A Practical Ply-Based Appearance Model of Woven Fabrics

ZAHRA MONTAZERI, Luxion, Inc. and University of California, Irvine
SØREN B. GAMMELMARK, Luxion, Inc.
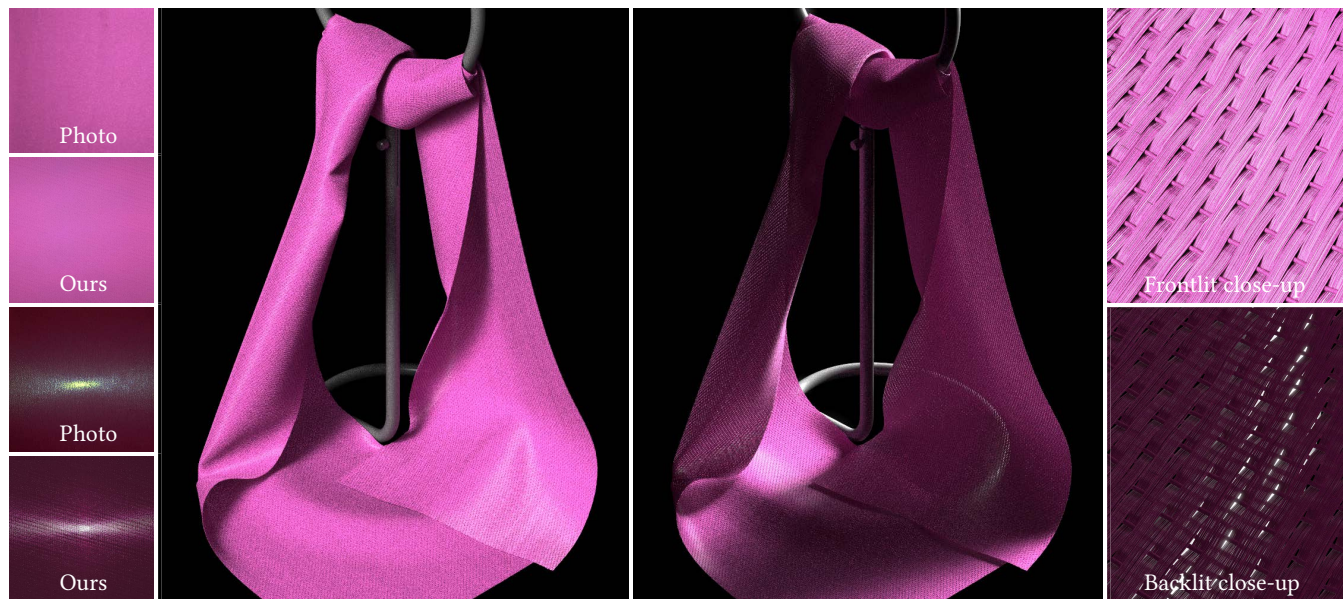SHUANG ZHAO, University of California, Irvine
HENRIK WANN JENSEN, Luxion, Inc.

Fig. 1. Our ply-based appearance model is able to simulate both reflected and transmitted light through complex weave patterns as seen on the left and our rendered results match with photographs of a real sample. Furthermore, it has a compact representation that makes it efficient to use for large pieces of clothing such as the scarf shown on the right. In addition to the anisotropic highlights, the color due to front and back lighting changes considerably and our model reproduces the color shift seen in the photographs. The images on the right are close-ups of the scarf showing the individual yarns. Please see the supplementary video for the full light rotation.

Simulating the appearance of woven fabrics is challenging due to the complex interplay of lighting between the constituent yarns and fibers. Conventional surface-based models lack the fidelity and details for producing realistic close-up renderings. Micro-appearance models, on the other hand, can produce highly detailed renderings by depicting fabrics fiber-by-fiber, but become expensive when handling large pieces of clothing. Further, neither surface-based nor micro-appearance model has not been shown in practice to match measurements of complex anisotropic reflection and transmission simultaneously.

In this paper, we introduce a practical appearance model for woven fabrics. We model the structure of a fabric at the ply level and simulate the local appearance of fibers making up each ply. Our model accounts for both reflection and transmission of light and is capable of matching physical measurements better than prior methods including fiber based techniques. Compared to existing micro-appearance models, our model is light-weight and scales to large pieces of clothing.

CCS Concepts: • **Computing methodologies → Rendering**.

Additional Key Words and Phrases: cloth modeling, cloth rendering

Authors' addresses: Zahra Montazeri, Luxion, Inc. and University of California, Irvine, zmontaze@uci.edu; Søren B. Gammelmark, Luxion, Inc., soren@luxion.com; Shuang Zhao, University of California, Irvine, shz@ics.uci.edu; Henrik Wann Jensen, Luxion, Inc., henrik@luxion.com.

1 INTRODUCTION

Fabrics are essential in our daily lives and used for clothing as well as many functional cloths from window curtains to table-sheets. Thus, modeling the appearance of fabrics in a physically faithful

fashion has applications in many areas such as design, online retail, and entertainment.

Modeling cloth appearance, an active research topics in computer graphics for decades, generally boils down to representing the geometry and optics of a fabric. Existing models can be broadly classified into two categories. *Surface-based models* depict cloth geometry using smooth 2D sheets (expressed as polygonal meshes) and optics using specialized reflectance models (e.g., [Adabala et al. 2003; Irawan and Marschner 2012; Sadeghi et al. 2013]). *Micro-appearance models*, on the other hand, express cloth geometry at the microscale down to individual micron-diameter fibers (using high-resolution volumes [Zhao et al. 2011, 2012] or fiber meshes [Khungurn et al. 2016]) and optics by utilizing volumetric [Zhao et al. 2011] or fiber-based [Khungurn et al. 2016; Schroder et al. 2015] light scattering models.

In practice, both categories have their own advantages and disadvantages. The surface-based models are typically light-weight, easy to edit, and can faithfully reproduce the appearance of a fabric at the macroscale. Thus, they have been widely used in the computer graphics industry [Google 2019; McAuley et al. 2012]. However, these models lack the fine-grained details that are crucial for cloth rendering under close-up views. The micro-appearance models, in contrast, are capable of generating renderings with remarkably high fidelity and details. Unfortunately, as these models are highly data-intensive and usually difficult to manipulate, their practical use has been quite limited. Furthermore, both the surface-based and the micro-appearance techniques have mostly focused on modeling the reflection of incident light. Validation against measurements of the transmission of light, on the other hand, has been largely neglected at the macroscale.

There are many reasons why previous models have been unsuccessful in capturing both reflection and transmission properties of textiles. In the quest to obtain very high fidelity renderings micro-appearance models relies on a very detailed representation of the textile. This is very useful for understanding how different aspects of the micro-geometry affect the overall appearance, but makes the models inflexible for dealing with textiles that do not fall within the idealized micro-properties of these models. Aliaga et al. [2017] for example, demonstrate a wide range of optical properties for different textile fibers.

In the quest for more realistic cloth appearance models the addition of highly detailed fiber geometry can be seen as a form of reductionism [Anderson 1972]. While these models with increasingly detailed geometry offer incredible detail they also require more complex controls (fiber arrangement and shape as well as local fiber shading models). This added complexity makes the models difficult to control and edit, and previous work [Khungurn et al. 2016] as well as our results show that fiber based models cannot easily match both reflected and transmitted lighting in measured cloth samples. In this paper we argue that an aggregate model offers a compromise where the appearance is easier to manipulate, while retaining most of the details from the fiber level.

Concretely, our main contributions include:

- A new geometric model (§3) that depicts a fabric as a collection of *plies* that are further enriched using procedurally generated fiber details. By working at the ply level, this representation offers a good balance between model complexity and fidelity.

- A new light scattering model (§4) that expresses how light interacts with a ply by capturing how, on average, light scatters off the constituent fibers. Thanks to this aggregate formulation, our cloth models can be rendered much more efficiently compared to the previous micro-appearance models. Additionally, by decoupling reflection and transmission properties, our light scattering model offers the flexibility to accurately reproduce both effects.

To demonstrate the practical usefulness of our technique, we have integrated it into an industrial rendering system that provides user-friendly material parameter tweaking.

2 RELATED WORK

Surface-based cloth models. Traditionally, cloth has been modeled as thin 2D surfaces with light-cloth interactions depicted using bidirectional reflectance distribution functions (BRDFs) (e.g., [Adabala et al. 2003; Irawan and Marschner 2012; Sadeghi et al. 2013]). These models are light-weight and can produce high-quality results. However, they usually lack fine-grained details such as thread and fiber structures necessary for close-up views.

Micro-appearance cloth models. A number of *micro-appearance models* have been developed in the recent years that explicitly describe a fabric's micro-geometry down to the fiber level and have offered a new level of fidelity and details [Khungurn et al. 2016; Loubet and Neyret 2018; Montazeri et al. 2019; Schröder et al. 2011; Schroder et al. 2015; Zhao et al. 2011, 2016]. Due to the high complexity of these models, rendering them efficiently is challenging. To address this problem, several precomputation-based rendering methods have been developed. Zhao et al. [Zhao et al. 2013] proposed a technique that exploits the modularity of woven cloth models by efficiently combining precomputed light-transport information on the fly. Wu and Yuksel [2017] developed a GPU-based method for procedurally described fabrics by leveraging 'core fibers' with precomputed shadowing information. Khungurn et al. [2017] extended precomputed radiance transfer [Sloan et al. 2002] to handle micro-appearance cloth models. Unfortunately, these methods either lack performance [Khungurn et al. 2017; Zhao et al. 2013] or physical accuracy [Wu and Yuksel 2017] for predictive rendering and editing micro-appearance cloth models. Furthermore, even though these models are very detailed they have not demonstrated the ability to simulate both reflection and transmission through cloth accurately.

Data-driven models. Cloth appearance can also be modeled in data-driven ways, such as using Bidirectional Texture Functions (BTFs) [Dana et al. 1999]. However, these representations are cumbersome and difficult to manipulate. Further, they are normally too data-intensive to capture the level of details offered by micro-appearance models.

Fiber scattering models. The reflectance profile used by our technique is related to various fiber [Aliaga et al. 2017; Khungurn et al. 2016; Marschner et al. 2003; Zinke and Weber 2007] and fur [Yan et al. 2017, 2015] scattering models. Additionally, there are several multi-scattering studies for hair rendering [Moon and Marschner

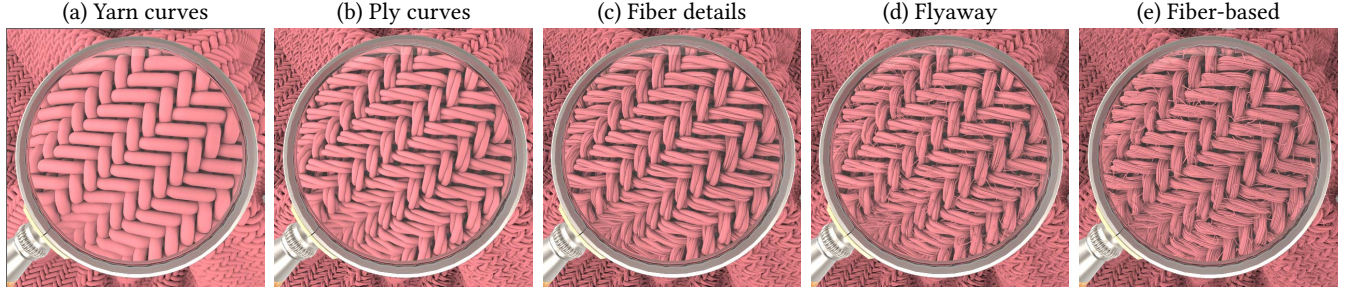| (a) Yarn curves	(b) Ply curves	(c) Fiber details	(d) Flyaway	(e) Fiber-based

Fig. 2. **Cloth geometry generation.** Given a weave pattern and a base mesh as the inputs to our system (a) Optimization of the yarn curves to simulate the 3D yarn structure. (b) Plies are generated as helix curves around each yarn curve. (c) Fiber details are added as normal/tangent map as well as fiber shadowing. (d) Flyaway fibers representing fibers breaking away from the yarns are generated. (e) Matching the geometry of fiber-based model (with 240 fibers per yarn) given the same procedural parameters.

2006; Moon et al. 2008; Zinke et al. 2008]. These models utilize bidirectional curve scattering distribution functions (BCSDFs) to describe how light scatters off individual fibers or fur strands.

Granular materials models. Similar to our aggregation ply BSDF, granular materials are composed of many dielectric grains where their aggregate behaviour determines the overall appearance [Meng et al. 2015; Moon et al. 2007; Müller et al. 2016]. They study the decomposition of a complex modeling/appearance problem to decouple individual grains from their arrangement to achieve high performance.

Micro-geometry models. Rendering glints on micro-geometry surfaces liked brushed metal shares some of the complexity seen with bundles of fibers in cloth rendering [Iwasaki et al. 2012; Kuznetsov et al. 2019; Raymond et al. 2016; Wu et al. 2011, 2013]. These methods have not been used for the fiber structures seen in clothing, but the ideas could likely be extended to increase the efficiency of highly detailed cloth models.

3 MODELING CLOTH GEOMETRY

We now describe how our model depicts detailed geometry of woven cloth. Unlike most prior models that either lack sub-yarn-level information [Irawan and Marschner 2012; Sadeghi et al. 2013] or rely on explicit fiber structures [Khungurn et al. 2016; Zhao et al. 2011, 2016], our expression offers a good balance between performance and fidelity by mainly working at the ply level.

Preliminaries. The geometric structure of a textile is complex at many scales. Fabrics are constructed from multiple *yarns* via manufacturing technologies like weaving and knitting. Each yarn, in turn, consists of few substrand bundles, or *plies*, that are created by twisting hundreds of micron-diameter fibers. The variations of these fiber details can impact the overall appearance [Irawan and Marschner 2012; Khungurn et al. 2016; Zhao et al. 2011].

We consider two types of yarns, *round* and *flat*, as illustrated in the adjacent figure. Round yarns have cylindrical shapes and are
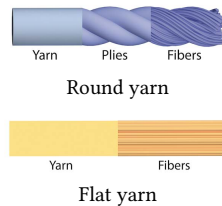
Table 1. A list of important symbols.

	P	Weave pattern
	Y	Yarn curve
	s	Parameter along yarn curve
	s_k	Piecewise yarn polynomial knot locations
Geometry (§3)	α_k, h_k	Yarn polynomial control points
	S_i	Length of yarn period
	h	Yarn height polynomial
	o	Yarn offset
	$\hat{\mathbf{t}}, \hat{\mathbf{b}}, \hat{\mathbf{w}}$	Yarn direction, offset and height basis
	k, k_c	Yarn strain and collision coefficient
	E	Yarn and weave pattern energy
	E_C	Yarn collision energy
	f	BSDF
	$f_{\boldsymbol{x}}^{\mathrm{S}}$	Specular reflection at $\boldsymbol{x}$
	$f_{\boldsymbol{y}}^{\mathrm{S}}$	Specular transmission at $\boldsymbol{y}$
	$f_{\boldsymbol{x}}^{\mathrm{B}}$	Body reflection at $\boldsymbol{x}$
	$f_{\boldsymbol{y}}^{\mathrm{B}}$	Body transmission at $\boldsymbol{y}$
	$\boldsymbol{x}, \boldsymbol{y}$	entrance and exit point
	τ	Volumetric transmittance of light
	σ_{t}	Material extinction coefficient (optical density)
	η	Material refractive index
Appearance (§4)	$\omega_{\mathrm{i}}, \omega_{\mathrm{o}}$	Incoming and outgoing light direction
	ω_{h}	Half-way vector (normalized)
	$\boldsymbol{n}$	Surface normal
	$\mathrm{G_f}$	Fiber-shadowing function
	G_2	Smith non-correlated masking-shadowing function
	F^{S}	Reflection fresnel term
	D	Microfacet Normal distribution function (NDF)
	β	Anisostropic roughness
	k^{S}	Specular tint
	k^{B}	Base color
	α	Reduced single-scattering albedo convertor
	B	Body energy
	r	Body reflected coefficient
	h	Thickness of virtual medium for flat yarns

Round yarn

Flat yarn

very common in woven fabrics. Flat yarns, which are usually used in light and glossy fabrics such as silk satin, are ribbon-like and consist of well aligned fibers.

Overview. We take as input a polygonal mesh (with per-vertex UV coordinates), which captures a fabric's macro-geometry, and a 2D weave pattern that specifies the arrangement of the yarns. Using a simple physical model, we compute a tilable set of yarn centerline curves (§3.1) that can be mapped to the input geometry in a manner

analogous to displacement mapping. The plies are constructed as helices mapped from a yarn centerline space to fabric space similar to previous procedural techniques [Luan et al. 2017; Zhao et al. 2016] (§3.2). Finally, fiber-level details are added in a procedural texture-based fashion on individual plies (§3.3). We summarize this pipeline in Figure 2.

3.1 Yarn Centerlines

One strategy for obtaining yarn centerlines is to ensure the curves pass over and under each-other as dictated by the weaving pattern. While this may give reasonably looking results the resulting distribution of slopes will not match those computed by more complex simulations [Leaf et al. 2019] nor those found by investigating the specular highlights [Sadeghi et al. 2013]. In the following, we describe a physical model that can reproduce the slope distributions of yarns reported in [Leaf et al. 2019; Sadeghi et al. 2013] and at the same time produce tileable curves.

Relaxation of yarn curves. As [Leaf et al. 2019], we assume that the textile is flat compared to the size of the yarns and perform yarn relaxation in flat space. In this space, the yarns in the pattern run along two primary orthogonal axes: the warp (x) and weft (y) directions. Each yarn centerline is a curve along one of these primary directions with the displacement (z) away from the tiling plane as the unknown quantity we need to compute. A weave pattern P with N warps and M wefts are placed in the rectangle $[0, N] \times [0, M]$ and the yarn offsets at half-integer coordinates. For example, the first warp curve is placed at $y = 0.5$ and runs along the x-axis from $[0, M]$. Further, we assume this pattern to be tiled over the entire plane, yielding periodic boundary conditions as detailed below.

In general, a yarn Y_i has a primary direction $\hat{\mathbf{t}}_i$ and a transverse direction $\hat{\mathbf{b}}_i$: for warps, $\hat{\mathbf{t}}_i$ is the x-direction and $\hat{\mathbf{b}}_i$ is the y-direction; for wefts, the opposite holds. Under our formulation, a yarn curve can be expressed as

$$Y_i(s) = s\,\hat{\mathbf{t}}_i + o_i\,\hat{\mathbf{b}}_i + h_i(s)\,\hat{\mathbf{w}}, \tag{1}$$

where o_i is the half-integer offset of the yarn centerline and h_i is a function describing the deviation of the centerline from a straight line and $\hat{\mathbf{w}}$ is the direction orthogonal to both weave directions. Each yarn satisfy the periodic boundary conditions $h_i(s + S_i) = h_i(s)$ and $h'_i(s + S_i) = h'_i(s)$ where $S_i = M$ for warps and $S_i = N$ for wefts. This allows tiling the cell $[0, N] \times [0, M]$ with C^1-continuity.

For $h_i(s)$ we choose a piece-wise cubic polynomial with Q_i knots equidistantly placed between $s = 0$ and the period length ($S_i = N$ or $S_i = M$). At each knot s_k we specify the height $h_{i,k}$ and the slope $\alpha_{i,k}$ which uniquely defines the cubic polynomial within each segment[1] $[s_k, s_{k+1}[$. Note that the periodic boundary conditions imply that $h_{i,0} = h_{i,Q_i-1}$ and $\alpha_{i,0} = \alpha_{i,K_i-1}$ for all yarns Y_i.

Self-energy of a single yarn. We opt for a simple strain-energy function of the deformed yarn given by

$$E[Y_i] = \frac{k}{2} \int_0^{S_i} \left(\frac{\partial h_i(s)}{\partial s}\right)^2 \, \mathrm{d}s, \tag{2}$$

---

[1]We use '$[s_k, s_{k+1}[$' to denote an interval that is closed on s_k and open on s_{k+1}.

where k is the yarn stiffness coefficient. Since h_i is a piece-wise cubic polynomial the energy $E[Y_i]$ and its derivatives can be expressed analytically in terms of the control points $\{h_{i,j}, \alpha_{i,j}\}$ which is detailed in Appendix A.

Collision energy. To handle yarn-yarn collisions, we select the density of knots such that there is at least one knot at each yarn crossing and one knot in between each crossing. Let P_C be the set of pairs of knots located at yarn crossing. For example if $(A, B) \in P_C$ denotes a yarn crossing index pair $A = (i, k)$, $B = (j, l)$ then yarn i and j intersect at control points k and l respectively.

We model the collision energy of these two control points as soft spheres. We denote the effective collision radii of these spheres by R_A and R_B which for round yarns is the radius of the yarns, but for flat yarns is reduced to allow the flat yarns to approach each other more closely. The collision energy is then given by

$$E_C(A, B) = k_c\, f\left[\left(\frac{h_A - h_B}{R_A + R_B}\right)^2\right], \tag{3}$$

where k_c controls the magnitude of the collision energy and f is a collision potential [Kaldor et al. 2008] that equals $d^{-2} + d^2 - 2$ if $d < 1$ and is zero otherwise.

Energy of a woven fabric. Adding the strain- and collision-energies up for all yarns gives the full energy functional for the pattern

$$E = \sum_{i \in P} \frac{k_i}{2} \int_0^{S_i} \left(\frac{\partial h_i(s)}{\partial s}\right)^2 \, \mathrm{d}s + \sum_{A,B \in P_C} E_C(A, B), \tag{4}$$

where we have allowed the stiffness coefficient to depend on the yarn i. Using a good non-linear optimization technique such as L-BFGS [Fletcher 1987], this energy can be minimized at interactive speeds for small weaving patches even for a single-threaded CPU computation. As noted previously, yarn centerlines obtained this way matches the slopes reported by [Leaf et al. 2019; Sadeghi et al. 2013] well and can therefore also reproduce specular highlights at the same locations.

With the yarn centerlines computed in a unit-cell, we can map the curves to a UV-mapped mesh. Let $p : \mathbb{R}^2 \mapsto S$ be the mapping from the mesh UV space to the surface and $\boldsymbol{n} : S \mapsto S^2$ be the normal field of the surface. Then each point (x, y, z) in yarn-space can be mapped to the mesh space by computing

$$\phi(x, y, z) = p(x, y) + H z\, \boldsymbol{n}(p(x, y)), \tag{5}$$

where H is a mesh- and uv-mapping dependent scale factor to ensure ϕ is approximately isometric. In other words, we displace the yarn height along the normal of the mesh. In general, each point of the surface has a frame defined by the UV-tangent plane basis and the normal. This frame allows us to map any vector-field in weave-space to mesh space, allowing a natural transformation of yarn tangent and normal information to mesh space.

3.2 Ply-Level Geometry

Recall that a real yarn can consist of multiple twisted plies. After obtaining the yarn centerlines (§3.1), we generate the ply curves following the individual yarn curves. The plies add both geometric
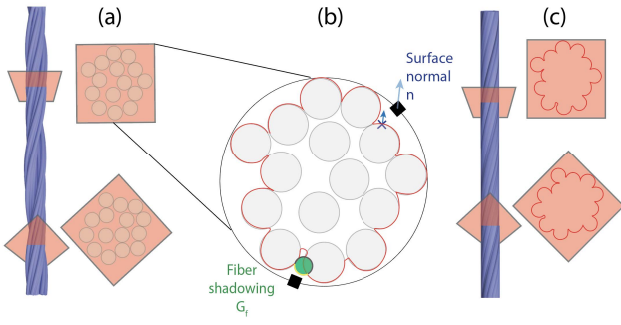
Fig. 3. **Cross section processing**. (a) We generate the cross section of a ply using a procedural model. (b) Within this cross section we perform 2D ray tracing to detect fibers visible from the exterior. We find the normals for these visible fibers using the same ray tracing process, and tangents are computed using derivative of the corresponding helix in 3D. Furthermore, local fiber shadowing is calculated for the visible fibers using 2D ray tracing. (c) Using the fiber normals, tangents, and shadowing textures we can shade the ply curve and account for shading details due to the individual fibers. Lastly, we add fiber irregularities as stated in §3.3.

and shading complexity that greatly increases the realism compared to just using a single yarn.

Similar to previous works [Max 1990; Schroder et al. 2015; Zhao et al. 2016], we model ply curves as helices around each yarn curve with user-controllable parameters (e.g., ply count and pitch). This is shown in Figure 2-c. The ply-curves are subsequently rendered as piecewise cylinders. Flat yarns are not constructed from plies and are instead rendered as bilinear patches.

3.3 Fiber-Level Details

Our model enriches the ply surfaces by adding procedurally generated fiber-level details with normal, tangent, and shadowing maps. To this end, we utilize a procedural yarn model [Schroder et al. 2015; Zhao et al. 2016] to drive the generation of these maps.

Specifically, we generate the cross section of a ply using the procedural model. Within this cross section, each fiber is a small 2D disc (with its location and radius specified by the model). Provided the cross-sectional arrangement of fibers, we perform ray tracing within the cross section to detect *visible fibers*—those visible from the exterior (see Figure 3-b).

Further, we extract the normal information of all visible fibers via the same ray tracing process and store the results as a 1D texture (see Figure 3-b). Tangent information of each visible fiber is calculated using the derivative of corresponding helix in 3D using the fiber distance from the ply center. Lastly, we compute local fiber shadowing, as shown in 3-b, that approximates the occlusion from neighboring fibers. Specifically, our inter-fiber shadowing term, which is a scalar ranging from 0 to 1, captures the ratio of incident directions that are not occluded. For fast computation, we consider the fibers to be locally straight and approximate the

$$G_f(\boldsymbol{x}) = \ell_2/\pi$$
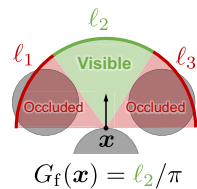
shadowing term by computing within ply cross sections as the ratio of the angle of the visible arc to 2π (see the adjacent figure). This gives the exact result for a point on the bottom of an infinite V-groove and is otherwise an approximation.

Introducing irregularities. For round plies the fibers are circular helices around the ply centerlines, the cross-sectional normal and tangent information is almost invariant along a ply up to a global rotation (see Figure 3-a). This

Without migration

With migration

allows us to use a 1D texture, which adds minimal computational and storage overhead. However, the use of rotationally invariant ply cross sections can lead to results that suffer from visually disturbing periodic patterns. We tackle this challenge by switching randomly among a few precomputed fiber arrangements that are distributed randomly along the ply. We interpolate between two 1D fiber arrangements using a smoothstep function. This effectively mimics the migration of fibers (see the adjacent figure).
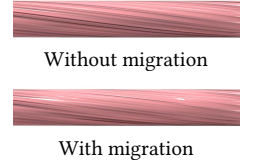
Flyaway fibers. Flyaway fibers are those that fray out from the textile. We use a technique similar to Zhao et al. [2011] by randomly modeling these fibers as separate pieces of geometry depicted as B-Spline curve (with user-controlled densities). Each flyaway is generated at a random location within a given ply. The initial direction of the flyaway is a randomly chosen angle away from the ply direction. The maximum angle and length are user controlled parameters. Subsequent flyaway segments are generated as continuations of the previous flyaway segment with a new random angle to create a flyaway curve with irregularities.

4 MODELING CLOTH APPEARANCE

To simulate how light is scattered by the generated ply-level geometry (§3) we build on recent work on appearance modeling of hair and fur [d'Eon et al. 2011; Yan et al. 2017; Zinke and Weber 2007]. One major difference to hair and fur rendering is that we compute light scattering with an aggregate bundle of cloth fibers rather than considering how light is scattered by each individual fiber.

Overview. Our model works at the ply level with decoupled reflection and transmission components. This decoupling offers great flexibility and control than previous fiber-level micro-appearance models (this is demonstrated in Figure 9). We use a near field bidirectional scattering distribution function (BSDF) to simulate light interaction at a specific location on a ply.

At a high level, our model depicts how light is scattered (i.e., reflected and refracted) at two surface points jointly (see Figure 4). The first point $\boldsymbol{x}$ is the intersection between an incident light ray and a ply surface, which arises from standard path tracing. The second point, $\boldsymbol{y}$, is where a (randomly sampled) refracted light ray originating at $\boldsymbol{x}$ intersects the ply surface (from the interior). At each point, our model utilizes a specialized BSDF that combines a *specular* and a *body* component. The specular term (denoted using a superscript S) captures (anisotropic) specular reflection and transmission, and the body term (denoted with a superscript B) approximates multiple scattering within the bundle of fibers.
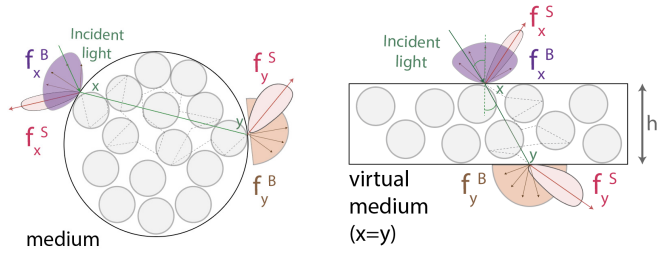
Fig. 4. **Illustration of our Ply BSDF model for round (left) and flat (right) plies.** At entrance point $\boldsymbol{x}$ and exit point $\boldsymbol{y}$, our model utilizes a specialized BSDF comprised of a *specular* and a *body* component where the former (denoted using a superscript S) captures (anisotropic) specular reflection and transmission, and the latter (denoted with a superscript B) approximates multiple-scattered light.
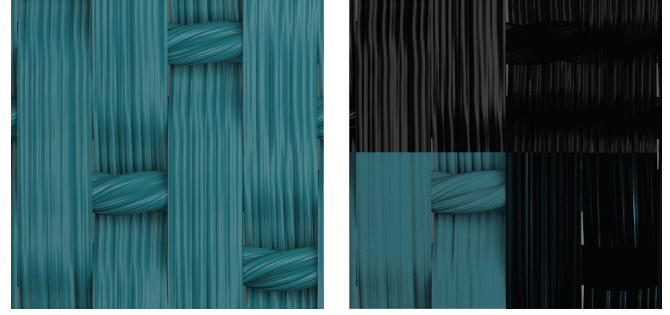

Full   Components

Fig. 5. **Components of our Ply-BSDF:** specular reflection (top left), specular transmission (top right), body reflection (bottom left), body transmission (bottom right)

When a light ray intersects a ply from the exterior at $\boldsymbol{x}$, we use the following BSDF:

$$f_{\boldsymbol{x}} = f_{\boldsymbol{x}}^{S} + f_{\boldsymbol{x}}^{B}. \tag{6}$$

Subsequently, when the light ray follows the transmission component of $f_{\boldsymbol{x}}^{S}$ and intersects the ply from the interior at $\boldsymbol{y}$, we use another BSDF:

$$f_{\boldsymbol{y}} = f_{\boldsymbol{y}}^{S} + f_{\boldsymbol{y}}^{B}. \tag{7}$$

In the following, we define in detail the specular and body components at both points.

4.1 Specular Components

Fabrics have evident specular and directional highlights due to the immediate reflection as well as light transmission from the other side of the cloth. We use the rough dielectric BSDF [Heitz 2018; Walter et al. 2007] to model the specular components. The specular component of the BSDF at an entrance point $\boldsymbol{x}$ can be written as:

$$f_{\boldsymbol{x}}^{S}(\boldsymbol{\omega}_{i}, \boldsymbol{\omega}_{o}) = k_{\boldsymbol{x}}^{S} \cdot \tag{8}$$

$$\begin{cases} \dfrac{F_{\boldsymbol{x}}\, G_2(\boldsymbol{\omega}_i,\boldsymbol{\omega}_o,\boldsymbol{\omega}_h)\, D(\boldsymbol{\omega}_h;\boldsymbol{\beta}_{\boldsymbol{x}})}{4\,|\boldsymbol{\omega}_i\cdot\boldsymbol{n}(\boldsymbol{x})|\,|\boldsymbol{\omega}_o\cdot\boldsymbol{n}(\boldsymbol{x})|}, & (\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_o > 0) \\[2ex] \dfrac{|\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_t|\,|\boldsymbol{\omega}_o\cdot\boldsymbol{\omega}_t|}{|\boldsymbol{\omega}_i\cdot\boldsymbol{n}(\boldsymbol{x})|\,|\boldsymbol{\omega}_o\cdot\boldsymbol{n}(\boldsymbol{x})|}\dfrac{\eta^2\,(1-F_{\boldsymbol{x}})\,G_2(\boldsymbol{\omega}_i,\boldsymbol{\omega}_o,\boldsymbol{\omega}_t)\,D(\boldsymbol{\omega}_t;\boldsymbol{\beta}_{\boldsymbol{x}})}{[(\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_t)+\eta\,(\boldsymbol{\omega}_o\cdot\boldsymbol{\omega}_t)]^2}, & (\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_o < 0) \end{cases}$$

where $\boldsymbol{n}(\boldsymbol{x})$ is the surface normal at $\boldsymbol{x}$, $F_{\boldsymbol{x}}$ indicates the Fresnel reflection coefficient, η is the material's refractive index,[2] $\boldsymbol{\omega}_h$ is the normalized half vector given by

$$\boldsymbol{\omega}_h = (\boldsymbol{\omega}_i+\boldsymbol{\omega}_o)/\|\boldsymbol{\omega}_i+\boldsymbol{\omega}_o\|, \qquad \boldsymbol{\omega}_t = (-\boldsymbol{\omega}_i-\eta\boldsymbol{\omega}_o)/\|-\boldsymbol{\omega}_i-\eta\boldsymbol{\omega}_o\|, \tag{9}$$

$D(\cdot;\boldsymbol{\beta})$ denotes the normal distribution function (NDF) where we use GGX distribution with width parameter $\boldsymbol{\beta} \in \mathbb{R}^2$ that captures the (anisotropic) roughness of a surface, and G_2 is the Smith uncorrelated masking-shadowing function. Additionally, $k_{\boldsymbol{x}}^{S}$ in Eq. (8) is a user-controlled specular albedo.

Similarly, at the exit point $\boldsymbol{y}$, for all $\boldsymbol{\omega}_o \cdot \boldsymbol{n}(\boldsymbol{y}) > 0$, the specular component uses the transmitted component of the rough dielectric

model:

$$f_{\boldsymbol{y}}^{S}(\boldsymbol{\omega}_{i}, \boldsymbol{\omega}_{o}) = k_{\boldsymbol{y}}^{S} \cdot \tau(\boldsymbol{x}, \boldsymbol{y}) \cdot \tag{10}$$

$$\begin{cases} 0, & (\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_o > 0) \\[2ex] \dfrac{|\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_t|\,|\boldsymbol{\omega}_o\cdot\boldsymbol{\omega}_t|}{|\boldsymbol{\omega}_i\cdot\boldsymbol{n}(\boldsymbol{y})|\,|\boldsymbol{\omega}_o\cdot\boldsymbol{n}(\boldsymbol{y})|}\dfrac{(1-F_{\boldsymbol{y}})\,G_2(\boldsymbol{\omega}_i,\boldsymbol{\omega}_o,\boldsymbol{\omega}_t)\,D(\boldsymbol{\omega}_t;\boldsymbol{\beta}_{\boldsymbol{y}})}{[\eta\,(\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_t)+(\boldsymbol{\omega}_o\cdot\boldsymbol{\omega}_t)]^2}, & (\boldsymbol{\omega}_i\cdot\boldsymbol{\omega}_o < 0) \end{cases}$$

where $\boldsymbol{\omega}_t = (-\eta\boldsymbol{\omega}_i-\boldsymbol{\omega}_o)/\|-\eta\boldsymbol{\omega}_i-\boldsymbol{\omega}_o\|$, $F_{\boldsymbol{y}}$ indicates the Fresnel reflection coefficient, and $k_{\boldsymbol{y}}^{S}$ is a user controlled specular transmission color, $\tau(\boldsymbol{x}, \boldsymbol{y})$ accounts for the transmission of light through the fiber volume between $\boldsymbol{x}$ and $\boldsymbol{y}$:

$$\tau(\boldsymbol{x}, \boldsymbol{y}) := \exp(-\sigma_t\,\|\boldsymbol{x} - \boldsymbol{y}\|), \tag{11}$$

where σ_t is the material's extinction coefficient.

In Eq. (8), we use slightly higher surface roughness $\boldsymbol{\beta}_{\boldsymbol{y}}$ to allow the transmission lobe (and the subsequent exiting location $\boldsymbol{y}$) to have a wider spread, approximating multiple internal reflections. This is similar to previous work on hair rendering [Marschner et al. 2003].

4.2 Body Components

We approximate multiple scattering of light within a ply with body reflection and transmission terms $f_{\boldsymbol{x}}^{B}$ and $f_{\boldsymbol{y}}^{B}$ in Eqs. (6) and (7). We represent both terms using diffuse-like distributions.[3]

We model the body of the ply as a homogeneous isotropic medium and approximate the scattering within using a BSDF approximation. This resulting BSDF body reflection term $f_{\boldsymbol{x}}^{B}$ at the entrance point $\boldsymbol{x}$ is then comprised of two terms: The first term is a Lommel-Seeliger (LS) term that approximates single-scattered light within the ply medium [Hanrahan and Krueger 1993; Jensen et al. 2001] and the second term approximates multiple scattering within the ply (that is, light enters the ply from $\boldsymbol{x}$, scattered many times internally before reemerging at $\boldsymbol{x}$); We observe that the introduction of the LS term to our reflection component is important for reproducing the smooth appearance of real-world plies.

---

[2]We have assumed the external refractive index to be one.

[3]Although this is a rough approximation to the true single- and multiple-scattering phenomenon, we find it adequate in practice since cloth fibers are usually disorganized, causing light to diffuse quickly inside a ply. A similar approach was used by Yan et al. [2017] to account for scattering by the medulla in fur.

$$f_{\boldsymbol{x}}^{\mathrm{B}}(\boldsymbol{\omega}_{\mathrm{i}}, \boldsymbol{\omega}_{\mathrm{o}}) = r\, B(\boldsymbol{\omega}_{\mathrm{i}}, \boldsymbol{\omega}_{\mathrm{o}})\, \mathrm{G}_{\mathrm{f}}(\boldsymbol{x}) \Bigg[\frac{\alpha}{4\pi(|\boldsymbol{\omega}_{\mathrm{i}} \cdot \boldsymbol{n}(\boldsymbol{x})| + |\boldsymbol{\omega}_{\mathrm{i}} \cdot \boldsymbol{n}(\boldsymbol{x})|)} +$$
$$\frac{\alpha}{2\pi}\left(1 + \mathrm{e}^{-\frac{4}{3}A\sqrt{3(1-\alpha)}}\right)\mathrm{e}^{-\sqrt{3(1-\alpha)}} \Bigg], \quad (12)$$

where α is the (reduced) single-scattering albedo, which is computed from a user controlled diffuse base color, $k_{\boldsymbol{x}}^{\mathrm{B}}$, as described by Jensen and Buhler [2002]. $\mathrm{G}_{\mathrm{f}}(\boldsymbol{x})$ is the fiber-shadowing term discussed in §3.3 and $r \in [0, 1]$ is the portion of the *total body scattering energy* B that exits at point $\boldsymbol{x}$; assuming that all light is scattered within the fibers this value approximates forward and backward scattering. For details on the other terms in the equation please see the work by Jensen et al. [2001].

Multiple scattered light, $f_{\boldsymbol{y}}^{\mathrm{B}}$, that exits at the backside, y, is modeled using a simple Lambertian term:

$$f_{\boldsymbol{y}}^{\mathrm{B}}(\boldsymbol{\omega}_{\mathrm{i}}, \boldsymbol{\omega}_{\mathrm{o}}) = (1 - r)\, B\, \mathrm{G}_{\mathrm{f}}(\boldsymbol{y})\, \frac{k_{\boldsymbol{y}}}{\pi}. \quad (13)$$

where, $(1 - r)$ is the transmitted portion of body energy B (whose dependency on incident and outgoing directions at $\boldsymbol{x}$ are omitted for notational convenience), and $k_{\boldsymbol{y}}$ and $\mathrm{G}_{\mathrm{f}}(\boldsymbol{y})$ are respectively, a user controlled diffuse transmission color and the fiber-shadowing term for point $\boldsymbol{y}$.

B is the fraction of the incident light that is multiple scattered. By assuming that absorption can be ignored we can find this fraction as the transmitted energy that is not also specularly transmitted at the backside:

$$B(\boldsymbol{\omega}_{\mathrm{i}}, \boldsymbol{\omega}_{\mathrm{o}}) \approx (1 - F_{\boldsymbol{x}})\,(1 - \tau(\boldsymbol{x}, \boldsymbol{y}) \cdot (1 - F_{\boldsymbol{y}})) \quad (14)$$

Note, that we have included user controlled specular and diffuse reflectance and transmission terms. Technically, only the fiber absorption and scattering should be used to control the color of the reflected and the transmitted light, but given the fact that the actual scattering within the fiber volume is approximated and that the non-linear behavior of scattering media is difficult to control we decided to add these explicit user controls. This decoupled control of the different emerging lobes is similar to current approaches for hair rendering [Sadeghi et al. 2010].

4.3 Round and Flat Yarns

We use the BSDFs expressed in §4.1 and §4.2 for both round and flat yarns. For round yarns where the plies are modeled as cylindrical pipes, we use ray tracing to find entrance and exiting points $\boldsymbol{x}$ and $\boldsymbol{y}$. For flat yarns where the plies are modeled as ribbon-like surfaces, on the other hand, we use the

Energy-conservation test for round yarn

Energy-conservation test for flat yarn

same location (i.e., where a light ray intersects a ply) for both $\boldsymbol{x}$ and $\boldsymbol{y}$. In this case, to model the volumetric transmittance of Eq. (11), we assume the ply to have some small thickness h and set

$$\tau = \exp(-\sigma_{\mathrm{t}}\, h/|\boldsymbol{\omega}_{\mathrm{o}} \cdot \boldsymbol{n}(\boldsymbol{x})|), \quad (15)$$

where $\boldsymbol{\omega}_{\mathrm{o}}$ is refracted direction of the light ray at $\boldsymbol{x}$. Our model satisfies energy conservation (i.e., not creating energy), which we
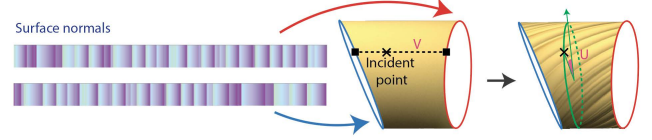

Fig. 6. **Ply surface parameterization:** We parameterize the surface of a ply so that points with identical V coordinates correspond to individual ply cross sections. By assigning 1D textures specifying fiber normal and tangent information to a discrete set of cross sections, we can then obtain continuous fiber details on the ply surface by linearly interpolating the two 1D textures.

verify by rendering plies under uniform white illuminations (see the adjacent figure).

5 IMPLEMENTATION DETAILS

To render our ply-based models, we implement a path tracer in KeyShot [Luxion 2020] that supports environmental and local lighting, on-the-fly geometry generation, and next-event estimation (NEE). To accelerate ray-ply-intersection computations, we generate the ply curves and build a bounding volume hierarchy (BVH) during preprocessing. At render time, we use a custom shader to both evaluate and importance sample our shading model described in §4. In what follows, we provide more details on our rendering process.

5.1 Ray-Ply Intersection

We store plies as polylines and store position, curve normal, and curve length information at each vertex. At render time, we compute ray-ply intersections using ply surfaces defined implicitly around their centerlines as follows.

- We treat *round plies* as cylindrical pipes. Specifically, we define a cylinder around each segment of a ply centerline. To ensure that two consecutive segments meet seamlessly, we make the two cylinders to join at an ellipse perpendicular to the average tangent of the two segments. To obtain shading frames continuous across the ply surface, we define these frames on the joints of adjacent segments and interpolate them for the interior points (see Figure 6).

- For *flat plies*, each segment is rendered as a bilinear surface whose orientation is determined using the curve normal. We use a standard bilinear patch intersection technique [Ramsey et al. 2004] to ray-trace these segments.

5.2 Fiber-Level Details

As explained in §3.3, we enrich the ply surfaces with fiber details using normal and tangent mapping. To this end, we UV-map the ply surface such that the V-direction is aligned with the ply tangent and all points belonging to the joint of two segments have identical V coordinates. In this way, all interior points along a segment with the same V coordinate give the boundary of a ply cross section. For round plies this is an ellipse and a line segment for flat plies. This

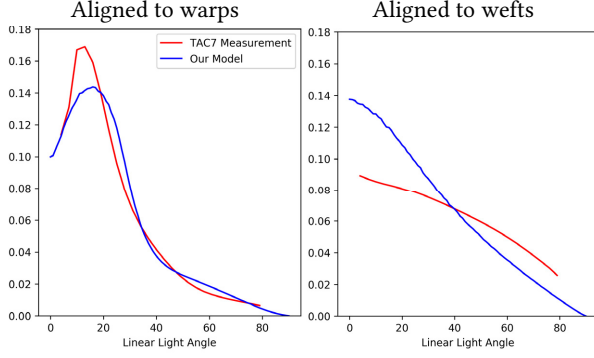**Aligned to warps** **Aligned to wefts**

Fig. 7. **Comparison with X-Rite TAC7 measurements:** Comparison of linear light measurements using an X-Rite TAC7 scanner. The sample was rotated such that the warps faced the linear light on the left plot whereas the wefts faced the linear light on the right. In this example, the pattern is satin.

(a) Black plain (b) Pink satin (c) Blue twill (d) Ivory satin
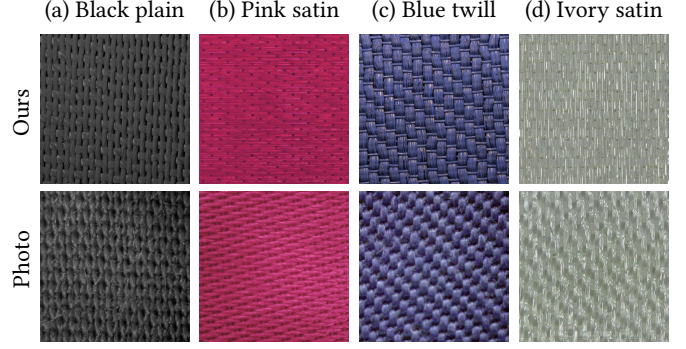
Ours

Photo

Fig. 8. **Matching photographed samples at the micro scale:** Our system is capable of representing multi-ply yarns along with the fiber-level details using an interactive user interface. Once the geometry and appearance parameters are set at this level, large scale models can be rendered with the same material as shown in Figure 9.

ellipse or line determines our U coordinate as the phase around the ellipse or the parameter along the line segment.

Based on this surface parameterization, we use multiple 1D textures to specify the fiber normal and tangent at a few discrete locations (each with fixed V coordinates) along each ply. By interpolating these textures along the ply, as illustrated in Figure 6, we obtain normal and tangent directions that are C^0-continuous over the ply surfaces.

6 RESULTS

In what follows, we show rendered results generated with our technique. Specifically, we validate the physical accuracy of our model in §6.1 and §6.2. Then, we demonstrate the effectiveness of our method by comparing to prior models in §6.3. Lastly, we show additional renderings generated with our model in §6.4.

6.1 Validation with Measurements

We compare our model against a detailed scan of a silk sample using an X-Rite TAC7 scanner, which is a sophisticated scanning device with built-in color calibration, multiple cameras, as well as directional and a linear illuminations. The device reconstructs the scanned sample in 3D and uses this reconstruction to re-project all camera images onto a unified coordinate system for all camera angles. By moving a linear light source in fixed angular increments, the device is capable to accurately recovering reflectance profiles of highly specular materials.

The pixel values $I(x, y)$ of the image output from the TAC7 (after re-projection) is determined by the ratio between radiance exiting each surface point and the incident irradiance at that point:

$$I(x, y) = \frac{L_o(\mathbf{p}, \omega_o)|\mathbf{n} \cdot \omega_i|}{E(\mathbf{p})} = \frac{L_o(\mathbf{p}, \omega_o)|\mathbf{n} \cdot \omega_i|}{\int |\mathbf{n} \cdot \omega_i| L_i(\mathbf{p}, \omega_i) \, d\omega(\omega_i)}. \quad (16)$$

To validate the physical accuracy of our model, we rebuild this measurement configuration virtually. Specifically, we use orthographic cameras to bypass the re-projection step needed by the TAC7 scanner. In this way, the rendered image directly gives the radiance

L_o at the sample's surface. To obtain the incident irradiance $E(\mathbf{p}) = \int |\mathbf{n} \cdot \omega_i| L_i(\mathbf{p}, \omega_i) \, d\omega(\omega_i)$, we render a unit-white diffuse material, producing an image with the value $\frac{1}{\pi} \int |\mathbf{n} \cdot \omega_i| L_i(\omega_i) \, d\omega(\omega_i) = E(\mathbf{p})/\pi$.

In Figure 7, we plot the mean intensity of the measured and rendered image as a function of illumination direction. In this figure, the red curve are obtained using the scanned data and exhibits off-specular peaks due to the arrangement of the warps. The blue curve shows our model, which faithfully reproduces the same off-specular peak although the measured value is slightly higher and narrower. The weft orientation, on the other hand, is almost diffuse in its appearance, confirming the highly anisotropic nature of this material.

Given the difficulty of matching the unknown yarn-geometry of physical cloth samples and the corresponding light-scattering parameters, we consider the match between the measured data and the our model being quite reasonable. Recall that our yarn geometry is obtained using the relaxation discussed in §3.1, which is not expected to match the exact shape of the real-world sample.

6.2 Validation with Photographs

To further validate the physical accuracy of our model, we create virtual replica of a few real-world fabrics using our model. We manually model the yarn-level geometry (by examining the samples under a digital microscope) and rely on the procedural yarn model [Zhao et al. 2016] (of the corresponding materials) to add fiber details. Given the cloth geometry, we search for the values of our light-scattering model parameters (§4) automatically by solving an inverse-rendering problem. Specifically, we minimize the L_2 distance between the average intensity of the rendered images and the photographs using the sequential least squares programming (SLSQP) method [Kraft 1994] implemented in SciPy [Virtanen et al. 2020]. The parameter fitting process can be improved by using gradient-based methods like Adam [Kingma and Ba 2014] with differentiable rendering of our model, although this is beyond the scope of the paper.

| | Ours | Photo | Fiber-based (BCSDF) | | | Ours | Photo | Fiber-based (BCSDF) | | |
| | | | Fit both | Fit front | Fit back | | | Fit both | Fit front | Fit back |

(a) Black plain (b) Pink satin
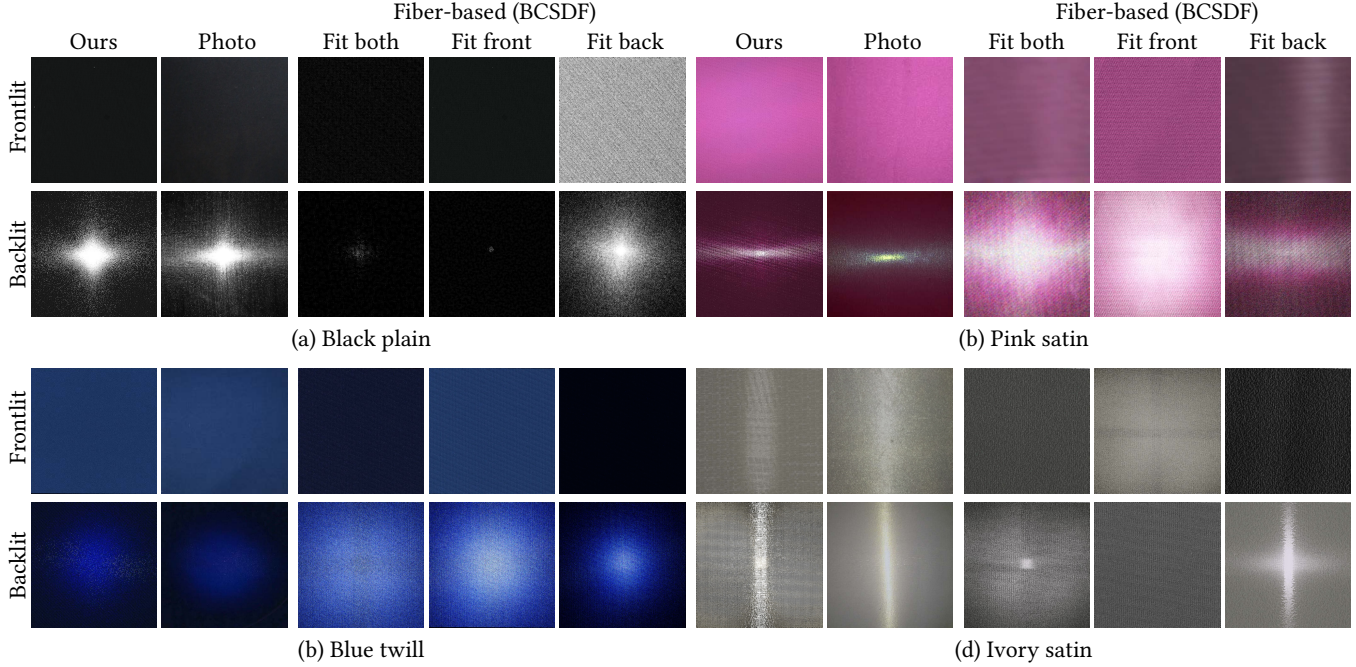
(b) Blue twill (d) Ivory satin

Fig. 9. **Matching photographs at the macro scale for Ours and Fiber-based models:** Using the same samples and parameters as in figure 8 our model is able to simulate the anisotropic highlights seen in the photographs at macro scale. Note, how both the plain weave and the satin has very different appearances when the sample is front lit compared to the back lit sample; our model is able to simulate this difference and match the photographs. In the last three columns we use fiber-based geometry and BCSDF [Khungurn et al. 2016]. While BCSDF model can not match the front and back lighting simultaneously, our system found a descent fit for matching only one side.

Figures 8 and 9 respectively compare renderings generated using our model and photographs at two greatly varying physical scales.

6.3 Comparison with Previous Methods

We demonstrate the flexibility of our model by comparing to the fiber-based representation introduced by Khungurn et al. [2016], which depicts a fabric as a collection of micron-diameter fibers, in Figure 9.

We generate the fiber geometry using the same yarn centerlines. Unlike our method that expresses fiber details via normal and tangent mapping on ply surfaces, this fiber-based model directly stores all the procedurally generated fibers. To obtain the fiber-scattering parameters proposed in the same work, we solve inverse-rendering problems using three different objectives. In the "fit both" scenario, we minimize the difference of both front-lit and backlit renderings. In the "fit front" and "fit back" cases, we consider only the front-lit and the backlit rendering, respectively. Our experiment indicates that the fiber-based model, in its current form, has difficulties matching both reflection and transmission. This could be due to a number of simplifying assumptions being violated in reality (e.g., fibers having spatially invariant circular cross sections). In other words, although the fiber-based model is more detailed than our formulation, its provides insufficient flexibility for correcting numerous approximation errors introduced via the modeling process. Our model, on the other hand, is much more flexible and can easily match both front- and backlit photographs, thanks to its aggregate nature.
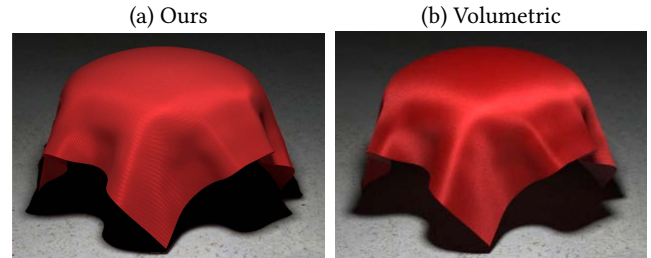

(a) Ours (b) Volumetric

Fig. 10. **Appearance-matching volumetric models**. Despite the great difference between our model and the volumetric one that relies on anisotropic radiative transfer, our technique is able to produce appearances that are quite similar, thanks to its flexibility.

Lastly, we further demonstrate the flexibility of our model by fitting to the rendering of a volumetric cloth model expressed as a micro-flake-based anisotropic medium [Zhao et al. 2011]. Notice that this model has been demonstrated to be less flexible than the fiber-based model [Khungurn et al. 2016]. This is mainly because the model relies on the radiative transfer framework to simulate light scattering in cloth. Without explicit fiber surfaces, phenomena like Fresnel reflection cannot be accurately modeled by this approach, yielding the lack of white highlights on colored fabrics. Additionally, this model has to model glossy appearances using high

optical density, drastically decreasing the physical plausibility of the modeling of light transmission.

Despite the significant difference between our model and the volumetric one, we manage to obtain a close match, as demonstrated in Figure 10.

6.4 Additional Results

We now show additional rendered results produced using our model.

In Figure 11, we show two frames of an animation where a large piece of cloth drapes onto two spheres. Our model is capable of producing temporally consistent results down to individual flyaway fibers. Please refer to the accompanying video for the full animation.

In Figure 12 we show renderings of a lamp. When the lamp is off, our model shows mainly the reflected lighting, thus visualizing the texture that is mapped to the surface color. When the lamp is on, it shows the anisotropic transmission of the light as well as the change of color which is derived from the medium. We note that the backlit highlight orientation is perpendicular to the flat yarns as appeared in Figure 9-b and d as well. This is further demonstrated in the supplementary video.

Lastly, Figure 13 demonstrates that our model can be easily manipulated. Starting with the fitted parameters of the pink satin in Figure 9, we edit the model in the following steps: (a) changing yarn color, (b) reducing roughness, and (c) changing weave pattern (from satin to basket) and increasing the amount of flyaway fibers.

7 DISCUSSION AND CONCLUSION

Limitations and future work. Our technique works only for woven fabrics due to the simple weave pattern with a warp and weft structure. Similar methodology can be employed for knitted fabrics to handle more complex patterns. Generalizing our practical ply-based model to knit textiles should be fairly straight forward with the main challenge being modeling the knit yarn geometry.

We would also like to explore a multi-resolution representation to further compress the geometry for distant views where the ply and fiber detail is not visible. Ideally, this would be a gradual transition from our model to a coarse surface-based cloth model.

Our model follows previous procedural cloth models (e.g. [Zhao et al. 2016]) and inherit the simplifying assumptions (i.e. fiber cross-sections are assumed to be circular while in reality they vary and may have concave shapes). Improving fiber details could lead to a more realistic detailed appearance.

Likewise, it would be interesting to explore a more accurate model for volume scattering within a bundle of fibers.

Conclusion. We have presented an appearance model for woven fabrics that takes a surface mesh and a weave pattern as input and represents the clothing using interwoven ply curves with flyaway fibers. The ply curves are rendered using a BSDF model that accounts for both reflection and transmission. Fiber details are added to the ply BSDF model using precomputed fiber normals and shadows. Our results show that our model is able to match measured data for both reflection and transmission with intuitive physics-based parameters and artistically pleasing tool. Furthermore, our model matches previous surface-, volume-, and fiber-based models

Table 2. **Performance statistics** for renderings of our models in Figures 9–12. All rendering have been done using a dual Xeon E5-2699 system.

Fig. 9-a	Fig. 9-b	Fig. 9-c	Fig. 9-d	Fig. 10	Fig. 11
139.8 MB	364.6 MB	100.3 MB	142.3 MB	308.3 MB	906.1 MB
191 sec	60 sec	93 sec	80 sec	52 sec	41 sec
Fig. 13-a	Fig. 13-b	Fig. 13-c	Fig. 13-d	Fig. 13-e	Fig. 12
449.9 MB	449.9 MB	899.9 MB	899.9 MB	99.2 MB	7.9 MB
76 sec	81 sec	81 sec	62 sec	68 sec	84 sec

for reflection, and is at the same time practical for simulating transmission including the anisotropic highlights often seen through clothing. Our ply based representation is easy to edit, compact, and stable as the cloth surface is deformed making it practical to use for woven fabrics.

ACKNOWLEDGMENTS

REFERENCES

Neeharika Adabala, Nadia Magnenat-Thalmann, and Guangzheng Fei. 2003. Real-Time Rendering of Woven Clothes. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '03)*. Association for Computing Machinery, New York, NY, USA, 41–47. https://doi.org/10.1145/1008653.1008663

Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. 2017. An appearance model for textile fibers. In *Computer Graphics Forum*, Vol. 36. 35–45.

P W Anderson. 1972. More is different. *Science (New York, N.Y.)* 177, 4047 (1972), 393–396. https://doi.org/10.1126/science.177.4047.393 arXiv:arXiv:1011.1669v3

Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. 1999. Reflectance and Texture of Real-World Surfaces. *ACM Trans. Graph.* 18, 1 (Jan. 1999), 1–34. https://doi.org/10.1145/300776.300778

Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. 2011. An Energy-Conserving Hair Reflectance Model. In *Proceedings of the Twenty-Second Eurographics Conference on Rendering (EGSR '11)*. Eurographics Association, Goslar, DEU, 1181–1187. https://doi.org/10.1111/j.1467-8659.2011.01976.x

R. Fletcher. 1987. *Practical Methods of Optimization; (2nd Ed.)*. Wiley-Interscience, USA.

Google. 2019. Physically Based Rendering in Filament. https://google.github.io/filament/Filament.md.html.

Pat Hanrahan and Wolfgang Krueger. 1993. Reflection from Layered Surfaces Due to Subsurface Scattering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. Association for Computing Machinery, New York, NY, USA, 165–174. https://doi.org/10.1145/166117.166139

Eric Heitz. 2018. Sampling the GGX Distribution of Visible Normals. *Journal of Computer Graphics Techniques (JCGT)* 7, 4 (30 November 2018), 1–13. http://jcgt.org/published/0007/04/01/

Piti Irawan and Steve Marschner. 2012. Specular Reflection from Woven Cloth. *ACM Trans. Graph.* 31, 1, Article Article 11 (Feb. 2012), 20 pages. https://doi.org/10.1145/2077341.2077352

Kei Iwasaki, Yoshinori Dobashi, and Tomoyuki Nishita. 2012. Interactive Bi-Scale Editing of Highly Glossy Materials. *TOG* 31, 6 (Nov. 2012), 144:1–144:7. https://doi.org/10/f96zzf

Henrik Wann Jensen and Juan Buhler. 2002. A Rapid Hierarchical Rendering Technique for Translucent Materials. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '02)*. Association for Computing Machinery, New York, NY, USA, 576–581. https://doi.org/10.1145/566570.566619

Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A Practical Model for Subsurface Light Transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. Association for Computing Machinery, New York, NY, USA, 511–518. https://doi.org/10.1145/383259.383319

Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2008. Simulating Knitted Cloth at the Yarn Level. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 1–9. https://doi.org/10.1145/1360612.1360664

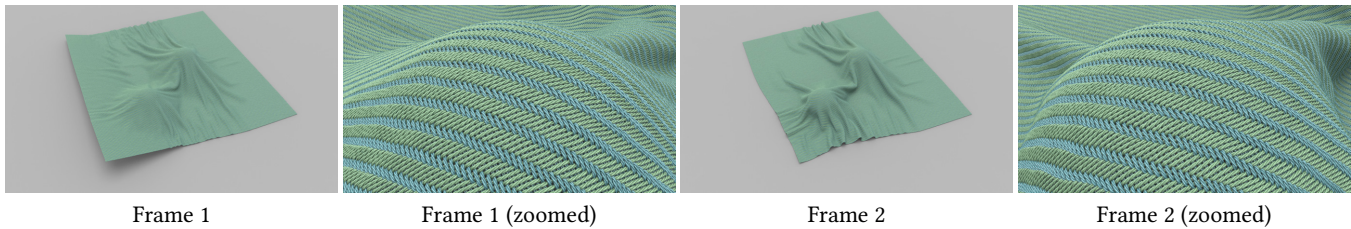Frame 1   Frame 1 (zoomed)   Frame 2   Frame 2 (zoomed)

Fig. 11. **Temporal consistency:** We show two frames of a rendered animation where a cloth drapes onto two spheres and the ground. Our technique manages to produce stable and temporally consistent renderings. Please see the supplementary video for the full animation.
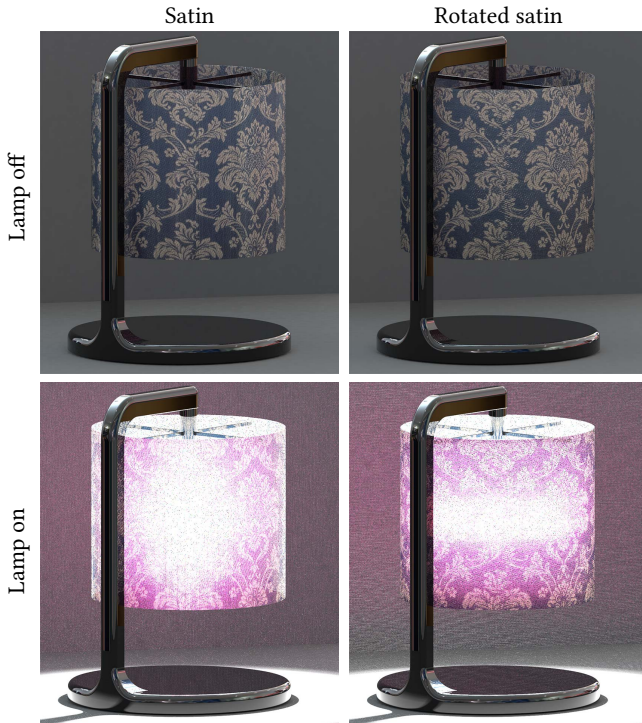

Fig. 12. **Anisotropic transmission for satin pattern in two orientations**. When the lamp is off, our model shows mainly the reflected lighting, showing the texture that is mapped to the surface color. And when the lamp is on, it shows the anisotropic transmission of the light as well and the change of color. Note the backlit highlight orientation is perpendicular to the flat yarns as also expressed in Figure 9-b and d. This is further demonstrated in the supplementary video.

Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2016. Matching Real Fabrics with Micro-Appearance Models. *ACM Trans. Graph.* 35, 1, Article Article 1 (Dec. 2016), 26 pages. https://doi.org/10.1145/2818648

Pramook Khungurn, Rundong Wu, James Noeckel, Steve Marschner, and Kavita Bala. 2017. Fast Rendering of Fabric Micro-Appearance Models under Directional and Spherical Gaussian Lights. *ACM Trans. Graph.* 36, 6, Article Article 232 (Nov. 2017), 15 pages. https://doi.org/10.1145/3130800.3130829

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

Dieter Kraft. 1994. Algorithm 733: TOMP–Fortran modules for optimal control calculations. *ACM Transactions on Mathematical Software (TOMS)* 20, 3 (1994), 262–281.

Alexandr Kuznetsov, Miloš Hašan, Zexiang Xu, Ling-Qi Yan, Bruce Walter, Nima Khademi Kalantari, Steve Marschner, and Ravi Ramamoorthi. 2019. Learning Generative Models for Rendering Specular Microgeometry. *ACM Trans. Graph.* 38, 6, Article Article 225 (Nov. 2019), 14 pages. https://doi.org/10.1145/3355089.3356525

Jonathan Leaf, Rundong Wu, Eston Schweickart, Doug L James, and Steve Marschner. 2019. Interactive design of periodic yarn-level cloth patterns. *ACM Transactions on Graphics* 37, 6 (jan 2019), 1–15. https://doi.org/10.1145/3272127.3275105

Guillaume Loubet and Fabrice Neyret. 2018. A new microflake model with microscopic self-shadowing for accurate volume downsampling. In *Computer Graphics Forum*, Vol. 37. 111–121.

Fujun Luan, Shuang Zhao, and Kavita Bala. 2017. Fiber-Level On-the-Fly Procedural Textiles. *Comput. Graph. Forum* 36, 4 (July 2017), 123–135. https://doi.org/10.1111/cgf.13230

Luxion. 2020. KeyShot. https://www.keyshot.com.

Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light Scattering from Human Hair Fibers. *ACM Trans. Graph.* 22, 3 (July 2003), 780–791. https://doi.org/10.1145/882262.882345

Nelson Max. 1990. Cone-Spheres. In *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '90)*. Association for Computing Machinery, New York, NY, USA, 59–62. https://doi.org/10.1145/97879.97885

Stephen McAuley, Stephen Hill, Naty Hoffman, Yoshiharu Gotanda, Brian Smits, Brent Burley, and Adam Martinez. 2012. Practical Physically-Based Shading in Film and Game Production. In *ACM SIGGRAPH 2012 Courses (SIGGRAPH '12)*. Association for Computing Machinery, New York, NY, USA, Article Article 10, 7 pages. https://doi.org/10.1145/2343483.2343493

Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. 2015. Multi-Scale Modeling and Rendering of Granular Materials. *SIGGRAPH* 34, 4 (July 2015), 49:1–49:13. https://doi.org/10/gfzndr

Zahra Montazeri, Chang Xiao, Yun Fei, Changxi Zheng, and Shuang Zhao. 2019. Mechanics-Aware Modeling of Cloth Appearance. *IEEE Transactions on Visualization and Computer Graphics* PP (08 2019), 1–1. https://doi.org/10.1109/TVCG.2019.2937301

Jonathan T. Moon and Stephen R. Marschner. 2006. Simulating Multiple Scattering in Hair Using a Photon Mapping Approach. *SIGGRAPH* 25, 3 (July 2006), 1067–1074. https://doi.org/10/cd4qf5

Jonathan T. Moon, Bruce Walter, and Steve Marschner. 2008. Efficient Multiple Scattering in Hair Using Spherical Harmonics. *SIGGRAPH* 27, 3 (Aug. 2008), 31:1–31:7. https://doi.org/10/d6r3zt

Jonathan T. Moon, Bruce Walter, and Stephen R. Marschner. 2007. Rendering Discrete Random Media Using Precomputed Scattering Solutions. In *EGSR*. 231–242. https://doi.org/10/gfzp5n

Thomas Müller, Marios Papas, Markus Gross, Wojciech Jarosz, and Jan Novák. 2016. Efficient Rendering of Heterogeneous Polydisperse Granular Media. *SIGGRAPH* 35, 6 (Nov. 2016), 168:1–168:14. https://doi.org/10/f9cm65

Shaun D. Ramsey, Kristin Potter, and Charles Hansen. 2004. Ray Bilinear Patch Intersections. *Journal of Graphics Tools* 9, 3 (2004), 41–47. https://doi.org/10.1080/10867651.2004.10504896 arXiv:https://doi.org/10.1080/10867651.2004.10504896

Boris Raymond, Gaël Guennebaud, and Pascal Barla. 2016. Multi-Scale Rendering of Scratched Materials Using a Structured SV-BRDF Model. *ACM Trans. Graph.* 35, 4, Article Article 57 (July 2016), 11 pages. https://doi.org/10.1145/2897824.2925945

Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A Practical Microcylinder Appearance Model for Cloth Rendering. *ACM Trans. Graph.* 32, 2, Article Article 14 (April 2013), 12 pages. https://doi.org/10.1145/2451236.2451240

Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. 2010. An artist friendly hair shading system. *ACM Transactions on Graphics, SIGGRAPH 2010* 29, 4 (July 2010), 56:1–56:10.

Kai Schröder, Reinhard Klein, and Arno Zinke. 2011. A Volumetric Approach to Predictive Rendering of Fabrics. In *Proceedings of the Twenty-Second Eurographics Conference on Rendering (EGSR '11)*. Eurographics Association, Goslar, DEU, 1277–1286.

https://doi.org/10.1111/j.1467-8659.2011.01987.x

Kai Schroder, Arno Zinke, and Reinhard Klein. 2015. Image-Based Reverse Engineering and Visual Prototyping of Woven Cloth. *IEEE Transactions on Visualization and Computer Graphics* 21 (02 2015), 188–200. https://doi.org/10.1109/TVCG.2014.2339831

Peter-Pike Sloan, Jan Kautz, and John Snyder. 2002. Precomputed Radiance Transfer for Real-time Rendering in Dynamic, Low-frequency Lighting Environments. *ACM Trans. Graph.* 21, 3 (2002), 527–536.

Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods* 17, 3 (2020), 261–272.

Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. Eurographics Association, Goslar, DEU, 195–206.

Hongzhi Wu, Julie Dorsey, and Holly Rushmeier. 2011. Physically-Based Interactive Bi-Scale Material Design. *SIGGRAPH Asia* 30, 6 (Dec. 2011), 1. https://doi.org/10/cqtdrq

Hongzhi Wu, Julie Dorsey, and Holly Rushmeier. 2013. Inverse Bi-Scale Material Design. *SIGGRAPH Asia* 32, 6 (Nov. 2013), 163:1–163:10. https://doi.org/10/gbd5kf

Kui Wu and Cem Yuksel. 2017. Real-time Cloth Rendering with Fiber-level Detail. *IEEE Transactions on Visualization and Computer Graphics* PP, 99 (2017), 1–1. https://doi.org/10.1109/TVCG.2017.2731949

Ling-Qi Yan, Henrik Wann Jensen, and Ravi Ramamoorthi. 2017. An Efficient and Practical Near and Far Field Fur Reflectance Model. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2017)* 36, 4 (2017).

Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. Physically-Accurate Fur Reflectance: Modeling, Measurement and Rendering. *ACM Trans. Graph.* 34, 6, Article Article 185 (Oct. 2015), 13 pages. https://doi.org/10.1145/2816795.2818080

Shuang Zhao, Miloš Hašan, Ravi Ramamoorthi, and Kavita Bala. 2013. Modular Flux Transfer: Efficient Rendering of High-Resolution Volumes with Repeated Structures. *ACM Trans. Graph.* 32, 4, Article Article 131 (July 2013), 12 pages. https://doi.org/10.1145/2461912.2461938

Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building Volumetric Appearance Models of Fabric Using Micro CT Imaging. *ACM Trans. Graph.* 30, 4, Article Article 44 (July 2011), 10 pages. https://doi.org/10.1145/2010324.1964939

Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-Aware Synthesis for Predictive Woven Fabric Appearance. *ACM Trans. Graph.* 31, 4, Article Article 75 (July 2012), 10 pages. https://doi.org/10.1145/2185520.2185571

Shuang Zhao, Fujun Luan, and Kavita Bala. 2016. Fitting Procedural Yarn Models for Realistic Cloth Rendering. *ACM Trans. Graph.* 35, 4, Article Article 51 (July 2016), 11 pages. https://doi.org/10.1145/2897824.2925932

A. Zinke and A. Weber. 2007. Light Scattering from Filaments. *IEEE Transactions on Visualization and Computer Graphics* 13, 2 (March 2007), 342–356. https://doi.org/10.1109/TVCG.2007.43

Arno Zinke, Cem Yuksel, Andreas Weber, and John Keyser. 2008. Dual Scattering Approximation for Fast Multiple Scattering in Hair. *SIGGRAPH* 27, 3 (Aug. 2008), 32:1–32:10. https://doi.org/10/bt33s4

A SELF-ENERGY OF A YARN CURVE

A segment of the height function h_i between two knots s_k and s_{k+1} can be expressed simply in terms of the control points at the knots $h_{i,k}, h_{i,k+1}, \alpha_{i,k}, \alpha_{i,k+1}$:

$$h_i(s) = C_0(\tilde{s})h_{i,k} + C_1(\tilde{s})\frac{\alpha_{i,k}}{\Delta s_k} + C_2(\tilde{s})\frac{\alpha_{i,k+1}}{\Delta s_k} + C_3(\tilde{s})h_{i,k+1}, \quad (17)$$

where $\Delta s_k = s_{k+1} - s_k$, $\tilde{s} = (s - s_k)/\Delta s_k$ and C_i are the cubic Hermite basis functions on the unit interval. In the following we omit the yarn index i for brevity. After some straightforward but tedious algebra, the contribution to the strain-energy of the segment in $[s_k, s_{k+1}[$ is:

$$\frac{15 E_k}{T\Delta s_k} = 2\alpha_k^2 \Delta x_i^2 - \alpha_k \alpha_{k+1} \Delta x_i^2 + 3\alpha_k \Delta x_i h_k$$
$$- 3\alpha_k \Delta x_i h_{k+1} + 2\alpha_{k+1}^2 \Delta x_i^2 + 3\alpha_{k+1}\Delta x_i h_k$$
$$- 3\alpha_{k+1}\Delta x_i h_{k+1} + 18 h_k^2 - 36 h_k h_{k+1} + 18 h_{k+1}^2, \quad (18)$$

which, in turn, gives us the gradient of the energy as:

$$\frac{15}{T\Delta x_i}\frac{\partial E_k}{\partial h_k} = 3\Delta x_i \alpha_k + 3\Delta x_i \alpha_{k+1} + 36 h_k - 36 h_{k+1}, \quad (19)$$

$$\frac{15}{T\Delta x_i}\frac{\partial E_k}{\partial h_{k+1}} = -3\Delta x_i \alpha_k - 3\Delta x_i \alpha_{k+1} - 36 h_k + 36 h_{k+1}, \quad (20)$$

$$\frac{15}{T\Delta x_i}\frac{\partial E_k}{\partial \alpha_k} = 4\Delta x_i^2 \alpha_k - \Delta x_i^2 \alpha_{k+1} + 3\Delta x_i h_k - 3\Delta x_i h_{k+1}, \quad (21)$$

$$\frac{T\Delta x_i}{15}\frac{\partial E_k}{\partial \alpha_{k+1}} = -\Delta x_i^2 \alpha_k + 4\Delta x_i^2 \alpha_{k+1} + 3\Delta x_i h_k - 3\Delta x_i h_{k+1}. \quad (22)$$

Fitted to measurement	(a) Color	(b) Glossiness	(c) Weave pattern

Fig. 13. **Editing geometry and appearance**. Starting from the Figure 9-b where we use fitted parameters to measurement: (a) Change the color (b) edit glossiness value, (c) change the weave pattern to basket with pronounced flyaways. Note, the subtle appearance of flyaways on the basket pattern which greatly increases realism along the edges.