiming to model the multiple scatterings of a bundle of fibers implicitly. Montazeri et al. [MGZJ20; MGJZ21] pioneered an aggregated technique that encapsulates the light scatterings of individual fibers, approximating the overall appearance at the ply level for woven and knitted fabrics, respectively; later followed by the yarn-level extension [KZP*24]. However, their model, while being fast and practical, is predominantly observation-driven and not efficient for yarns with a high number of plies.

Zhu et al. [ZZW*22] advanced the field by proposing a technique to aggregate the scatterings of a bundle of straight fur fibers in a data-driven manner. They then parameterize the aggregated scattering by fitting analytical lobes, followed by the training of a neural network to predict the parameters for the lobes. This model does not accommodate twisted fibers and, being a far-field model, cannot represent yarn-level highlights at close-up views. In a subsequent study [ZMA*23], the authors introduced an analytical solution designed to accurately approximate the multi-scattering of yarn by utilizing dual scattering theory. However, this model relies heavily on the assumptions inherited from dual scattering theory and also imposes additional assumptions on the fiber shading model. In contrast, our work, while employing similar fiber scattering models and micro-geometry, presents a more generalized model capable of fitting any yarn without necessitating specific assumptions.

Neural BRDF Representation - [CNN22; SRRW21] was one of the firsts to leverage machine learning to represent BRDFs and achieve a high compression rate while preserving the fidelity of the BRDF. In this paper, we improve on Sztrajman et al.'s [SRRW21] framework to support aggregated yarn scattering, as we demonstrate that using their framework in a naïve manner do not produce optimal results.

3. Preliminaries

Every yarn is made up of twisted plies, which in turn consist of hundreds of strands called fibers. In our study, the primary aim is

Table 1: List of important symbols for our neural yarn shading

Notation	Definition
N	number of fibers in a yarn
α	twist factor
ρ	fiber density in a yarn
C_R	attenuation of fiber reflection
C_{TT}	attenuation of fiber transmission
β_{f_R}	longitudinal roughness of fiber reflection
$\beta_{f_{TT}}$	longitudinal roughness of fiber transmission
$\gamma_{f_{TT}}$	azimuthal roughness of fiber transmission
ω_i	incoming direction relative to yarn frame
ω_o	outgoing direction relative to yarn frame
$\tilde{\omega}_i$	incoming direction relative to surface fiber frame
$\tilde{\omega}_o$	outgoing direction relative to surface fiber frame
θ_i	incoming longitudinal angle
θ_o	outgoing longitudinal angle
ϕ_i	incoming azimuthal angle
ϕ_o	outgoing azimuthal angle
S	our aggregated yarn scattering function
S_R	our aggregated yarn direct reflection component
S_T	our aggregated yarn direct transmission component
S_M	our aggregated yarn multi-scattering component
S_f	fiber scattering function

to aggregate a single-ply geometry while explicitly tracing interactions between plies for multi-ply yarn. The arrangement of the fibers around the yarn is characterized by the parameters N , ρ , and α . Here, N represents the number of fibers in the yarn, ρ represents the fiber density, and α describes the twist factor.

$$\rho = \frac{Nr^2}{R^2}, \quad \alpha = \frac{2R\Delta n}{\Delta l}, \quad (1)$$

where r is the fiber radius, R is the yarn radius, n is the number of revolutions, and l denotes the length along the yarn. Importantly, these parameters are defined such that they are invariant to the yarn's overall scale, allowing us to use our fitted model on all scales of the yarn with the same parameters, without having to re-train the neural networks or re-fit the parameters. The list of all parameters is detailed in Table 1.

3.1. Fiber Shading Model

In this paper, our fiber shading model is based on the method by Khungurn et al. [KSZ*15], where fibers are modeled as glass-like tubes and the scatterings are split into Reflection (R) and Transmission (TT) lobes.

$$S_f(\omega_i, \omega_o) = S_f(\theta_i, \theta_o, \phi_i, \phi_o) = \sum_k M_k(\theta_i, \theta_o) N_k(\phi_i, \phi_o), \quad (2)$$

where $k = \{R, TT\}$. The incident and outgoing directions ω_i , ω_o are parameterized into the longitudinal angle θ and azimuthal angle ϕ using the coordinate system defined in Marschner et al. [MJC*03]. M represents the scattering in the longitudinal plane,

and N represents the scattering in the azimuthal plane. They are defined as:

$$\begin{aligned} M_R(\theta_i, \theta_o) &= F_R(\theta_i) \bar{g}(\theta_o; -\theta_i, \beta_{f_R}) \\ M_{TT}(\theta_i, \theta_o) &= C_{TT}(1 - F_R(\theta_i)) \bar{g}(\theta_o; -\theta_i, \beta_{f_{TT}}) \\ N_R(\phi_i, \phi_o) &= \frac{1}{2\pi} \\ N_{TT}(\phi_i, \phi_o) &= f(\phi_o; \phi_i + \pi, \gamma_{f_{TT}}^{-2}), \end{aligned} \quad (3)$$

where C_R , C_{TT} represents the attenuation of each component, β_{f_R} , $\beta_{f_{TT}}$ represents the longitudinal roughness, and $\gamma_{f_{TT}}$ represents the azimuthal roughness. \bar{g} is the normalized Gaussian function defined in Khungurn et al. [KSZ*15] and f denotes the von Mises distribution. Furthermore, F_R is the Fresnel term and is approximated via Schlick's approximation [Sch94]:

$$F_R(\theta_i) = C_R + (1 - C_R)(1 - \cos \theta_i)^5. \quad (4)$$

3.2. Yarn Shading Frame

In this work, we found it useful to describe the light scattering in terms of two separate shading frames, the yarn shading frame, and the surface fiber shading frame. The yarn shading frame is defined as a traditional anisotropic surface shading frame on the yarn cylinder, with the incident and outgoing directions parameterized with longitudinal angle θ and azimuthal angle ϕ . The normal of the frame is aligned with the normal of the cylinder surface, while the tangent of the frame is aligned in the direction of the yarn tangent. We chose this in contrast to existing hair literature, where a longitudinal angle θ and an azimuthal offset h are used, to make the process of finding the surface fiber shading frame easier. The surface fiber shading frame describes the fiber shading frame on the surface of the yarn, and using the coordinates system of Marschner et al. [MJC*03] with $\phi = 0^\circ$ when pointing towards the surface normal. The frame is rotated around the surface normal due to the fiber twist. In our paper, we denote the directions relative to the yarn shading frame with ω_i and ω_o , while the directions relative to the surface fiber shading frame as $\tilde{\omega}_i$ and $\tilde{\omega}_o$ which can be defined as:

$$\tilde{\omega}_i = M\omega_i, \quad \tilde{\omega}_o = M\omega_o, \quad M = \begin{bmatrix} \cos \tilde{\phi} & -\sin \tilde{\phi} & 0 \\ \sin \tilde{\phi} & \cos \tilde{\phi} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (5)$$

where $\tilde{\phi} = \tan^{-1}(\pi\alpha)$. Derivation of the angle $\tilde{\phi}$ can be found in the appendix below.

4. Our Aggregated Shading Model

In alignment with the methodology proposed by Zhu et al. [ZZW*22], the aggregation of yarn fibers is achieved by encapsulating them within a closely bounded cylinder. We aggregate the yarn scattering by simulating many light rays into the yarn and recording their exiting radiance and direction to obtain the Radiance Distribution Map (RDM). The RDM is a 4-dimensional map of the exiting radiance for a given ω_i and ω_o and is parameterized by θ , ϕ . Further details on obtaining the RDM are given in §5.3. We then propose a framework to model the RDM by observation

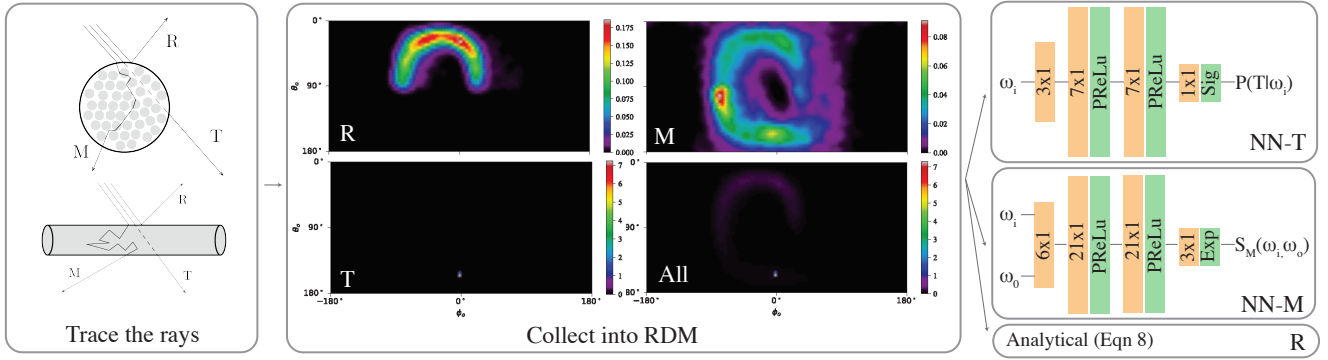


Figure 2: Overview of our pipeline. The first step is to explicitly trace the rays and label them into three components to gather the data (direct transmission T , direct reflection R , multi-scattering M). Next, collect them into Radiance Distribution Maps (RDM). Here, we separate each component of the RDM (T , R , M) to demonstrate the vastly different scales and distributions of each component and visualize for when $\theta_i = 45^\circ$ and $\phi_i = 0^\circ$. Lastly, the networks to learn the T and M components are visualized while R is being computed analytically.

and analysis and split the RDM into 3 components to model them individually.

It might intuitively seem advantageous to model the RDM directly by fitting a neural network to it. However, our experiments suggest that this approach is not optimal. Initially, it was observed that at certain incoming angles, specifically at grazing azimuthal angles, a substantial amount of light traverses through the yarn cylinder without interacting with any fibers. In such instances, most of the light is directly transmitted and exhibits Dirac delta distribu-

tions, resulting in the corresponding RDM displaying sharp lobes with pronounced peak values. Such distributions pose a significant learning challenge for the neural network due to their high values.

Furthermore, a considerable fraction of the brightness within the RDM is attributed to the paths characterized by a single bounce. These paths, interacting with a single fiber on the yarn's surface before exiting, create the highlights of the yarn and introduce abrupt alterations in brightness within the RDM. By isolating these paths into a distinct component, we achieved more precise highlights and facilitated the learning process for the neural network regarding the remaining data. Consequently, we introduce the subsequent shading model as a mixture of separate components T , R , and M , corresponding to the Direct Transmission Component, Direct Reflection Component, and Multi-Scattering Component respectively:

$$S(\omega_i, \omega_o) = S_T(\omega_i, \omega_o) + S_R(\omega_i, \omega_o) + S_M(\omega_i, \omega_o), \quad (6)$$

where the T component models the light paths that directly pass through the yarn without intersecting many fibers, the R component models the light paths that hit a single fiber and are reflected away, and the M component models the multiple scattering of light within the yarn before exiting. The components T and M are more complex and hence modelled by a neural network, while the R component can be found analytically. By splitting the shading model into separate components, we can better fit each lobe more accurately, whilst using fewer parameters for the neural network, increasing computational efficiency. The first column in Fig. 2 illustrated the pathways associated with each component, followed by the visualization of the distributions of each component. Fig. 3 visualize the appearance of each component to showcase their contribution individually.

4.1. Direct Transmission Component

The Direct Transmission Component of our model represents the fraction of incoming light that directly passes through the yarn without intersecting any fibers, given a specific incident direction ω_i . It becomes particularly prominent in yarn assemblies with lower




Figure 3: The contribution of each component (direct transmission T , direct reflection R , multi-scattering M) to the final appearance of the yarn.

fiber densities, where a high proportion of light rays pass through directly, resulting in a more translucent appearance. Its influence is also more pronounced at grazing azimuthal angles. Consequently, we incorporate this component into the final scattering function. This component can be mathematically expressed as:

$$S_T(\omega_i, \omega_o) = P(T|\omega_i)\delta(\omega_i + \omega_o), \quad (7)$$

where $\delta(\omega_i + \omega_o)$ is the Dirac delta distribution, which is zero except when $\omega_i = -\omega_o$. The probability $P(T|\omega_i)$ is multiplied with the Dirac delta distribution $\delta(\omega_i + \omega_o)$ to determine the radiance of the transmission component. Instead of fitting $S_T(\omega_i, \omega_o)$ directly with a neural network, we fit $P(T|\omega_i)$. This component is a two-dimensional map and can easily be modelled by a lightweight neural network.

4.2. Direct Reflection Component

Since the Direct Reflection Component is the reflection of fibers with a single bounce, this component corresponds to a bright highlight on the yarn surface. Therefore, this component contributes to a sharp change in radiance in the RDM. Hence, it would be beneficial to model this component analytically as opposed to fitting it with a neural network as this would allow us to achieve more accurate highlights, while simultaneously allowing the neural networks to converge at a faster rate with the other parameters. We model this component as a single fiber scattering relative to the surface fiber shading frame on the upper hemisphere of the surface.

$$S_R(\omega_i, \omega_o) = \begin{cases} (1 - P(T|\omega_i))S_f(\tilde{\omega}_i, \tilde{\omega}_o) & \text{if } \omega_i \cdot \mathbf{n} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

4.3. Multi-Scattering Component

The Multi-Scattering Component captures the detailed interactions among fibers within the ply or yarn and is represented as $S_M(\omega_i, \omega_o)$ in the scattering function. By using a neural network, we can effectively learn the distribution of these interactions, creating a robust multi-scattering model. This method is especially useful for modeling yarn aggregation to capture the scatterings of more complex yarn geometries, such as twist, a feature that the existing studies overlooked [ZZW*22]. For more detailed information and specific details about the network, please refer to §5.3.

4.4. Importance Sampling

Given that pieces of cloth are composed of numerous yarns, the inter-reflection amongst the yarns significantly influences the overall visual appearance. It is important to employ an advanced importance sampling scheme to reduce variance as showcased in Fig. 4. Nonetheless, due to the complexity of light scattering within a yarn when utilizing a neural network in our approach, we are precluded from using the Bidirectional Scattering Distribution Function (BSDF) for importance sampling. Consequently, we chose to fit an invertible analytical approximation of the data to enhance the sampling of the distribution. Sztrajman et al. [SRRW21] utilized Blinn-Phong lobes to fit the distribution of their Neural BRDFs. However, given that the scattering of light within a yarn does not center around the half angle, the Blinn-Phong lobe is a poor fit for




Figure 4: Comparison of uniform sampling vs our proposed importance sampling scheme. The images are rendered at 64spp and demonstrate that our importance sampling scheme significantly reduces the variance with less noise.

our model. From our observations of the multi-scattering component, we found that the light mainly scatters around the upper half of the cone centered at the fiber tangent at the yarn surface. Thus, we propose the following importance sampling scheme:

- *Sample lobe k* - We sample the lobes proportional to the energy of their lobes. Since $S_T(\omega_i, \omega_o)$ is comprised of light passing-through with a probability of $P(T|\omega_i)$, the proportion of energy can be described as $P(T|\omega_i)$ directly. The remaining portion of samples can be split proportionally according to the energy of $S_R(\omega_i, \omega_o)$ and $S_M(\omega_i, \omega_o)$, which can be approximated by a constant κ_R which is pre-computed beforehand based on the computed RDM.
- *Sample outgoing direction* - For the direct transmission component, we sample in the direction $\omega_o = -\omega_i$ to simulate the light passing through the yarn. The direct reflection is sampled similarly to the fiber's distribution on the yarn surface. It is done by sampling the longitudinal angles via a normalized Gaussian around $-\theta_i$ with the standard deviation corresponding to the fiber reflection's longitudinal roughness β_{f_r} , while the azimuthal angle is uniformly distributed on the upper cone in the range $[-\pi/2, \pi/2]$. The remaining multi-scattering component is sampled via two lobes which are derived from careful observations of the RDM. The first lobe is comprised of a distribution similar to the direct reflection component but with a different longitudinal and azimuthal roughness. It is defined by a longitudinal Gaussian distribution $g(\theta_i, -\theta_o, \beta_M)$ and azimuthal von Mises distribution $f(\phi_i, 0^\circ, \gamma_M)$, where azimuthal angle is zero at \mathbf{n} . The second distribution is described by a simple uniform sphere to capture the remaining directions not covered by the first lobe. The two lobes are split with a parameter κ_M . The parameters β_M , γ_M , and κ_M are to be fitted beforehand.
- *Compute the PDF* - The PDF can be described as a mixture of the lobes and can be computed as:

$$\text{pdf}(\omega_i, \omega_o) = \sum_{k=\{T,R,M\}} p_k \text{pdf}_k(\tilde{\omega}_i, \tilde{\omega}_o) \quad (9)$$

Here, the PDF for each component is defined as:

$$\begin{aligned} \text{pdf}_T(\tilde{\omega}_i, \tilde{\omega}_o) &= \delta(\tilde{\omega}_i + \tilde{\omega}_o) \\ \text{pdf}_R(\tilde{\omega}_i, \tilde{\omega}_o) &= \frac{g(\theta_o; -\theta_i, \beta_{f_R})}{\pi \cos \theta_o} \\ \text{pdf}_M(\tilde{\omega}_i, \tilde{\omega}_o) &= \frac{\kappa_M g(\theta_o; -\theta_i, \beta_M) f(\phi_o; 0^\circ, \gamma_M^{-2})}{\cos \theta_o} + \frac{1 - \kappa_M}{4\pi} \end{aligned} \quad (10)$$

with their proportions:

$$\begin{aligned} p_T &= P(T|wi) \\ p_R &= (1 - p_T)\kappa_R \\ p_M &= (1 - p_T)(1 - \kappa_R) \end{aligned} \quad (11)$$

where $g(\theta_o; -\theta_i, \beta)$ represents a Gaussian normalized in the range $[-\pi/2, \pi/2]$ with mean $-\theta_i$ and standard deviation β . $f(\phi_o; 0^\circ, \gamma_M^{-2})$ represents the von Mises distribution with mean 0° and roughness γ_M .

5. Our Neural Approach

Due to the intricate nature of light scattering within a yarn, we opted to model it using neural networks, inspired by the success of Sztrajman et al. [SRRW21] in accurately and efficiently modeling measured BRDFs. Additionally, this approach provides the generality and flexibility to model various yarn types without making assumptions about the underlying geometry. Beyond the enhancements described in §4 to boost the neural network's performance, we have also refined their base architecture to achieve higher accuracy with minimal runtime costs by employing the channel-wise Parametric Rectified Linear Unit (PReLU) [HZRS15] activation function instead of the Rectified Linear Unit (ReLU) [NH10] activation presented in their paper. A comparison with this model naïvely is shown in Fig. 5.

In brief, channel-wise PReLU allows each channel of the input its own learnable parameter, which provides the model with additional flexibility to learn more complex representations without a substantial increase in computational cost and mitigates issues related to the "dying ReLU" problem. The "dying ReLU" problem refers to the phenomenon where neurons in a network become inactive and only output zero during training, essentially ceasing to learn or update and thereby reducing the capacity of the model. This often occurs when a large gradient flows through a ReLU neuron, updating the weights in such a way that the neuron will always output zero. PReLU helps to avoid this issue by maintaining active learning and adapting its negative slope to the learned features of the input data.

Additionally, We chose the channel-wise PReLU activation function over ReLU because it introduces additional trainable parameters for the negative values of ReLU, allowing the neural network more flexibility to overfit with nearly no extra runtime cost, while avoiding the instability of the dying ReLU problem, which is more prevalent in smaller neural networks. Please refer to §6 for additional details on the performance of various neural network architectures and activation functions.




Figure 5: We compare our neural network approach with the naïve approach by contrasting them with the reference. Our method models each component of the RDM as described in §4, while the naïve approach models the RDM directly using the framework described in [SRRW21]. Our approach successfully models the reference, including the subtle multi-scatterings, while the latter does not.

5.1. Data Generation

To generate data for computing the RDM and preparing the training data for the neural networks, we initially establish our foundational single-ply yarn geometry, as previously detailed in §3. A bounding cylinder is defined around the yarn, and light rays, each possessing an initial weight of 1, are projected at random directions ω_i , uniformly distributed over a hemisphere, into the yarn. Monte Carlo random walks are subsequently utilized to trace the interactions of each ray with the fibers until it exits the yarn cylinder. For each sample, variables such as the incident angle, outgoing angle, outgoing weight, and the number of bounces (depth) are documented. Our dataset consists of 50-100 million sampled rays that are fully traced for each five yarn materials, with a maximum bounce depth of 200 on average. This sample collection process persists until convergence is attained.

5.2. Direct Transmission Neural Network

To acquire the training data for this network, we compute the probability of transmission for a given incoming direction with the gathered samples:

$$P(T|\omega_i) = \frac{\text{Count}(T|\omega_i)}{\text{Count}(All|\omega_i)}. \quad (12)$$

We gather samples using the method outlined in §5.1, then organize the data into two 22x90 histograms, representing ω_i with θ_i and ϕ_i bins across a range of 90x360 degrees. The first histogram calculates the number of direct transmission paths, while the second histogram counts the total number of paths. Subsequently, we

divide the first histogram by the second to derive the probability map $P(T|\omega_i)$.

Next, we train a lightweight neural network on the probability map. The network, which takes ω_i as a unit Cartesian vector and predicts $P(T|\omega_i)$, is configured with two hidden layers and follows a 3-7-7-1 structure. The hidden layers utilize the channel-wise Parametric Rectified Linear Unit (PReLU) activation function, and the output layer employs the Sigmoid activation function. The model is trained using the Mean Squared Error (MSE) loss function, coupled with the Adam optimizer. Our network architectures are illustrated in the last column of Fig 2.

5.3. Multi-Scattering Neural Network

The neural network is trained on the multi-scattering component of the RDM. To prepare the data for the neural network, it is necessary to isolate the multi-scattering component from the collected samples. Initially, samples with a depth of 0 are removed to exclude the direct transmission samples, along with samples having a depth of $1 \cap (\omega_o \cdot \mathbf{n}) > 0$ to exclude the direct reflection samples. Subsequently, a weighted 4D histogram is computed from the remaining data into $22 \times 90 \times 45 \times 90$ bins of θ_i , ϕ_i , θ_o , and ϕ_o , each spanning across the respective ranges of $90 \times 360 \times 180 \times 360$ degrees. The data is then divided by the number of samples in each incident bin and the solid angle in each outgoing bin to obtain the radiance at each bin. With the multi-scattering component RDM available, samples are randomly drawn from it to generate our training data.

The neural network is configured to accept Cartesian unit vectors ω_i and ω_o as inputs and to output r, g, b radiance values. The model incorporates two hidden layers with a 6-21-21-3 structure. The hidden layers utilize channel-wise PReLU activation functions, while the final layer employs the exponential activation function. The model is trained using the Mean Squared Error (MSE) loss function and optimized with the Adam optimizer.

As previously noted in §4.3, the multi-scattering component is represented by $S_M(\omega_i, \omega_o)$. However, in practice, the neural network was configured to model the product of $S_M(\omega_i, \omega_o)$ and $(\omega_i \cdot \mathbf{n})$. Here, $(\omega_i \cdot \mathbf{n})$ represents the cosine foreshortening factor and is inherently included in the RDM as we record the radiance for each (ω_i) and (ω_o) directly.

6. Model Analysis and Ablation

In this section, we perform an ablation study about the neural network used in the multi-scattering component by comparing the performance of the model with different architectures. The model is trained on polyester until convergence (40 epoch). Fig. 6 shows the final loss of various model architectures along with different activation functions. It can be seen that channel-wise PReLU consistently outperforms other activation functions with the same model architecture. This is due to the additional trainable parameters of channel-wise PReLU, which gives the model more flexibility at a negligible increase in runtime cost. It also can be seen that increasing the model weights from our base model to 6-21-64-64-21-3 increases the model size by 10.6 times while only offering a 9% decrease in loss. From this, we can see that the model does not need

to be overly large, and performs well even with a smaller number of weights.

7. Results

In this section, we validate our model and evaluate its performance by comparing renderings with our model to reference images generated by rendering the explicit fiber geometry [KSZ*15] as well as the hierarchical yarn-based model [ZMA*23]. For all the materials presented, besides polyester, we have used the fiber shading parameters given in Khungurn et al. [KSZ*15] which were computed by fitting the parameters to match real-life photographs. The parameters of polyester are determined ad hoc to demonstrate the flexibility of our framework. We then wrap these fibers into yarns with given fiber geometry parameters N , ρ , and α . A summary of the parameters can be found in Table 3. All images were rendered with path tracing on Mitsuba 3 [JSR*22], including neural network inference, using an Intel Core i7-10750H 6 Core Processor 2.60GHz machine, while neural network training was done on an NVIDIA GeForce RTX4080 (Mobile). The computation time required to gather the RDM is around a minute on an RTX4080 (Mobile). The average time it takes to train a neural network per material for the direct transmission and multi-scattering components are 30 seconds and 30 minutes respectively.

7.1. Reference Comparisons

3-Ply Knitted Glove - In this section, we rendered a scene with a 3-ply knitted glove. The base yarn curves defining the glove were

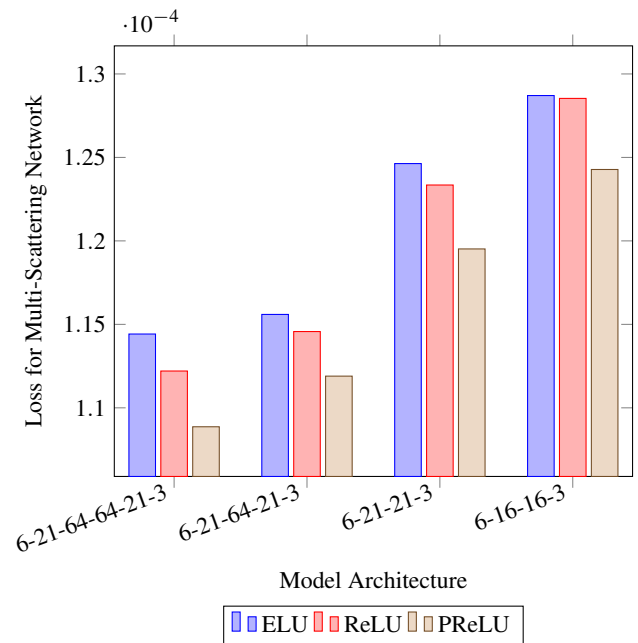


Figure 6: The performance of various neural network architectures on the multi-scattering component. We demonstrate the effectiveness of channel-wise PReLU compared to common activation functions.

taken from [Yuk20; YKJM12; WY17] and wrapped with 3-ply. The plies are then wrapped with fibers procedurally to generate the ground truth image [ZLB16]. For Fig. 1 we rendered the scene at a resolution of 1080x1080. Our model matches the ground truth very well and performs 23 times faster while using around 300 times less memory. The scene is lit with an environmental map along with two spherical lights on the top-right and bottom-left corners. We also rendered the scene with different fiber parameters in Fig. 7 to highlight the flexibility of our model.

Close-up Yarn - In Fig. 8, we compare our model against the reference on a close-up view of a yarn with varying fiber parameters. The scene is rendered at a resolution of 512x512 and then cropped to an appropriate size. The reference is rendered at 1024 spp, while our model is rendered at 64spp and performs on average 30 times faster. Our model can match the overall yarn appearance despite not having explicit fiber geometry. Please note for a multi-ply yarn, our model aggregates the fiber bundle of a single ply and we rely on the renderer to take the ply-ply interactions into account.

Woven and Knitted Fabric - In Fig. 9, we rendered our images using the dataset of yarn curves by Leaf et al. [LWS*18]. The curves were interpolated and tiled into an appropriate size. All the images were rendered at a resolution of 720x720. All the reference images were rendered at 1024spp except for silk and cotton which were rendered at 4096spp as they take longer to converge due to their very high albedo. From our comparisons, our model matches the reference images very well and can accurately recreate yarn-level details even in the absence of explicit fiber geometry. However, although still visually accurate, we do note that cotton has difficulty matching the reference which is discussed further in the limitations section in §8. Our model performs around 11-17 times faster while using around 200-600 times less memory. Please refer to Table 2 for the full statistics.

7.2. Comparisons with Zhu et al. 2023

As depicted in Fig. 10, we demonstrate that our approach not only achieves faster rendering speeds, as detailed in Table 2, but also more accurately replicates the reference fiber-based appearance model by Khungurn et al. [KSZ*15]. Our model's superiority is due to our neural data-driven methodology that adapts more flexibly, allowing for an exact fit to the reference. In contrast, [ZMA*23] uses an approximated fiber appearance model, which does not model Fresnel effects, and often requiring manual adjustments to align with the reference model. Notably, we use the exact same set of parameters and values across the three models (reference, ours, and [ZMA*23]) without any post-tweaking.

8. Conclusion and Discussion

Limitations - Our final aggregated ply shading model assumes that the light scattering enters and exits from the same spot and does not exhibit subsurface scattering. Based on our experiments, this is true for most fabrics except for fibers with very high albedo (such as cotton with 0.999) as they exhibit significantly more bounces per sample and hence travel more throughout the yarn, causing the exit point to be far from the the enter point. While this assumption satisfies most of our cloth types, we left a more accurate distribution

Table 2: Performance Statistics for Fig. 9. All rendering times were counted at equal quality (EQ).

Material	Time (s)			Memory (MB)		
	Ref	[ZMA*23]	Ours	Ref	[ZMA*23]	Ours
fleece	12631	3263	952	6032	20	20
silk	31223	20831	2908	6216	20	20
polyester	8229	11932	612	5935	29	29
cotton	44836	24167	2626	4391	7	7
gabardine	11427	13024	785	9340	20	20

of the exit point as a future study. Furthermore, our model assumes the appearance of the yarn is not spatially varying, and is unable to handle spatially varying yarn colors such as dyed cloth. Lastly, our model requires re-training to alter the yarn parameters, which might limit its use in interactive design and modelling for artists.

Future Works - Besides addressing the limitations above, a straightforward extension can include the training and fitting of more complex fiber distributions and scattering as the neural network has the potential to learn any complex distributions. Additionally, we would like to develop and leverage an auto-encoder architecture similar to [SRRW21] to instantly interpolate our fitted yarn models with different fiber parameters to provide additional flexibility to designers and artists. Although our model performs well with an analytically fitted importance sampling lobe, we are interested in seeing if neural importance sampling methods could be used to further improve convergence [XWH*23]. We also would like to extend our method to support efficient level-of-detail simplification. This involves simplifying our model into a 3-dimensional BCSDf using a smaller neural network for far-field views, specifically when the width of the yarn is less than a pixel.

Conclusions - In this paper, we presented a novel aggregated shading framework by leveraging the flexibility and generality of neural networks to model the light interactions with a bundle of fibers i.e. ply. Our model can replicate the appearance of many fabrics while running significantly faster and requiring less memory. Through observations of the RDM, we also derived an analytical approximation and importance sampling scheme to further improve the rate of convergence of our model. Finally, our fitted model can be applied to any yarn geometry instantly, providing greater flexibility in designing fabrics.

References

- [ACG*17] ALIAGA, CARLOS, CASTILLO, CARLOS, GUTIERREZ, DIEGO, et al. "An appearance model for textile fibers". *Computer Graphics Forum*. Vol. 36. 4. Wiley Online Library. 2017, 35–45 2.
- [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. URL: <https://doi.org/10.1145/1008653.1008663> 2.
- [CNN22] CHEN, ZHE, NOBUHARA, SHOHEI, and NISHINO, KO. "Invertible Neural BRDF for Object Inverse Rendering". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44.12 (2022), 9380–9395. DOI: 10.1109/TPAMI.2021.3129537 2.




Figure 7: Rendering results of a 3-ply glove with fiber parameters; fleece material; fleece with half of the fiber density ρ ; fleece with double twist factor α ; polyester material. Please refer to the supplementary video for a comprehensive overview.




Figure 8: Single-yarn comparisons between our aggregated model and the fiber-based ground truth with varying fiber parameters. In this figure, we compare against both 1-ply and 3-ply yarn and demonstrate that our model can accurately recreate the ply level highlights.

[HZRS15] HE, KAIMING, ZHANG, XIANGYU, REN, SHAOQING, and SUN, JIAN. “Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification”. *2015 IEEE International Conference on Computer Vision (ICCV)*. 2015, 1026–1034. DOI: [10.1109/ICCV.2015.1236](https://doi.org/10.1109/ICCV.2015.1236).

[IM12] IRAWAN, PITI and MARSCHNER, STEVE. “Specular reflection from woven cloth”. *ACM Transactions on Graphics (TOG)* 31.1 (2012), 1–20 [2](#).

[JSR*22] JAKOB, WENZEL, SPEIERER, SÉBASTIEN, ROUSSEL, NICOLAS, et al. *Mitsuba 3 renderer*. Version 3.1.1. <https://mitsuba-renderer.org>. 2022 [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](#).

[KMX*21] KUZNETSOV, ALEXANDR, MULLIA, KRISHNA, XU, ZEXIANG, et al. “NeuMIP: multi-resolution neural materials”. *ACM Transactions on Graphics (TOG)* 40.4 (2021), 1–13 [2](#).

[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](#), [3](#), [7](#), [8](#), [10](#), [11](#).

[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](#).

[KZP*24] KHATTAR, APOORV, ZHU, JUNQUI, PADOVANI, EMILIANO, et al. *A Multi-scale Yarn Appearance Model with Fiber Details*. 2024. arXiv: [2401.12724](https://arxiv.org/abs/2401.12724) [cs.GR] [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.




Figure 9: The SSIM comparison of our aggregated yarn shading model vs explicit fiber-based models [KSZ*15] on knitted and woven fabrics with equal quality. The EQ and ET references are highlighted to showcase their slow and noisy performance, respectively. Our model can accurately replicate the appearance of fiber-based models at a fraction of the time and memory costs.




Figure 10: A comparison of ours with state-of-the-art model [ZMA*23] in equal quality, using reference's parameters for aggregation

Table 3: Fiber parameters used throughout our paper. The shading parameters are based on matched fibers from [KSZ*15], and the geometrical parameters are set on an ad hoc basis

	N	ρ	α	C_R	C_{TT}	β_R	β_{TT}	γ_{TT}
fleece	300	0.30	0.24	0.040, 0.087, 0.087	0.452, 0.725, 0.948	7.238	10.000	25.989
silk	300	0.20	0.00	0.745, 0.008, 0.070	0.620, 0.553, 0.562	1.000	10.000	19.823
polyester	200	0.40	0.20	0.700, 0.700, 0.700	0.600, 0.000, 0.800	5.238	20.000	25.000
cotton	600	0.35	0.06	0.989, 0.959, 0.874	0.999, 0.999, 0.999	1.000	27.197	38.269
gabardine	450	0.25	0.12	0.185, 0.047, 0.069	0.999, 0.330, 0.354	2.141	10.000	23.548

DOI: [10.1111/cgf.13346](https://doi.org/10.1111/cgf.13346). URL: <https://hal.science/hal-01702000/2>.

[LWS*18] LEAF, JONATHAN, WU, RUNDONG, SCHWEICKART, ESTON, et al. "Interactive Design of Periodic Yarn-Level Cloth Patterns". *ACM Trans. Graph.* 37.6 (Dec. 2018). ISSN: 0730-0301. DOI: [10.1145/3272127.3275105](https://doi.org/10.1145/3272127.3275105). URL: <https://doi.org/10.1145/3272127.3275105>.

[LZB17] LUAN, FUJUN, ZHAO, SHUANG, and BALA, KAVITA. "Fiber-Level On-the-Fly Procedural Textiles". *Computer Graphics Forum*. Vol. 36. 4. Wiley Online Library, 2017, 123–135 2.

[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.

[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". *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).

[NH10] NAIR, VINOD and HINTON, GEOFFREY E. "Rectified Linear Units Improve Restricted Boltzmann Machines". *Proceedings of the 27th International Conference on International Conference on Machine Learning*. ICML'10. Haifa, Israel: Omnipress, 2010, 807–814. ISBN: 9781605589077 6.

[RJGW19] RAINER, GILLES, JAKOB, WENZEL, GHOSH, ABHIJEET, and WEYRICH, TIM. "Neural BTF compression and interpolation". *Com-*

- puter Graphics Forum. Vol. 38. 2. Wiley Online Library. 2019, 235–244 2.
- [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.
- [Sch94] SCHLICK, CHRISTOPHE. “An Inexpensive BRDF Model for Physically-based Rendering”. *Computer Graphics Forum* (1994). ISSN: 1467-8659. DOI: [10.1111/1467-8659.13302333](https://doi.org/10.1111/1467-8659.13302333).
- [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.
- [SRRW21] SZTRAJMAN, ALEJANDRO, RAINER, GILLES, RITSCHEL, TOBIAS, and WEYRICH, TIM. “Neural BRDF Representation and Importance Sampling”. *Computer Graphics Forum* (2021). DOI: <https://doi.org/10.1111/cgf.14335>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14335> 2, 5, 6, 8.
- [SSK03] SATTLER, MIRKO, SARLETTE, RALF, and KLEIN, REINHARD. “Efficient and realistic visualization of cloth”. *Rendering techniques*. 2003, 167–178 2.
- [WY17] WU, KUI and YUKSEL, CEM. “Real-time Fiber-level Cloth Rendering”. *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D 2017)*. San Francisco, CA: ACM, 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> 2, 8.
- [XWH*23] XU, BING, WU, LIWEN, HASAN, MILOS, et al. “NeuSample: Importance Sampling for Neural Materials”. *ACM SIGGRAPH 2023 Conference Proceedings*. 2023, 1–10 8.
- [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> 8.
- [Yuk20] YUKSEL, CEM. *Yarn-level Cloth Models*. <http://www.cemyuksel.com/research/yarnmodels>. 2020 8.
- [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.
- [ZJA*23] ZHU, JUNQIU, JARABO, ADRIAN, ALIAGA, CARLOS, et al. “A Realistic Surface-based Cloth Rendering Model”. *ACM SIGGRAPH 2023 Conference Proceedings*. 2023, 1–9 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.
- [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 8.
- [ZMA*23] ZHU, JUNQIU, MONTAZERI, ZAHRA, AUBRY, JEAN-MARIE, et al. “A Practical and Hierarchical Yarn-based Shading Model for Cloth”. *Computer Graphics Forum* (2023). ISSN: 1467-8659. DOI: [10.1111/cgf.14894](https://doi.org/10.1111/cgf.14894) 2, 7, 8, 11.
- [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> 2, 3, 5.

Appendix A: Derivation of Rotation Angle.

Borrowing concepts from rotational dynamics, the tangential velocity, v , on the surface of the yarn can be defined as:

$$v = R\omega \quad (13)$$

where R represents the radius of the yarn and ω denotes the rotational velocity. The rotational velocity can be expressed as:

$$\omega = \frac{\Delta\theta}{\Delta l} \quad (14)$$

where $\Delta\theta$ is the change in angular position and Δl is change in position along the yarn tangent. From Equation 1 and 14, we can obtain:

$$\alpha = \frac{2R\Delta n}{\Delta l} = \frac{R\Delta\theta}{\pi\Delta l} = \frac{R}{\pi}\omega \quad (15)$$

Substituting the previous equation with 13, we obtain:

$$v = \pi\alpha \quad (16)$$

Given that v represents the rate of change of the bi-tangent with respect to the yarn tangent, and considering that they are perpendicular to each other, forming a right-angle triangle, we deduce:

$$\begin{aligned} \tan \tilde{\phi} &= v \\ \tilde{\phi} &= \tan^{-1}(\pi\alpha) \end{aligned} \quad (17)$$