

A Realistic Surface-based Cloth Rendering Model

Junqiu Zhu
University of California, Santa
Barbara
Meta Reality Labs Research
USA

Adrian Jarabo
Meta Reality Labs Research
Spain

Carlos Aliaga
Meta Reality Labs Research
USA

Ling-Qi Yan
University of California, Santa
Barbara
USA

Matt Jen-Yuan Chiang
Meta Reality Labs Research
USA




Figure 1: The curtain here consists of different types of cloth materials (see Table 1 for details). State-of-the-art surface-based models ([Irawan and Marschner 2012] and [Jin et al. 2022]) only consider the *local* surface reflection. In contrast, our method accounts for *non-local* effects including shadowing-masking, cloth surface transmission, and delta transmission. These are important for accurately represent a wide range of fabrics, as we show for a thick polyester satin (top) and a thin cotton plain weave (bottom) fabrics.

ABSTRACT

We propose a surface-based cloth shading model that generates realistic cloth appearance with ply-level details. It generalizes previous surface-based models to a broader set of cloth including knitted and thin woven cloth. Our model takes into account the most dominant visual features of cloth, including anisotropic S-shaped reflection highlight, cross-shaped transmission highlights, delta transmission, and shadowing masking. We model these elements via a comprehensive micro-scale BSDF and a meso-scale effective BSDF formulation. Then, we propose an implementation that leverages the Monte Carlo sampler of path tracing for reducing precomputation to the bare minimum, by evaluating the effective BSDF as a Monte Carlo estimate, and encoding visibility using anisotropic spherical Gaussians. We demonstrate our model by replicating a set of woven and knitted fabrics, showing good match with respect to captured photographs.

CCS CONCEPTS

• Computing methodologies → Rendering.

KEYWORDS

cloth rendering, microflake

ACM Reference Format:

Junqiu Zhu, Adrian Jarabo, Carlos Aliaga, Ling-Qi Yan, and Matt Jen-Yuan Chiang. 2023. A Realistic Surface-based Cloth Rendering Model. In

Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings), August 6–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3588432.3591554>

1 INTRODUCTION

In this work, we propose a surface-based scattering model for cloth, which accurately simulates both the reflection and transmission due to interactions at the different scales of cloth. Cloth rendering remains challenging due to its complex structure – a piece of cloth is a complex aggregation of fibers, twisted into plies, then twisted into yarns, then woven or knitted together. Each level of aggregation forms a distinct type of geometric structure, therefore resulting in a unique kind of appearance for different types of cloth.

Different methods have been proposed for reproducing the appearance of cloth. These might be roughly categorized on curve-based and surface-based. Curve-based methods explicitly model individual fibers/plies/yarns as curves, and model their scattering using the *bidirectional curve scattering distribution function* (BCSDF) similarly to hair rendering [Zhao et al. 2011a; Schroder et al. 2011; Aliaga et al. 2017]. While they represent cloth with high precision, they are slow to render, require heavy storage, are difficult to edit, and cannot be efficiently filtered in render time. On the other hand, surface-based methods model cloth as an explicit surface, and encode the scattering at a given point using the *bidirectional scattering distribution function* (BSDF) [Irawan and Marschner 2012; Sadeghi

et al. 2013]. Surface-based methods are significantly more efficient to represent and render, but they result in a flat and opaque appearance, fail at representing high-frequency details such as highlights from individual fibers, and usually lack physical accuracy. Moreover, efficient surface-based models rely on assumptions that are specific for woven cloth, thus handling knitted cloth is not straightforward.

Therefore, these two families of works represent a trade off between the accuracy and richness of curve-based representations, and the efficiency of surface-based models based on BSDFs. The goal of our work is to develop a model that gets the best of these two worlds. Namely, the ability to resolve geometric and light scattering details at a fine scale, for both reflectance and transmission, also accurately accounting for self shadowing and masking effects. These two aspects are critical for matching far field anisotropic highlights and color shifting often present in real cloth. In addition, it has to do so for a range of viewing distances and directions efficiently, that is, aggregating the reflectance of arbitrarily sized patches and distorted pixel footprints. To do that, we propose an efficient surface-based model that handles both woven and knitted cloth, resolving ply-level details with accurate light transport effects at such scale. It accounts for both the reflection and transmission components of real cloth, and requires minimal precomputation.

Our approach is built upon the recent *SpongeCake* scattering model [Wang et al. 2022], and models light interactions at two different scales: A micro-scale BSDF modeling ply-level scattering and transmission; and a meso-scale component that handles the shadowing-masking at mesoscopic level accurately, and that is crucial for correctly handling view-dependent scattering and transmission effects, as well as filtering the effective BSDF at the pixel footprint.

To conclude, we propose a novel method for accurately translating cloth appearance from ply-based descriptions to surfaces, which addresses several previously unconsidered effects. Our method’s primary contribution is the development of a model that incorporates the following features:

- include the shadowing effect between plies and yarns, which is approximated using anisotropic spherical gaussians (ASGs) for compact storage. Our shadowing term is spatially varying and can be applied to knitted cloth, unlike the microcylinder model proposed in [Sadeghi et al. 2013].
- include the masking and parallax effects described in [Wu2019], which involves mipmapping normal maps to simulate heightfields, thereby avoiding costly ray-tracing. This step is performed before rendering, adding no additional noise or rendering overhead.
- include the transmission effect, which accounts for both surface transmission through yarns and delta transmission through gaps, producing cross-shaped highlights under backlighting. This feature is crucial for modeling “thin” cloth, which can only be attained through curve-based models. Note that no previous surface-based model includes transmission.

We demonstrate our model against captured cloth samples, efficiently reproducing the appearance of cloth at both close and far views, for both reflection and transmission.

2 RELATED WORK

Cloth rendering involves different aspects including geometric representation, scattering models, appearance capture and matching, and level of details. In the following, we discuss the most related works to ours along two main representations of cloth appearance (based either on explicit curves or on surface representations). For a broader review, we refer to the survey by Castillo et al. [2019].

Curve-based models. A large number of works define the cloth as an assembly of fibers or plies, either through explicit geometry [Zhao et al. 2016a] or as heterogeneous volumetric statistical representations [Zhao et al. 2011b; Schroder et al. 2011]. Both representations can conceptually be categorized together, since both explicit fibers and volumes are proven to equally produce high-quality renderings [Khungurn et al. 2015]. The key advantage of cloth models based on curves is the accurate representation of micro-scale geometric details. However, rendering them efficiently is challenging given the amount of three-dimensional detail, and the large storage requirements. Several acceleration techniques have been proposed, ranging from exploiting the repeatability of the cloth structures [Zhao et al. 2012, 2013], to approximating high-order scattering or precomputing volumetric levels of detail [Khungurn et al. 2017; Zhao et al. 2016b]. These techniques are off-line and require expensive precomputation. Some real-time curve-based cloth models have been proposed by using ply-level representations, baking the fiber details as textures [Wu and Yuksel 2017] or procedurally generating them on-the-fly [Luan et al. 2017]. Nonetheless, the first model lacks physical accuracy by approximating light interactions, and the second remains impractical for the large number of yarns that compose real-world garments. More recent offline ply-level cloth models [Montazeri et al. 2020, 2021] improve over these previous works, including accurate light transport inside the ply and transmission, but still rely on coarse approximations of the volume scattering within fiber bundles. In our work, we propose a surface-based method for cloth up to ply-level detail, which exhibits the visual features of curve-based methods while being more compact and efficient to evaluate and edit.

Surface-based models. Inspired by the seminal paper of Westin et al. [1992] for predicting reflectance from micro-surface details, cloth has been traditionally modelled as a two-dimensional surface, embedding the complex light-geometry interactions in a bidirectional reflectance distribution function (BRDF) [Irawan and Marschner 2012]. Despite depicting important features such as anisotropic highlights, it lacks a number of important features including shadowing and masking effects and transmission. Later, Sadeghi et al. [2013] accounted for shadowing and masking for woven cloth only, and through a simplification that does not consider warp-weft and weft-warp interactions. In addition, their model assumes far field, and ply-level details are missing. Other approaches used the bidirectional texture function (BTF) [Sattler et al. 2003] for implicitly encoding the complex scattering of cloth and the shadowing-masking effects, at the cost of massive storage requirements and little representation flexibility. More recent approaches [Wu et al. 2019; Xu et al. 2019] bake the fiber-level cloth microgeometry into a height-field, and model their scattering using microfacet, microflake, or fiber scattering models. However, these methods cannot model transmission, require a high-resolution heightmap up to fiber-scale,

and needs for heavy precomputations in the order of minutes or hours. Differently, our method requires minimal precomputation and is able to reproduce the fine scale details and lighting effects present in the aforementioned methods. Closer to our work, Jin et al. [2022] used the SpongeCake model [Wang et al. 2022] for efficiently representing the cloth surface reflection in woven cloth, using a normal map to represent detail up to yarn level, at the cost of ignoring shadowing-masking from neighbor yarns. In contrast, our method shares the benefits of Jin et al.'s model, but it resolves up to ply-level details, accounts for accurate shadowing and masking effects for both woven and knitted cloth and handles transmission, which are important features to match real-world cloth.

3 MOTIVATION

Before we delve into our model, here we first analyze the visual features of real-world cloth, to motivate how they fit into our model. These visual characteristics are a consequence of the multiscale nature of cloth, and therefore are directly related with its structure.

Cloth fabrics are built by interlacing several yarns, in such a way that they form a continuous surface with different mechanical and optical characteristics depending on the underlying structure. That structure has two different scales. The cloth are made up of individual yarns: Yarns are built by twisting together a set of *plies*, which are themselves aggregations of several textile fibers, twisted to form a relatively uniform bundle. Depending on the number of plies and their twisting, each low-level yarn might exhibit different appearance properties, which translate to the overall appearance of the fabric. We call this scale "ply level", which is the minimal representation we employ in this work, whereas fiber-level is considered statistically.

At the coarser level, yarns can be interlaced in two main ways: In *woven* cloth, two perpendicular straight yarns (called *warps* and *wefts*) are interlaced following a particular weave pattern; woven cloth has generally a flatter appearance, and masking effects mainly manifests in the form of the anisotropic highlights, or the color shifts in the case of *two-tone fabrics* where the warp and weft yarns have different color. On the other hand, *knitted* cloth is usually made by a single, long yarn, following a rather complex interleaving pattern forming repeatable loops; in contrast to woven cloth, knitted cloth exhibits generally more volume and shadowing-masking is a prevalent effect due to the high variability of visible ply directions along such loop. We call this scale "pattern level".

Structural and optical properties of cloth at both coarse and fine level scales create different visual effects. In the following, we discuss the most important ones, which serve as motivation of our work

3.1 Visual features of cloth at ply level

Curved-shaped reflection highlights. The reflection from plies is strongly anisotropic, following the tangent direction of fibers forming the ply. While at distance the reflection anisotropy roughly follows the direction of the yarn [Sadeghi et al. 2013], at close views the fiber twisting results into a dominant feature, making the highlights from plies to appear curved shaped [Irawan and Marschner 2012]; such curved-shaped highlight is particularly noticeable on knitted cloth (Figure 2, a). The same applies for yarns with multiple plies, breaking up the main highlight into several ones. Handling

this curved-shaped aggregated behaviour motivates us to work on a ply-level basis.

Cross-shaped surface transmission highlight. Despite transmittance being an important feature in cloth, specially in thin fabrics, it has been generally ignored in surface-based cloth models. Similar to reflection, transmission manifests both as a diffuse-like color due to multiple scattering in fibers, and a low-frequency, directional anisotropic transmission highlight aligned to the direction of the yarns. The shape of such highlight depends on the pattern, being particularly characteristic the cross-shape for wovens, aligned to warp and weft yarn directions (Figure 2, b).

3.2 Visual features of cloth at pattern level

Repeatability. A key structural feature of cloth is its repeatable structure of the knit/weave pattern along the whole fabric. This tileability has been leveraged before to precompute light transport in fabric in minimal tiles [Zhao et al. 2013]. However, this repeatability is slightly perturbed in a per-ply basis: The fibers inside each ply are not perfectly and uniformly twisted, but they show irregularities affecting the overall ply's mean tangent direction. Previous works have introduced this spatially-resolved perturbation by adding a noise layer in the final render [Irawan and Marschner 2012], but that approach does not hold for close-ups. We leverage this repeatability by using minimally tileable patterns, which are perturbed in shading time.

Meso-scale shadowing and masking. While the scattering in cloth is mostly local, neighboring plies and yarns can have a strong influence on the global appearance of cloth, especially at grazing angles. This is due to the parallax that creates structured shadowing and masking effects, varying the visible fibers (of different tangent directions and optical properties like albedo) as the viewing direction changes. Curved-based representations implicitly account for such effects, as the geometry of yarns or plies are explicitly modeled and rendered. However, surface-based representations require to mathematically account for these effects, which is specially challenging in knitted cloth compared to wovens, which that can be abstracted into two main perpendicular directions for warp and weft [Sadeghi et al. 2013].

Delta transmission. A key visual feature of thin fabrics is their transparency. It is the result of a relatively sparse weave/knit pattern, which allows light to pass through unscattered. We call this component *delta transmission*, and its angular dependency is directly related to masking (Figure 2, c).

4 OUR METHOD

Our model takes as input a macro-scale cloth mesh whose UV layout aligns with the orientation of the pattern, along with spatially-varying texture maps and shading parameters to describe pattern and ply-level features, as detailed in Section 4.1. Then, we plug these maps into our micro-scale BSDF (Section 4.2): A ply-level model that accounts for both reflection and transmission. In Section 4.3, We propose the meso-scale effective BSDF, aggregating the scattering of the surface points over a pixel footprint. Finally, we detail our implementation.




Figure 2: Illustration of the different visual features of cloth. a) A knitted cloth patch shows distinct structural lines. Such asymmetric feature is caused by the curved-shaped highlight along the staple yarns. b) A back-lit piece of plain weave cloth demonstrates a cross-shaped transmission highlight that follows the direction of the yarns. c) Transparency of thin cloth as a result of a sparse weave pattern, which exhibits a strong angular dependency due to the masking effect. Note that all the images provided are real photographs.




Figure 3: Feature maps used by the model. From left to right, ID map, ply normal $\mathbf{n}_p(p)$, ply tangents $\mathbf{t}_p(p)$, and positions $x(p)$, for both jersey knit (top row) and twill patterns (bottom row).

4.1 Input features

Profiting from the repeatability of the cloth patterns, the input feature maps only require to cover the minimally repeating portion of the pattern and their sizes can be fairly small (under 256x256), specially for simple patterns like plain weave. Figure 3 shows example maps for two different patterns.

Geometry. Since our model supports cloth transmission, we allow two sets of input geometric texture maps for both the front and back faces of the cloth. The input set of maps include: ply normal $\mathbf{n}_p(p)$; ply tangent directions $\mathbf{t}_p(p)$; 3D position $x(p)$; and an ID map specifying the yarn type (warp or weft for woven patterns) as well as the free space between yarns $I(p)$. In addition, we require the twist angle of the fibers inside the plies ω , which we use to compute the fibers tangent direction $\mathbf{t}(p)$ by rotating $\mathbf{t}_p(p)$ around the ply normal. Finally, we introduce variability of the geometric features along the tileable map by adding a noise term $N(p)$ on top of ω , introducing a small perturbation due to fiber irregularity in a ply. We model $N(p)$ using Perlin noise with average amplitude of 5 degrees.

Shading. In addition to the spatially-varying geometric texture maps, we also input shading parameters including albedo, and roughness α . The latter is an expressive parameter that embeds both the roughness of the scattering lobe at the fiber level and models per fiber deviation along the main twisting angle, following Sadeghi et al. [2013].

We synthesize the maps by projecting in 2D the features of a custom made 3D procedural ply-level tool. By taking these maps (as opposed to through implicit generation [Irawan and Marschner

2012]), we avoid assumptions about the shape of the yarns and plies (e.g. elliptical arc [Jin et al. 2022]). This widens the range of cloth types that can be represented by the model, providing the user complete freedom of authorship.

4.2 Micro-scale BSDF

Our micro-scale BSDF builds upon the one proposed by Jin et al. [2022], which is based on the SpongeCake BSDF model [Wang et al. 2022], and adds a diffuse term for simulating diffuse scattering in the cloth. However, Jin et al.’s model is limited to reflection, missing some important features in thin and translucent cloth. Our micro-scale BSDF adds three additional terms to account for transmission, and define our BSDF $f(x, \mathbf{i}, \mathbf{o})$ as

$$\begin{aligned} f(x, \mathbf{i}, \mathbf{o}) = & I(x) \left(f^{r,s}(x, \mathbf{i}, \mathbf{o}) + f^{r,d}(x, \mathbf{i}, \mathbf{o}) \right) \\ & + I(x) \left(f^{t,s}(x, \mathbf{i}, \mathbf{o}) + f^{t,d}(x, \mathbf{i}, \mathbf{o}) \right) \\ & + (1 - I(x)) f^\delta(x, \mathbf{i}, \mathbf{o}), \end{aligned} \quad (1)$$

where \mathbf{i} and \mathbf{o} are the incident and outgoing directions respectively. The first line accounts for reflection in a similar fashion of [Jin et al. 2022]. The second line accounts for transmission and is the sum of two lobes: One SpongeCake transmission lobe $f^{t,s}(x, \mathbf{i}, \mathbf{o})$ that models the directional scattering through the fabric, and that depends on the direction of the fibers at x , and a diffuse transmission $f^{t,d}(x, \mathbf{i}, \mathbf{o})$ that approximates the effect of multiple scattering, in a similar fashion as $f^{r,d}(x, \mathbf{i}, \mathbf{o})$. Finally, the last term $f^\delta(x, \mathbf{i}, \mathbf{o})$ models the delta transmission. All the terms are modulated by an indicator function $I(x)$ which returns 0 if x is in a gap and 1 otherwise. Both SpongeCake-based terms $f^{r,s}(x, \mathbf{i}, \mathbf{o})$ and $f^{t,s}(x, \mathbf{i}, \mathbf{o})$ are defined by a fiber-like normal distribution function following the SGGX distribution [Heitz et al. 2015], aligned to the fiber tangent $\mathbf{t}(x)$ and with roughness α_r and α_t for the reflection and transmission part, respectively. Details on these terms can be found in the Supplementary.

Delta Transmission. An important effect that is missing in previous models is defined by the light paths that traverse the cloth unscattered, which we model as

$$f^\delta(x, \mathbf{i}, \mathbf{o}) = \frac{\delta(\mathbf{i} + \mathbf{o})}{\langle \mathbf{i} \cdot \mathbf{n}_s \rangle}. \quad (2)$$

Note that there is an extra cosine $\langle \mathbf{i} \cdot \mathbf{n}_s \rangle$ in this equation. This is because the energy of delta transmission is not related to the surface area, similar to other singular BSDFs.

Summary. The proposed *micro-scale* BSDF is a ply-level scattering model accounting for all *fiber-to-fiber* interactions inside a ply. It can accurately match reflectance and transmission highlights observed in real-world clothes. However, such micro-scale BSDF is only defined by the local shading information—ply normal and fiber tangent—at an infinitesimal point, thus neglecting important *ply-to-ply* interactions from neighboring regions. In the next section, we introduce a *meso-scale* BSDF to address this shortcoming.

4.3 Meso-scale BSDF

We describe in the following how to accurately account for non-local effects not handled by our *micro-scale* BSDF. To this end, we rely on the *effective BSDF* formulation [Wu et al. 2011], that describes the aggregated reflectance over micro-geometry covered by a pixel footprint. Let \mathcal{P} be a flat geometric surface patch covered by the pixel footprint, with area normalized by the kernel $\kappa_{\mathcal{P}}$ so that $\int_{\mathcal{P}} \kappa_{\mathcal{P}}(p) dp = 1$. Our micro-geometry $\mathcal{G}(\mathcal{P})$ is defined by the set $\mathcal{G}(\mathcal{P}) = \{\mathbf{n}_p(p), \mathbf{t}_p(p), x(p) \mid p \in \mathcal{P}\}$ with $\mathbf{n}_p(p)$ the ply normal, $\mathbf{t}_p(p)$ the ply tangent and $x(p)$ the position. Each point $p \in \mathcal{P}$ have its associated microscale BSDF, defined following Equation (1). With that in place, we define the effective BSDF $f^{\text{eff}}(\mathbf{i}, \mathbf{o})$ as [Wu et al. 2011]

$$f^{\text{eff}}(\mathbf{i}, \mathbf{o}) = \frac{1}{A_{\mathcal{G}}(\mathbf{o})} \int_{\mathcal{P}} f(x(p), \mathbf{i}, \mathbf{o}) \langle \mathbf{n}_p(p) \cdot \mathbf{i} \rangle V(x(p), \mathbf{i}) A_p(p, \mathbf{o}) \kappa_{\mathcal{P}}(p) dp, \quad (3)$$

where $V(x, \mathbf{d})$ is the binary visibility function of point x in direction \mathbf{d} , $\langle \cdot \cdot \rangle$ is the dot product operator clamped to positive values, and $A_{\mathcal{G}}(x, \mathbf{o})$ is the visible projected area along \mathbf{o} given by [Wu et al. 2019]

$$A_p(p, \mathbf{o}) = \frac{\langle \mathbf{o}, \mathbf{n}_p(p) \rangle}{\langle \mathbf{n}_s, \mathbf{n}_p(p) \rangle} V(x(p), \mathbf{o}), \quad (4)$$

where $\langle \mathbf{o}, \mathbf{n}_p(p) \rangle$ is the projection of $x(p)$ in the view direction \mathbf{o} , and $\frac{1}{\langle \mathbf{n}_s, \mathbf{n}_p(p) \rangle}$ is the Jacobian $\left| \frac{dx(p)}{dp} \right|$. Finally, $A_{\mathcal{G}}$ is the total projected area of $\mathcal{G}(\mathcal{P})$ along \mathbf{o} , defined as [Dupuy et al. 2013]

$$A_{\mathcal{G}}(\mathbf{o}) = \int_{\mathcal{P}} \frac{\langle \mathbf{o}, \mathbf{n}_p(p) \rangle}{\langle \mathbf{n}_s, \mathbf{n}_p(p) \rangle} V(x(p), \mathbf{o}) \kappa_{\mathcal{P}}(p) dp = \frac{\langle \mathbf{o}, \mathbf{n}_f(\mathcal{P}) \rangle}{\langle \mathbf{n}_s, \mathbf{n}_f(\mathcal{P}) \rangle}, \quad (5)$$

where $\mathbf{n}_f(\mathcal{P})$ is the average visible microscale normal over the pixel footprint \mathcal{P} defined as through the average slope in \mathcal{P} .

4.4 Implementation

Here we detail our implementation, which is critical for allowing an efficient rendering with very little precomputation and minimal storage, as opposed to previous precomputation-heavy approaches.

4.4.1 Computing the effective BSDF. Previous works building upon the effective BSDF [Dupuy et al. 2013; Wu et al. 2019] preintegrate the integral in Equation (3), at the cost of long precomputation times and significant storage. This also limits the potential application of a small pattern texture, perturbed with high-resolution noise. Instead, inspired by the work of Chiang et al. [2016] on hair rendering, we

calculate stochastically this integral on the fly, by using the Monte Carlo estimate

$$\langle f^{\text{eff}}(\mathbf{i}, \mathbf{o}) \rangle = \frac{1}{N_f A_{\mathcal{G}}(\mathbf{o})} \sum_{i=1}^{N_f} \frac{f_p(p_i, \mathbf{i}, \mathbf{o})}{\text{pdf}(p_i)}, \quad (6)$$

with p_i a random sample in \mathcal{P} , $f_p(p_i, \mathbf{i}, \mathbf{o})$ the integrand in Equation (3), and $\text{pdf}(p_i)$ the probability of sampling p_i , which in our case is uniform in \mathcal{P} . For each intersection at the macroscopic surface, we compute the footprint \mathcal{P} using ray differentials [Igehy 1999], and compute Equation (6). In practice, we leverage on the pixel sampler of the renderer, and set $N_f = 1$; this translates some variance to the final pixel, but results into a significantly faster shading evaluation, which is specially important in the context of global illumination where each path might need evaluating the shader multiple times, and where variance from e.g. illumination might be more dominant than variance from Equation (6).

Note that computing $A_{\mathcal{G}}$ in Equation (6) requires computing $\mathbf{n}_f(\mathcal{P})$, which is an additional integral in the footprint \mathcal{P} ; we opt for precomputing it to reduce potential variance coming from this additional term in the integrand. We follow the approach of Wu et al. [2019], and downsample our position map to minimize the mesoscale average surface slopes for different mipmap levels; then we transform and store them as normal maps, which are queried in render time. Note that precomputing $\mathbf{n}_f(\mathcal{P})$ introduces some bias on the Monte Carlo estimate, since it does not account for the random perturbations on the base pattern maps.

4.4.2 Shadowing and Masking. We account for shadowing and masking via the visibility function $V(p, \mathbf{d})$ in Equations (3) and (4). This term accounts for the local visibility of point $x(p)$ based on the height of neighbor texels p in the meso-geometry, and can be approximated by doing an horizon search in texture space. However, computing $V(p, \mathbf{d})$ explicitly might be very expensive, and precomputing the 4D visibility function is impractical.

Instead, we approximate the visibility at each texel by using an anisotropic spherical gaussian (ASG) [Xu et al. 2013], as

$$V(p, \mathbf{d}) \approx \text{ASG}(\mathbf{d} \mid \boldsymbol{\mu}(p), \gamma(p), \sigma_x(p), \sigma_y(p), C(p)) = C \langle \boldsymbol{\mu}(p) \cdot \mathbf{d} \rangle \exp(-\sigma_x(p)(\mathbf{d} \cdot \mathbf{x}(p))^2 - \sigma_y(p)(\mathbf{d} \cdot \mathbf{y}(p))^2) \quad (7)$$

with $\boldsymbol{\mu}(p)$ the mean vector, and $\boldsymbol{\mu}(p)$, \mathbf{x} and \mathbf{y} forming an orthogonal. The angle $\gamma(p)$ is the rotation of \mathbf{x} and \mathbf{y} around $\boldsymbol{\mu}(p)$, $\sigma_x(p)$ and $\sigma_y(p)$ the bandwidth for the \mathbf{x} - and \mathbf{y} -axis, respectively, and $C(p)$ is the amplitude. Modeling visibility as an ASG allows us to approximate the complex visibility per texel storing only 5 float numbers.

Efficiently fitting ASG to visibility. While a regular non-linear optimization can be used for fitting the ASG, this would require a relatively long precomputation, to be added to the cost of computing $V(p, \mathbf{d})$ itself. Instead, fit the ASG by adding a marginal cost to the computation of the visibility. We compute $V(x, \mathbf{d})$ by using the horizon search approach common in real-time rendering: We radially sample around the normal at x , and for each radial sample ϕ , we compute the horizon elevation angle $\theta(\phi) \in [0, \pi/2]$ within a neighborhood with radius 20 pixels. Then, as we compute the horizon, we iteratively fit the ASG. The mean of the ASG $\boldsymbol{\mu}(p)$ is computed by centering it at the bent normal following [Jiménez et al.

2016]. For fitting the shape parameters $\sigma_x(p)$, $\sigma_y(p)$ and $\gamma(p)$, we observe that a 2D Gaussian in tangent space provides a similar fit to the ASG; therefore, instead of fitting the ASG directly, we compute the covariance matrix of the tangent-space 2D Gaussian, by using an online weighted approximation of the covariance, where we weight each sample by its subtended solid angle.



Figure 4: Measurement setup for cloth appearance capture. The cloth sample is wrapped onto a cylinder, and photographed from a cellphone with known camera intrinsics. The scene is illuminated by the flash light of a second cellphone with known position.

5 RESULTS

We show the results of our method over different scenes, demonstrating its effectiveness at reproducing characteristic light scattering effects in a synthetic scene, and also its ability to match the appearance of real cloth in terms of both reflectance and transmittance, captured with minimalistic set-ups. All scenes shown in the paper are rendered using path tracing on an AMD Ryzen 9 7950X 16-Core (4.50 GHz) machine. In all cases, the input maps have a resolution of 256×256.

Cloth cylinders. Similarly to Jin et al. [2022], we qualitatively demonstrate our model against photographs of cloth on a controlled setup. We take single pictures of the cloth samples wrapping onto a cylinder, lit by the flash of another cellphone, with known camera intrinsics, cylinder sizes, and distances to lights and camera sensors, as shown in Figure 4. Since the appearance of satin is sensitive to viewing and lighting orientation due to its asymmetric pattern and shiny nature, we take two pictures, aligning its warp yarns and then weft to the longitudinal axis of the cylinder. We model a total of three types of patterns: Two weaves (satin and twill) and one knit (jersey). For each fabric, we match the reference picture using the baseline *micro-scale* BSDF model of Jin et al. [2022], and our *meso-scale* BSDF model. Figure 7 shows the comparison for all three fabrics against the photograph.

The satin pattern showcases the highly anisotropic behaviour of certain shiny clothes: its overall appearance drastically changes from the first alignment configuration at 0° to the one at 90° due to the asymmetry of the weaving pattern [Sadeghi et al. 2013]. Despite previous work can reproduce the overall shape and location of the main highlight for the first configuration, the lack of accurate shadowing and masking is visible at grazing angles, and fails when

trying to match the 90°-rotated orientation of the pattern with the same model parameters. In contrast, our *meso-scale* effective BSDF, with accurate shadowing and masking, captures the directionally-varying reflectance of satin both at micro and meso scales.

The other two patterns (twill and knit) are rougher fabric types that scatter more light in the non-specular directions. In that context, the shadowing and masking effect becomes more prominent. Our model preserves sharp ply level details while achieving softer transitions to shadows than previous methods. However, the transition to shadows in the reference is softer than in both models due to a number of reasons (i.e. lack of fly-away fibers) further discussed in Section 6.




Figure 5: Reproducing transmission of real cloth samples. We show reference photographs of back-lit cloths (top row), and renders using our model (bottom row) for a variety of fabric types. From left to right: Plain, satin, twill and satin. Our method approximates well the overall anisotropic shape of the transmission highlights. Photo reference from Montazeri et al. [2020].

Cloth transmission. We showcase the ability of our model to reproduce the complex anisotropic transmittance exhibited by some fabrics in Figure 5. This is a key feature neglected by previous models. The shape and orientation of the highlights directly depend on the patterns, influenced by features like its asymmetry, the free space between yarns and their alignment (specially in the case of woven). For woven cloth, the vertical component of the highlights results from the light scattered and delta-transmitted through horizontally-aligned yarns, whereas the horizontal component is generated by the light interaction with the vertically-aligned yarns. The ratio of the two components depends on the relative coverage of the surface pattern by warp and weft yarns. Our *meso-BSDF* model benefits from accurate shadowing and masking to faithfully reproduce such effect, even in cases where only one component is visible (either vertical or horizontal) due to the shadowing of one kind of yarns over the other one (i.e. warp over weft).

Day-to-day scenes. The curtain in Figure 1 consists of three different types of cloth representing different levels of cloth transmission, from thicker to thinner. Ignoring both transmission and delta transmission, results into lower fidelity when representing this types of cloth. Note also the importance of masking on handling transparency. Finally, we show in Fig. 6 an example of rendering full garments with our model, for an outfit composed of a very shiny satin blouse with anisotropic highlights, and rough twill pattern pants, exhibiting typical real cloth effects at far distance while preserving details. See the supplemental video for a zoom-in of the

Table 1: Statistics of all scenes and performance break-down. We report samples per pixel (SPP), precomputation time (Pre) for both the A_G term and visibility (V), and render time. Note that the precomputation is not optimized. During rendering our meso-scale BSDF only requires around 10% additional computation compared to the baseline [Jin et al. 2022], while it accounts for shadowing-masking effect of neighboring plies and yarns.

Scenes	Type	Roughness (σ)	SPP	Pre (A_G)	Pre (V)	Render (ours)	Render [Jin et al. 2022]	Storage
Curtain	White plain	0.20	128	3 s	7 s	2.8 min	2.5 min	14.3 MB
Curtain	Yellow satin	0.05	128	3 s	7 s	2.8 min	2.5 min	13.6 MB
Curtain	Brown twill	0.20	128	3 s	6 s	2.8 min	2.5 min	17.2 MB
Cylinder	Satin 0°	0.04	256	3 s	7 s	1.4 min	1.3 min	13.6 MB
Cylinder	Satin 90°	0.04	256	3 s	8 s	1.3 min	1.3 min	13.6 MB
Cylinder	Twill	0.25	256	3 s	6 s	1.4 min	1.4 min	17.2 MB
Cylinder	Knitted	0.3	256	4 s	8 s	1.6 min	1.5 min	17.2 MB

scene, demonstrating the all-scales appearance representation of our model.



Figure 6: Example render of a day to day outfit with our model, composed of a polyester satin blouse and cotton twill pants.

6 LIMITATIONS AND FUTURE WORK

Our method has been demonstrated to provide good fidelity on both reflection and transmission, for a variety of cloth types. Moreover, its compactness and low precomputation is very promising for interactive authoring.

Authoring. Currently, we manually define the parameters for reflection and transmission, which might result into non-physical values. While these can be computed via inverse rendering, there is still no guarantee that they represent feasible parameters beyond the acquired data. A more interesting approach would be to fix a small set of parameters based on the cloth structure, and derive the rest (e.g. albedo or transmission roughness) based on the initial parameters.

Multiple scattering. Our binary visibility function introduces an energy loss by neglecting multiple scattering between neighbor points. This causes an overly sharp appearance in some parts of e.g. our knit example in Figure 7. Accounting for such multiple scattering would likely result into softer appearances.

Fly-aways. While we are able to represent the overall look in e.g. knitted cloth, our results are overall sharper than the real counterparts. Fly-away fibers are key to produce that softer look, specially at grazing angles. Incorporating these fly-aways in a way that they can be efficiently treated and filtered is still an open problem.

Level of detail. To apply our model to real-time rendering we need to derive an efficient level-of-detail solution that integrates the BSDF (Equation (6)) on a per-pixel basis to avoid aliasing. Our approach tackles level of detail by means of explicit Monte Carlo integration, which allows a small precomputation and storage overhead. However, for very specular cloth it might be particularly difficult to find a texel contributing to a particular half vector. Exploring the application of techniques targeting glint rendering [Yan et al. 2016] to accelerate rendering in these conditions, and generalizing them for accounting for shadowing and masking, is an exciting avenue for future work.

7 CONCLUSIONS

In this paper we have proposed a new surface-based method for cloth rendering that accounts up to ply level. It includes many important features that are generally ignored in surface-based methods, including shadowing-masking and transmission. We hope our method would help bridging the quality gap between curve-based methods and the more efficient surface-based approaches.

ACKNOWLEDGMENTS

We thank Christophe Hery and Olivier Maury for their valuable input. We also thank Natasha Devaud for her contribution in authoring procedural tools, generating ply-level feature maps, as well as her guidance on the garment selection process. This project is solely sponsored by Meta. And Ling-Qi Yan is also supported by gift funds from Adobe, Intel, Meta and XVerse.

REFERENCES

Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. 2017. An appearance model for textile fibers. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 36, 4 (2017), 35–45.

- Carlos Castillo, Jorge López-Moreno, and Carlos Aliaga. 2019. Recent advances in fabric appearance reproduction. *Computers & Graphics* 84 (2019), 103–121.
- Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. 2016. A Practical and Controllable Hair and Fur Model for Production Path Tracing. *Computer Graphics Forum* 2, 35 (2016), 275–283.
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, Fabrice Neyret, and Victor Ostromoukhov. 2013. Linear Efficient Antialiased Displacement and Reflectance Mapping. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 32, 6 (2013), 211:1–211:11.
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX microflake distribution. *ACM Trans. Graph.* 34, 4 (2015), 1–11.
- Homan Igehy. 1999. Tracing Ray Differentials. *SIGGRAPH* (1999), 179–186.
- Piti Irawan and Steve Marschner. 2012. Specular reflection from woven cloth. *ACM Trans. Graph.* 31, 1 (2012), 1–20.
- Jorge Jiménez, Xianchun Wu, Angelo Pesce, and Adrian Jarabo. 2016. Practical real-time strategies for accurate indirect occlusion. In *SIGGRAPH 2016 Courses: Physically Based Shading in Theory and Practice*.
- Wenhua Jin, Beibei Wang, Milos Hasan, Yu Guo, Steve Marschner, and Ling-Qi Yan. 2022. Woven Fabric Capture from a Single Photo. In *SIGGRAPH Asia 2022 Conference Papers*. 1–8.
- Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. *ACM Trans. Graph.* 35, 1 (2015), 1–1.
- Pramook Khungurn, Rundong Wu, James Noeckel, Steve Marschner, and Kavita Bala. 2017. Fast rendering of fabric micro-appearance models under directional and spherical gaussian lights. *ACM Trans. Graph.* 36, 6 (2017), 1–15.
- Fujun Luan, Shuang Zhao, and Kavita Bala. 2017. Fiber-Level On-the-Fly Procedural Textiles. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 123–135.
- Zahra Montazeri, Soren B Gammelmark, Henrik Wann Jensen, and Shuang Zhao. 2021. A Practical Ply-Based Appearance Modeling for Knitted Fabrics. In *Proceedings of Eurographics Symposium on Rendering 2021*.
- Zahra Montazeri, Soren B Gammelmark, Shuang Zhao, and Henrik Wann Jensen. 2020. A practical ply-based appearance model of woven fabrics. *ACM Trans. Graph.* 39, 6 (2020), 1–13.
- 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 (2013), 1–12.
- Mirko Sattler, Ralf Sarlette, and Reinhard Klein. 2003. Efficient and realistic visualization of cloth. In *Rendering techniques*. 167–178.
- Kai Schroder, Reinhard Klein, and Arno Zinke. 2011. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum* 30, 4 (2011), 1277–1286.
- Beibei Wang, Wenhua Jin, Miloš Hašan, and Ling-Qi Yan. 2022. SpongeCake: A Layered Microflake Surface Appearance Model. *ACM Trans. Graph.* 42, 1 (2022), 1–16.
- Stephen H Westin, James R Arvo, and Kenneth E Torrance. 1992. Predicting reflectance functions from complex surfaces. In *Proceedings of SIGGRAPH*, Vol. 26.
- Hongzhi Wu, Julie Dorsey, and Holly Rushmeier. 2011. Physically-based interactive bi-scale material design. *ACM Trans. Graph.* 30, 6 (2011), 1–10.
- Kui Wu and Cem Yuksel. 2017. Real-time Fiber-level Cloth Rendering. In *Proceedings of I3D*. <https://doi.org/10.1145/3023368.3023372>
- Lifan Wu, Shuang Zhao, Ling-Qi Yan, and Ravi Ramamoorthi. 2019. Accurate appearance preserving prefiltering for rendering displacement-mapped surfaces. *ACM Trans. Graph.* 38, 4 (2019), 1–14.
- Chao Xu, Rui Wang, Shuang Zhao, and Hujun Bao. 2019. Multi-scale hybrid micro-appearance modeling and realtime rendering of thin fabrics. *IEEE transactions on visualization and computer graphics* 27, 4 (2019), 2409–2420.
- Kun Xu, Wei-Lun Sun, Zhao Dong, Dan-Yong Zhao, Run-Dong Wu, and Shi-Min Hu. 2013. Anisotropic Spherical Gaussians. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 32, 6 (2013), 209:1–209:11.
- Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. 2016. Position-Normal Distributions for Efficient Rendering of Specular Microstructure. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 35, 4 (2016), 56:1–56:9.
- 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 (2013), 1–12.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011a. Building Volumetric Appearance Models of Fabric Using Micro CT Imaging. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 30, 4 (2011), 44:1–44:10.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011b. Building volumetric appearance models of fabric using micro CT imaging. *ACM Trans. Graph.* 30, 4 (2011), 1–10.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-aware synthesis for predictive woven fabric appearance. *ACM Trans. Graph.* 31, 4 (2012), 1–10.
- Shuang Zhao, Fujun Luan, and Kavita Bala. 2016a. Fitting procedural yarn models for realistic cloth rendering. *ACM Trans. Graph.* 35, 4 (2016), 1–11.
- Shuang Zhao, Lifan Wu, Frédo Durand, and Ravi Ramamoorthi. 2016b. Downsampling Scattering Parameters for Rendering Anisotropic Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35, 6 (2016), 166:1–166:11.




Figure 7: Reproducing reflectance of real cloth samples. We compare our method (meso-scale BSDF) against the work from Jin et al. [2022], which equals to our micro-scale BSDF in this reflection only scene. Four examples are evaluated here from top to bottom: polyester satin (rotated 0°), polyester satin (rotated 90°), cotton twill and cotton jersey knit. Our method models shadowing and masking accurately, revealing its effect on the grazing angles for satin, and crisper contact shadow with smoother transition around the shadow terminator for twill and knit.