
Realistic Cloth Rendering with a Ray-Wave Hybrid Shading Model

Supplemental Document

Official supplemental document for the paper, *Realistic Cloth Rendering with a Ray-Wave Hybrid Shading Model*. In this document, we discuss a list of topics that we did not expand on in the main paper, and also present a more comprehensive set of fabric samples that we studied and rendered.

1 ASG Lobe Fitting

Recall from Section 5.2 of the main paper that our ply model fits analytical lobes to both the reflection-only and higher-order scattering distributions using the Anisotropic Spherical Gaussian (ASG) function introduced by Xu et al. (2013). The ASG is defined over unit vectors $\mathbf{v}$ on the unit sphere as:

$$G(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mu_1, \mu_2, c; \mathbf{v}) = c \cdot \max(\mathbf{v} \cdot \mathbf{z}, 0) \cdot e^{-\mu_1 (\mathbf{v}\cdot\mathbf{x})^2 - \mu_2 (\mathbf{v}\cdot\mathbf{y})^2} \tag{S.1}$$

Here, $\mathbf{z}$ denotes the lobe center direction, while $\mathbf{x}$ and $\mathbf{y}$ are the tangent and bitangent axes defining the anisotropic falloff. The parameters μ_1 and μ_2 control the bandwidth along $\mathbf{x}$ and $\mathbf{y}$, and c scales the lobe amplitude. In our implementation, we represent the frame $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$ using three scalars $(\theta_z, \phi_z, \phi_{xy})$:

$$\begin{aligned}
\mathbf{z} &= \langle \cos\theta_z \cos\phi_z, \cos\theta_z \sin\phi_z, \sin\theta_z \rangle \\
\mathbf{x_0} &= \langle -\sin\phi_z, \cos\phi_z, 0 \rangle \\
\mathbf{y_0} &= \langle -\sin\theta_z \cos\phi_z, -\sin\theta_z \sin\phi_z, \cos\theta_z \rangle \\
\mathbf{x} &= \cos\phi_{xy} \cdot \mathbf{x_0} - \sin\phi_{xy} \cdot \mathbf{y}_0 \\
\mathbf{y} &= \sin\phi_{xy} \cdot \mathbf{x_0} + \cos\phi_{xy} \cdot \mathbf{y}_0
\end{aligned} \tag{S.2}$$

1.1 Fitting Higher-Order Scattering from Non-Absorptive Plies

We now elaborate on the higher-order ASG fitting scheme described in the paper. A key consideration is that ASG mixture fitting should not be performed fully independently for each 2D slice $f_k^{\text{high}}(\theta_i, h)$. Although accurate fits can be obtained by optimizing multiple ASG lobes for each slice $f_k^{\text{high}}(\theta i^m, h^n; \omega_o)$, doing so independently introduces a new challenge: it becomes difficult to establish consistent correspondence across the L lobes fitted to different slices. In contrast to single-lobe fitting, multi-lobe mixtures lack an obvious one-to-one mapping across incident directions, making it harder to track how each lobe evolves with changing parameters when each fit is produced in isolation. Also note that in this section, indices such as m and n are plain superscripts rather than exponents.

To address the correspondence problem, we propose a more structured fitting strategy and begin by fitting ASG mixtures to f_0^{high}, the scattering distribution for non-absorptive fibers (we always use $\sigma_0 = 0$). Rather than fitting each 2D slice independently, we perform ASG mixture fitting over each block of $3 \times 3 = 9$ slices, which we call a "fitting block."

Let the discrete set of simulated incident parameters be $\theta_i^0, \theta_i^1, \ldots, \theta_i^{M-1}$ and $h^0, h^1, \ldots, h^{N-1}$, where $M = 2M' + 1$ and $N = 2N' + 1$ (we always choose odd numbers of samples for both θ_i and h). For fitting block indices $\tilde{m} = 0, 1, \ldots, M' - 1$ and $\tilde{n} = 0, 1, \ldots, N' - 1$, we define the fitting block $(\tilde{m}, \tilde{n})$ as the set of 9 slices featuring incident parameter pairs:

$$\begin{array}{ccc} (\theta_i^{2\tilde{m}+0}, h^{2\tilde{n}+0}) & (\theta_i^{2\tilde{m}+0}, h^{2\tilde{n}+1}) & (\theta_i^{2\tilde{m}+0}, h^{2\tilde{n}+2}) \\ (\theta_i^{2\tilde{m}+1}, h^{2\tilde{n}+0}) & (\theta_i^{2\tilde{m}+1}, h^{2\tilde{n}+1}) & (\theta_i^{2\tilde{m}+1}, h^{2\tilde{n}+2}) \\ (\theta_i^{2\tilde{m}+2}, h^{2\tilde{n}+0}) & (\theta_i^{2\tilde{m}+2}, h^{2\tilde{n}+1}) & (\theta_i^{2\tilde{m}+2}, h^{2\tilde{n}+2}) \end{array} \tag{S.3}$$

Within each block, we perform a large-scale ASG mixture optimization, jointly fitting a shared set of L lobes across all 9 slices. The goal is to find ASG parameters such that:

$$f_0^{\text{high}}(\theta_i^{2\tilde{m}+a}, h^{2\tilde{n}+b}; \omega_o) \approx \sum_{l=0}^{L-1} G_l(\mathbf{q}_{l,a,b}^{\tilde{m}\tilde{n}}; \omega_o) \quad \text{for } a, b = 0, 1, 2 \tag{S.4}$$

Here, $\mathbf{q}$ is a shorthand for the six ASG parameters $(\theta_z, \phi_z, \phi_{xy}, \mu_1, \mu_2, c)$. Fig. 1 (reproduced from the main paper) illustrates this block-based fitting strategy.

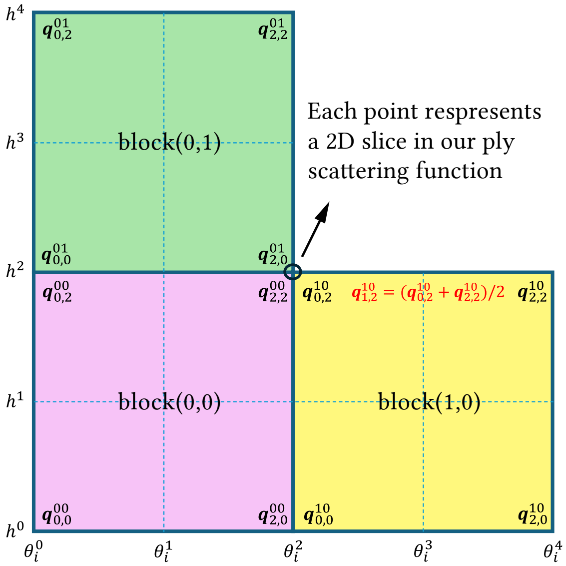

Figure 1: Illustration of fitting blocks, repeated from the main paper.

Among the ASG parameters in Eq. S.4, we treat $\mathbf{q}_{l,0,0}^{\tilde{m}\tilde{n}}$, $\mathbf{q}_{l,0,2}^{\tilde{m}\tilde{n}}$, $\mathbf{q}_{l,2,0}^{\tilde{m}\tilde{n}}$, and $\mathbf{q}_{l,2,2}^{\tilde{m}\tilde{n}}$ as variables to optimize. The remaining ASG lobe parameters within the block are then constrained via bilinear averaging:

$$\begin{aligned} \mathbf{q}_{l,0,1}^{\tilde{m}\tilde{n}} &= \frac{1}{2}\mathbf{q}_{l,0,0}^{\tilde{m}\tilde{n}} + \frac{1}{2}\mathbf{q}_{l,0,2}^{\tilde{m}\tilde{n}} \\ \mathbf{q}_{l,1,0}^{\tilde{m}\tilde{n}} &= \frac{1}{2}\mathbf{q}_{l,0,0}^{\tilde{m}\tilde{n}} + \frac{1}{2}\mathbf{q}_{l,2,0}^{\tilde{m}\tilde{n}} \\ \mathbf{q}_{l,1,1}^{\tilde{m}\tilde{n}} &= \frac{1}{4}\mathbf{q}_{l,0,0}^{\tilde{m}\tilde{n}} + \frac{1}{4}\mathbf{q}_{l,0,2}^{\tilde{m}\tilde{n}} + \frac{1}{4}\mathbf{q}_{l,2,0}^{\tilde{m}\tilde{n}} + \frac{1}{4}\mathbf{q}_{l,2,2}^{\tilde{m}\tilde{n}} \\ \mathbf{q}_{l,1,2}^{\tilde{m}\tilde{n}} &= \frac{1}{2}\mathbf{q}_{l,0,2}^{\tilde{m}\tilde{n}} + \frac{1}{2}\mathbf{q}_{l,2,2}^{\tilde{m}\tilde{n}} \\ \mathbf{q}_{l,2,1}^{\tilde{m}\tilde{n}} &= \frac{1}{2}\mathbf{q}_{l,2,0}^{\tilde{m}\tilde{n}} + \frac{1}{2}\mathbf{q}_{l,2,2}^{\tilde{m}\tilde{n}} \end{aligned} \tag{S.5}$$

This setup ensures that the ASG mixtures fitted to the four corner slices—indexed by $(0,0)$, $(0,2)$, $(2,0)$, and $(2,2)$—are fully optimized, while the mixtures for the five interior slices are defined as convex combinations of those corners. As a result, each fitting block yields a smoothly varying, piecewise-interpolatable representation of f_0^{high}.

To evaluate f_0^{high} at any unsimulated incident parameter pair (θ_i^*, h^*) that falls within a block $(\tilde{m}, \tilde{n})$, we first compute interpolation weights $r, s \in [0,1]$ such that

$$\begin{aligned} \theta_i^* &= (1-r) \cdot \theta_i^{2\tilde{m}+0} + r \cdot \theta_i^{2\tilde{m}+2} \\ h^* &= (1-s) \cdot h^{2\tilde{n}+0} + s \cdot h^{2\tilde{n}+2} \end{aligned} \tag{S.6}$$

We then define for each lobe $l = 0, 1, ..., L-1$:

$$\mathbf{q}_l^* = (1-r)(1-s) \cdot \mathbf{q}_{l,0,0}^{\tilde{m}\tilde{n}} + (1-r)s \cdot \mathbf{q}_{l,0,2}^{\tilde{m}\tilde{n}} + r(1-s) \cdot \mathbf{q}_{l,2,0}^{\tilde{m}\tilde{n}} + rs \cdot \mathbf{q}_{l,2,2}^{\tilde{m}\tilde{n}} \tag{S.7}$$

and approximate the scattering distribution as:

$$f_0^{\text{high}}(\theta_i^*, h^*; \omega_o) \approx \sum_{l=0}^{L-1} G_l(\mathbf{q}_l^*; \omega_o) \tag{S.8}$$

Since the ASG parameters encode each lobe's orientation, sharpness, and amplitude, interpolating them naturally results in smoothly evolving scattering lobes across incident configurations. This is precisely the behavior we seek for modeling higher-order scattering.

1.2 Fitting Absorptive Ply Scattering Functions

Our per-block ASG mixture fitting for f_0^{high} enables a much more compact representation of the original 4D raw data, while also supporting interpolation through smooth variation of ASG parameters. This framework further extends to the high-order scattering data for absorptive fibers—$f_1^{\text{high}}, f_2^{\text{high}}, \dots, f_{K-1}^{\text{high}}$—by reusing the ASG lobe parameters $\mathbf{q}_{l,0,0}^{\tilde{m}\tilde{n}}, \mathbf{q}_{l,0,2}^{\tilde{m}\tilde{n}}, \mathbf{q}_{l,2,0}^{\tilde{m}\tilde{n}}, \mathbf{q}_{l,2,2}^{\tilde{m}\tilde{n}}$ already optimized from f_0^{high}.

The data $f_1^{\text{high}}, ..., f_{K-1}^{\text{high}}$ are strongly correlated with f_0^{high}—they describe the same simulated ply instances, differing only in the level of energy absorption within the fibers. As absorption increases, less energy remains in the higher-order scattering, and the overall magnitude of f_k^{high} diminishes. For the largest absorption coefficient, f_{K-1}^{high} contains near-zero values throughout.

In practice, we find that these scattering distributions can be accurately approximated using the same ASG lobes fitted to f_0^{high}, if we rescale their amplitudes appropriately. Specifically, for each absorption index k and fitting block $(\tilde{m}, \tilde{n})$, we model:

$$f_k^{\text{high}}(\theta_i^{2\tilde{m}+a}, h^{2\tilde{n}+b}; \omega_o) \approx \sum_{l=0}^{L-1} s_{l,k}^{\tilde{m}\tilde{n}} \cdot G_l(\mathbf{q}_{l,a,b}^{\tilde{m}\tilde{n}}; \omega_o) \quad \text{for } a, b = 0, 1, 2 \tag{S.9}$$

The motivation behind this model design becomes clear when comparing Eq. S.4 and Eq. S.9. In Eq. S.4, each 2D slice of the non-absorptive data—$f_0^{\text{high}}(\theta_i^{2\tilde{m}+a}, h^{2\tilde{n}+b}; \omega_o)$—is approximated as a sum of L ASG scattering lobes. We assume that increasing fiber absorption does not substantially alter the overall multi-lobe structure, but primarily reduces the energy in each lobe. To capture this, we introduce scale-down parameters $sl, k^{\tilde{m}\tilde{n}}$ that modulate lobe amplitudes based on absorption.

In particular, we have $s_{l,0}^{\tilde{m}\tilde{n}} = 1$ for all l and block indices $(\tilde{m}, \tilde{n})$, corresponding to the baseline case of zero absorption. For each $k \geq 1$, we estimate $s_{l,k}^{\tilde{m}\tilde{n}}$ by fitting to f_k^{high}, while keeping the ASG lobe parameters $\mathbf{q}_{l,a,b}^{\tilde{m}\tilde{n}}$ fixed from the earlier optimization on f_0^{high} (see Eq. S.9).

These scale-down parameters allow us to interpolate high-order scattering functions across arbitrary absorption levels. Given incident parameters (θ_i^*, h^*) satisfying Eq. S.6 and an intermediate absorption coefficient σ^* such that

$$\sigma^* = (1 - t) \cdot \sigma_k + t \cdot \sigma_{k+1} \tag{S.10}$$

for some pair of simulated absorption coefficients σ_k and σ_{k+1}, we proceed as follows—

First, we compute the baseline ASG lobe parameters $\mathbf{q}_l^*$ using Eq. S.7.

Then, interpolate the scale-down weights:

$$s_l^* = (1 - t) \cdot s_{l,k}^{\tilde{m}\tilde{n}} + t \cdot s_{l,k+1}^{\tilde{m}\tilde{n}} \quad \text{for } l = 0, 1, \ldots, L - 1 \tag{S.11}$$

Finally, the scattering function for the desired triplet $(\theta_i^*, h^*, \sigma^*)$ is approximated as:

$$f^{\text{high}}(\theta_i^*, h^*, \sigma^*; \omega_o) \approx \sum_{l=0}^{L-1} s_l^* \cdot G_l(\mathbf{q}_l^*; \omega_o) \tag{S.12}$$

This formulation enables a fully analytical, smoothly interpolated representation of higher-order scattering over the entire space of incident angles and fiber absorption coefficients.

2 Multi-Ply Scattering

As described in Section 5.3 of the main paper, when a yarn containing multiple plies is illuminated by an incident beam with parameters (θ_i, h), the primary ply directly facing the light scatters according to $f^{\text{single}}(\theta_i, h, \sigma; \omega_o)$. Some of this scattered light, however, may strike neighboring plies before reaching the far field, where it can be redirected or absorbed. To capture these effects, our multi-ply scattering model decomposes the outgoing signal into three components:

- Direct scattering: unoccluded light escaping from the primary ply.
- Single secondary interaction: light scattered once more by a secondary ply before exiting.
- Residual multiple scattering: the remaining energy from higher-order inter-ply interactions, modeled as a diffuse term.

Formally, for a yarn with P plies, we write:

$$f^{\text{multi}}(\theta_i, h, \sigma; \omega_o) = f^{\text{direct}}(\theta_i, h, \sigma; \omega_o) + \sum_{p=1}^{P-1} f_p^{\text{conv}}(\theta_i, h, \sigma; \omega_o) + f^{\text{resi}}(\theta_i, h, \sigma; \omega_o) \tag{S.13}$$

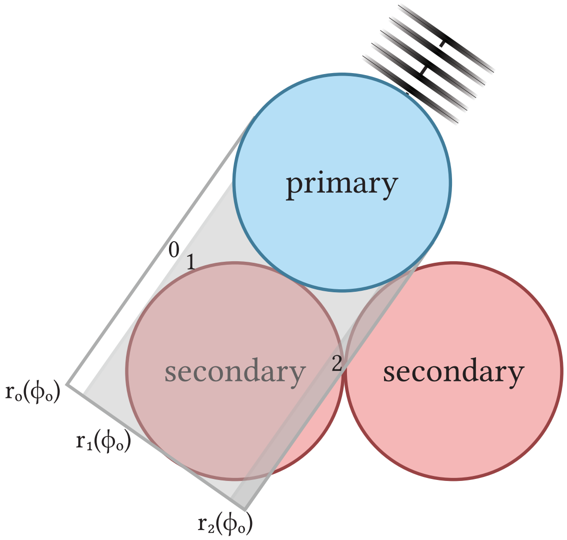

Figure 2: The ply illuminated by an incident beam is the primary ply, and other plies, which may rescatter the outgoing light from the primary ply, are secondary plies. Each secondary ply rescatters a different amount of energy.

This single-to-multi extension must be evaluated across all incident conditions (θ_i, h) as well as different absorption coefficients. As detailed in the paper, we begin by simulating 2D shadowing between plies, illustrated in Fig. 2. For each azimuthal angle ϕ_o, we trace rays originating from the primary ply (indexed as 0) and record the fraction $r_0(\phi_o)$ that escapes directly, along with fractions $r_p(\phi_o)$ that reach each secondary ply p. The direct contribution is then defined as:

$$f^{\text{direct}}(\theta_i, h, \sigma; \omega_o) = r_0(\phi_o) \cdot f^{\text{single}}(\theta_i, h, \sigma; \omega_o) \tag{S.14}$$

We now turn to the detailed computation of the other two terms in Eq. S.13.

2.1 Single Secondary Interaction

To approximate indirect scattering contributed by a secondary ply p, we simulate a process that effectively convolves the single-ply scattering functions of ply 0 (primary) and ply p. Our computation is as follows:

(1) Initialize a scattering function $\overline{f_p^{\text{conv}}}(\theta_i, h, \sigma; \omega_o)$ with zeros, where ω_o is measured in the primary ply's coordinate system.

(2) Loop over a dense set of sampled directions $\omega_m \sim (\theta_m, \phi_m)$. Each ω_m is an outgoing direction with respect to ply 0 and simultaneously an incoming direction for ply p. For each ω_m:

  (2.1) Use $r_p(\phi_m)$ (see main paper) to estimate the fraction of energy leaving ply 0 along ω_m that reaches ply p.

  (2.2) Model the secondary scattering from ply p with our single-ply formulation. We denote it by $\hat{f}_p^{\text{single}}(\omega_m; \omega_o)$, expressed in ply 0's coordinates. The mapping from ω_m to the corresponding incident angle on ply p depends on ply orientation. The precise incident position on ply p is determined by our 2D intersection test of rays from ply 0, but is omitted in the notation for clarity. (The dependence on absorption is likewise suppressed for brevity.)

(2.3) Perform careful coordinate transformation when computing $\hat{f}_p^{\text{single}}$, as we eventually need to express it with respect to directions in ply 0's coordinate system. This accounts for relative rotations and ply twists.

(2.4) Accumulate the contribution from ω_m into $\overline{f_p^{\text{conv}}}(\theta_i, h, \sigma; \omega_o)$, proportional to $f^{\text{single}}(\theta_i, h, \sigma; \omega_m) \cdot r_p(\phi_m) \cdot \hat{f}_p^{\text{single}}(\omega_m; \omega_o)$.

(3) The result $\overline{f_p^{\text{conv}}}(\theta_i, h, \sigma; \omega_o)$ represents the total contribution of secondary scattering from ply p, given primary scattering from ply 0 under incident parameters (θ_i, h).

(4) To obtain the portion that escapes to the far field, we again perform 2D shadow tests from ply p, analogous to the procedure used for the direct term. This yields $f_p^{\text{conv}}(\theta_i, h, \sigma; \omega_o)$, which represents the unoccluded secondary scattering from ply p.

This convolution-style procedure captures the additional scattering that arises when light scattered by the primary ply undergoes a single secondary interaction with its neighbors.

2.2 Residual Multiple Scattering

Since $f_p^{\text{conv}}(\theta_i, h, \sigma; \omega_o)$ models only the secondary scattering from ply p that reaches the far field without further occlusion, we must also include a residual term $f^{\text{resi}}(\theta_i, h, \sigma; \omega_o)$ in Eq. S.13 to capture higher-order inter-ply interactions. Without this term, our model would systematically lose energy. To compute f^{resi}, we define three scalar quantities derived from the previously introduced scattering functions:

$$\begin{aligned} s_1(\theta_i, h, \sigma) &= \int_{S^2} [f^{\text{single}}(\theta_i, h, \sigma; \omega_o) - f^{\text{direct}}(\theta_i, h, \sigma; \omega_o)] \, d\omega_o \\ s_2(\theta_i, h, \sigma) &= \sum_{p=1}^{P-1} \int_{S^2} \overline{f_p^{\text{conv}}}(\theta_i, h, \sigma; \omega_o) \, d\omega_o \\ s_3(\theta_i, h, \sigma) &= \sum_{p=1}^{P-1} \int_{S^2} f_p^{\text{conv}}(\theta_i, h, \sigma; \omega_o) \, d\omega_o \end{aligned} \tag{S.15}$$

Here, s_1 is the total energy from ply 0 that interacts with other plies, s_2 is the total secondary scattering from all plies before accounting for occlusions, and s_3 is the portion of that energy that escapes directly to the far field. We then define the fractions:

$$s_{\text{absorb}} = 1 - \frac{s2}{s1}; \quad s_{\text{farfield}} = \frac{s3}{s1}; \quad s_{\text{rescatter}} = \frac{s2 - s3}{s1} \tag{S.16}$$

which satisfiy $s_{\text{absorb}} + s_{\text{farfield}} + s_{\text{rescatter}} = 1$.

Assuming higher-order scattering can be modeled as an infinite process, the eventual fraction of energy (originating from s_1) that propagates to the far field is:

$$s_1 \cdot (s_{\text{farfield}} + s_{\text{rescatter}} \cdot s_{\text{farfield}} + s_{\text{rescatter}}^2 \cdot s_{\text{farfield}} + ...) = s1 \cdot \frac{s_{\text{farfield}}}{1 - s_{\text{rescatter}}} \tag{S.17}$$

Because the single secondary interaction terms integrate to s_3, the residual term must integrate to

$$s1 \cdot \frac{s_{\text{farfield}}}{1 - s_{\text{rescatter}}} - s_3 \tag{S.18}$$

to approximately conserve energy. We model f^{resi} as a broad, diffuse band:

$$f^{\text{resi}}(\theta_i, h, \sigma; \omega_o) = c(\theta_i, h, \sigma) \cdot f_1(\theta_i; \theta_o) \cdot f_2(\phi_o) \tag{S.19}$$

where f_1 is a Gaussian centered at $-\theta_i$, f_2 is uniform ($f_2(\phi_o) = 1/2\pi$), and $c(\theta_i, h, \sigma)$ is chosen so that f^{resi} integrates to Eq. S.18.

This energy accounting is approximate but reasonable, and it holds in special cases. For example, when fibers are non-absorptive ($\sigma = 0$), we always obtain $s_{\text{absorb}} = 0$ with $s_{\text{farfield}} + s_{\text{rescatter}} = 1$. In this case, Eq. S.17 reduces to s_1, ensuring that all light scattered from ply 0 and subsequently rescattered by other plies eventually reaches the far field without loss.

3 Double-Layer Scattering

In this section, we present our double-layer scattering modeling omitted from Section 5.4 of the main paper (due to page limit). This component in our model, which accounts for scattering between the two yarn layers in a woven cloth, is essential for characterizing the overall transmission through a fabric.

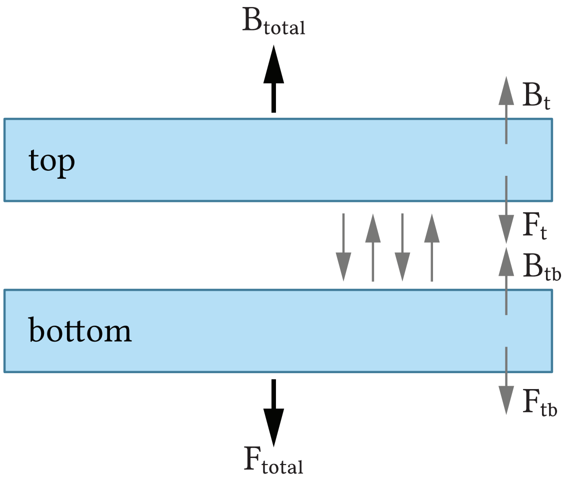

Figure 3: Illustration of our double-layer modeling. The total backward and forward scattering from a double-layer structure can be computed using estimated reflection and transmission factors of each layer, with inter-reflections modeled by a geometric sequence.

As shown in Fig. 3, a fabric is locally treated as a two-layer stack. Each layer has its corresponding BSDF, defined by averaging f^{single} or f^{multi} (depending on the number of plies) over a full range of incident positions h. Integrating away the incident position dependence results in two classic BSDFs $f_t(\xi_i, \xi_o)$ and $f_b(\xi_i, \xi_o)$, for the top and bottom layers. Both f_t and f_b are defined with respect to the fabric surface coordinate system, and we use ξ_i, ξ_o to denote surface-space directions as opposed to ply-space directions ω_i, ω_o. Note that since f_t and f_b are integrated from the same single-ply or multi-ply model, they should be identical up to some rotation, pertaining to the orientations of warps and wefts.

We then estimate a reflection factor and a transmission factor for each of these two layers, by simulating two steps of scattering through this double-layer stack. We compute these reflection and transmission factors as a function of the initial incident direction ξ_i on the top layer:

- For each ξ_i, we illuminate the top layer from this direction and first record the total amount of energy exiting backward and forward, $B_t(\xi_i)$ and $F_t(\xi_i)$. These quantities are computed by integrating $f_t(\xi_i, \cdot)$ in the upper and lower hemispheres.
- Light that proceeds along the forward directions interacts with the bottom layer, which scatters this light according to f_b. The resulting total backward and forward scattered light are recorded as $B_b(\xi_i)$ and $F_b(\xi_i)$.
- Since only light that propagates forward relative to the top layer can interact with the bottom layer, we compute the reflection and transmission factors of the bottom layer as $B_{tb}(\xi_i) = B_b(\xi_i)/F_t(\xi_i)$ and $F_{tb}(\xi_i) = F_b(\xi_i)/F_t(\xi_i)$.

We consider the scattering in between the top and bottom layers as an infinite process, as shown in Fig. 3. As an approximation, we assume that all further bounces in the inter-layer scattering share the same reflection and transmission factors as our characterized second step of scattering, i.e. B_{tb} and F_{tb}. The total backward and forward scattering from the double-layer structure, in response to the initial incident illumination from ξ_i, can be modeled through a geometric series. Specifically, we have:

$$\begin{aligned} B_{\text{total}}(\xi_i) &= B_t(\xi_i) + B_{tb}(\xi_i)\frac{F_t(\xi_i)F_{tb}(\xi_i)}{1 - B_{tb}^2(\xi_i)} \\ F_{\text{total}}(\xi_i) &= \frac{F_t(\xi_i)F_{tb}(\xi_i)}{1 - B_{tb}^2(\xi_i)} \end{aligned} \tag{S.20}$$

For rendering, we tabulate B_t, F_t, B_{total}, and F_{total} as functions of ξ_i. For each shading point, if our incident beam, incoming from ξ_i with respect to the fabric surface, intersects with yarns in both layers, we correct the backward or forward radiance, first computed from the ply model, using these tabulated factors. For backward scattering, we compute

$$L_b \to \frac{B_{\text{total}}(\xi_i)}{B_t(\xi_i)} \cdot L_b \tag{S.21}$$

Note that we always have $B_{\text{total}} > B_t$, as the double-layer structure increases backward scattering.

For forward scattering, we use a linear combination of two models with some blending factor s:

$$L_f \to (1 - s) \cdot \frac{F_{\text{total}}(\xi_i)}{F_t(\xi_i)} \cdot L_f + s \cdot F_{\text{total}}(\xi_i) \cdot L_d \tag{S.22}$$

where $F_{\text{total}}(\xi_i) \cdot L_d$ represents a diffuse BTDF that integrates to $F_{\text{total}}(\xi_i)$ in the lower hemisphere. We use this diffuse BTDF term to model the loss of directionality in light, after transmitting through two layers of warps and wefts. Note that $F_{\text{total}} < F_t$, and when the yarns are absorptive, this decrease in the total transmission can be significant.

4 Diffraction Modeling

In this section, we include more derivations and explanations on our diffraction modeling, presented in Section 6.1 of our main paper.

4.1 Model Derivation

We first provide a full, detailed derivation of Eq. 11 and Eq. 12 in the main paper, leveraging the original Kirchhoff diffraction integral Heurtley (1973).

The most classic and general form of the Kirchhoff diffraction integral is given by:

$$U(P) = \frac{1}{4\pi} \int_S U_0 \frac{\partial}{\partial \mathbf{n}} (\frac{e^{iks}}{s}) - \frac{e^{iks}}{s} \frac{\partial U_0}{\partial \mathbf{n}} dS \tag{S.23}$$

which computes the diffracted field at P, caused by an incident field U_0 that diffracts through an aperture S. Here $\mathbf{n}$ is the normal of the aperture, pointing to the side that aligns with the incident field's propagation direction. Eq. S.23 integrates over all the points on the aperture S, and s is the distance between each aperture point and the receiving point P. Moreover, k is the wavenumber of the field, given by $k = 2\pi/\lambda$, where λ is the wavelength.

Recall from our main paper that we formulate our diffraction modeling with respect to a rectangular aperture, defined by $[x_w, x_e] \times [y_s, y_n]$ in the xy-plane of a local coordinate system, as shown in Fig. 4.

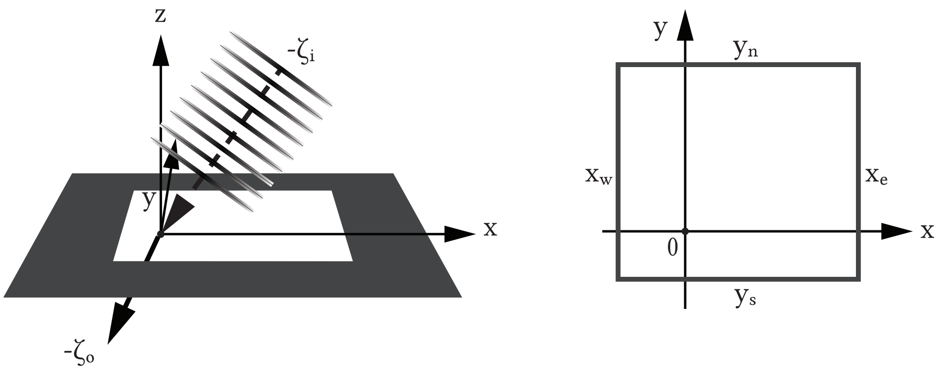

Figure 4: Rectangular aperture and coordinate system.

Thus, for our context, we can rewrite Eq. S.23 as:

$$U(\mathbf{P}) = \frac{1}{4\pi} \int_{x_w}^{x_e} \int_{y_s}^{y_n} U_0(x,y,0) \nabla (\frac{e^{iks}}{s}) \cdot (-\hat{\mathbf{z}}) - \frac{e^{iks}}{s} \nabla U_0(x,y,0) \cdot (-\hat{\mathbf{z}}) \, dydx \tag{S.24}$$

Note that the incident field propagates toward the aperture from above the xy-plane, making $\mathbf{n} = -\mathbf{z}$ in our case. Also, we have used the definition $\frac{\partial f}{\partial \mathbf{n}} = \nabla f \cdot \mathbf{n}$.

Our Gaussian beam incident field U_0 is approximated near the surface as:

$$U_0(x,y,z) = e^{-(x^2+y^2)/w_0^2} \cdot e^{ik(u_x \cdot x + u_y \cdot y + u_z \cdot z)} \tag{S.25}$$

i.e. a Gaussian-windowed plane wave, where w_0 is the beam waist. Here we use a different notation as in the paper, denoting the incident field direction with the unit vector $\mathbf{u} = \langle u_x, u_y, u_z \rangle$, where $u_z < 0$. Note that $\mathbf{u} = -\xi_i$ in the main paper.

Moreover, as we are interested in the diffracted field quantity along a direction $\mathbf{v} = \langle v_x, v_y, v_z \rangle$ in the far field limit, we take the receiving point P as $P = R\mathbf{v}$, for very large R.

For each point on the aperture $Q = (x, y, 0)$, the distance s is given by $s = |Q - P|$, and the quantity $\nabla(\frac{e^{iks}}{s})$, which takes the gradient w.r.t. Q, can be derived from vector calculus as:

$$\nabla(\frac{e^{iks}}{s}) = \frac{(iks-1)e^{iks}}{s^3}(Q-P) \tag{S.26}$$

Plugging in $Q = (x, y, 0)$ and $P = (Rv_x, Rv_y, Rv_z)$, we have:

$$\nabla(\frac{e^{iks}}{s}) \cdot (-\hat{\mathbf{z}}) = Rv_z \frac{(iks-1)e^{iks}}{s^3} \tag{S.27}$$

Moreover, we have $\nabla U_0(x, y, 0) \cdot (-\hat{\mathbf{z}}) = -iku_z \cdot U_0$. Combining into Eq. S.24, we obtain:

$$\begin{aligned} U(R\mathbf{v}) = \frac{1}{4\pi} \int_{x_w}^{x_e} \int_{y_s}^{y_n} e^{-\frac{x^2+y^2}{w_0^2}} e^{ik(u_x \cdot x + u_y \cdot y)} Rv_z \frac{(iks-1)e^{iks}}{s^3} \\ - \frac{e^{iks}}{s} \cdot (-iku_z) \cdot e^{-\frac{x^2+y^2}{w_0^2}} e^{ik(u_x \cdot x + u_y \cdot y)} \, dydx \end{aligned} \tag{S.28}$$

Since R is very large, we have $ks >> 1$. Thus, we can tighten the above formula as:

$$U(R\mathbf{v}) \approx \frac{1}{4\pi} \int_{x_w}^{x_e} \int_{y_s}^{y_n} e^{-\frac{x^2+y^2}{w_0^2}} e^{ik(u_x \cdot x + u_y \cdot y)} \cdot [Rv_z \frac{ike^{iks}}{s^2} + iku_z \frac{e^{iks}}{s}] \, dydx \tag{S.29}$$

Moreover, for amplitude variation, we can approximate s as $s \approx R$, allowing for a further cancellation in the s^2 term in a denominator and giving rise to:

$$\begin{aligned} U(R\mathbf{v}) &\approx \frac{1}{4\pi} \int_{x_w}^{x_e} \int_{y_s}^{y_n} e^{-\frac{x^2+y^2}{w_0^2}} e^{ik(u_x \cdot x + u_y \cdot y)} \cdot ik \frac{e^{iks}}{s} (u_z + v_z) \, dydx \\ &\approx \frac{ik}{4\pi R} \int_{x_w}^{x_e} \int_{y_s}^{y_n} (u_z + v_z) e^{-\frac{x^2+y^2}{w_0^2}} e^{ik(u_x \cdot x + u_y \cdot y)} e^{iks} \, dydx \end{aligned} \tag{S.30}$$

Furthermore, for approximating phase variation Gibson (2021), we have:

$$e^{iks} \approx e^{ikR} \cdot e^{-ik(v_x \cdot x + v_y \cdot y)} \tag{S.31}$$

Combining the two equations above finally gives us:

$$U(R\mathbf{v}) \approx \frac{ik}{4\pi R} e^{ikR} \int_{x_w}^{x_e} \int_{y_s}^{y_n} (u_z + v_z) e^{-\frac{x^2+y^2}{w_0^2}} e^{ik[(u_x - v_x) \cdot x + (u_y - v_y) \cdot y]} \, dydx \tag{S.32}$$

The unit vectors ξ_i, ξ_o introduced in the main paper, which both point away from the aperture following the rendering convention, can be related to $\mathbf{u}, \mathbf{v}$ via $\xi_i = -\mathbf{u}$ and $\xi_o = \mathbf{v}$. Also noting that $u_z, v_z < 0$, as well as $k = 2\pi/\lambda$, we rewrite Eq. S.32 as:

$$U(R\xi_o) \approx \left[-\frac{ie^{i\frac{2\pi}{\lambda}R}}{R}\right]\frac{1}{2\lambda}(|\xi_i^z| + |\xi_o^z|)\int_{x_w}^{x_e}\int_{y_s}^{y_n} e^{-\frac{x^2+y^2}{w_0^2}} \cdot e^{-2\pi i(\frac{\xi_i^x+\xi_o^x}{\lambda}\cdot x + \frac{\xi_i^y+\xi_o^y}{\lambda}\cdot y)}\, dydx \tag{S.33}$$

Denoting the aperture as $A = (x_w, x_e, y_s, y_n)$ and defining $f_x = (\xi_i^x + \xi_o^x)/\lambda$, $f_y = (\xi_i^y + \xi_o^y)/\lambda$ as in the main paper, we introduce

$$K(A; f_x, f_y) = \int_{x_w}^{x_e}\int_{y_s}^{y_n} e^{-\frac{x^2+y^2}{w_0^2}} \cdot e^{-2\pi i(f_x\cdot x + f_y\cdot y)}\, dydx. \tag{S.34}$$

which is identical to Eq. 12 in the main paper.

Finally, for rendering, we need radiometric intensity rather than field amplitude, so we omit the term in square brackets in Eq. S.33 and compute the squared magnitude of the rest. Using the definition in Eq. S.34, we arrive at the diffracted intensity in Eq. 11 of the main paper:

$$I(A, \lambda, \xi_i; \xi_o) = \frac{1}{4\lambda^2}(|\xi_i^z| + |\xi_o^z|)^2 K^2(A, f_x, f_y) \tag{S.35}$$

4.2 Aperture Position and Diffraction Pattern Color

As mentioned in the main paper, we find that with the same incident beam, apertures of the same size give rise to diffraction patterns with similar structures, and the relative position of the beam center to aperture edges mainly changes the pattern's color saturation, with the saturation decreasing as the beam center approaches the edge of the aperture. We illustrate this phenomenon with some examples.

Given a fixed beam waist of $w_0 = 10\mu m$, we consider two apertures of size $10\mu m \times 10\mu m$:

$$A_1 = (-5, 5, -5, 5); \qquad A_2 = (0, 10, 0, 10) \tag{S.36}$$

In other words, the incident beam center is aligned with the center of A_1, but passes through the corner of A_2. The diffraction patterns of the beam through these apertures have very similar shapes, but different details in the color patterns, as A_1 gives rise to a much more saturated color pattern.

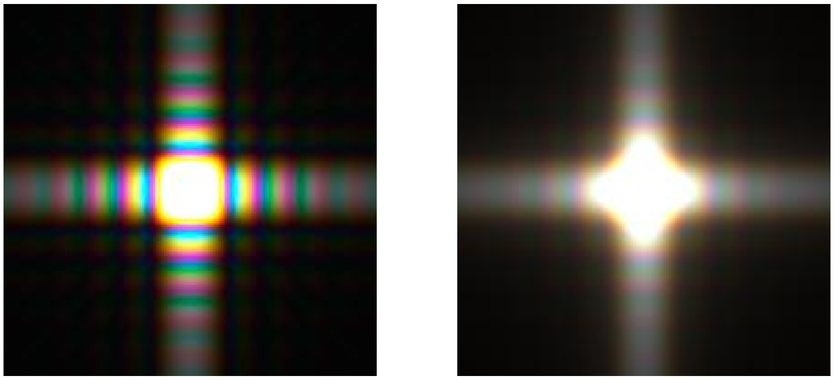

Figure 5: Patterns from apertures of the same size but different positions relative to the beam center.

4.3 Ply-Edge Diffraction Aperture

We also provide a simple sketch that illustrates Eq. 17 from our main paper, on the aperture A related to ply-edge diffraction—using a positive incident position h as an example, which satisfies $R - d_{\text{cut}} \leq h < R$.

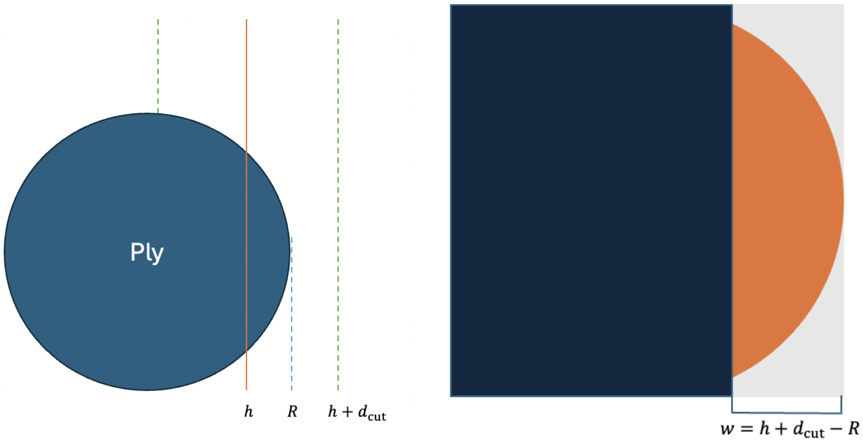

Figure 6: A simple sketch illustrating ply-edge diffraction.

As shown in Fig. 6, our incident beam center intersects the ply at position h ($h > 0$). With d_{cut} being the intensity die-off-to-zero distance of the beam, the right boundary of the beam is at position $h + d_{\text{cut}}$, relative to the ply center. Moreover, with the ply radius given by R, the beam is occluded by the ply for $x <= R$, and becomes unoccluded for $R < x < h + d_{\text{cut}}$. Thus, the aperture width along the x-direction is given by $h + d_{\text{cut}} - R$, and the aperture is considered unbounded along the y-direction.

Transforming to the aperture space where the beam center is located at the origin, since the right side of the beam is unoccluded, we have $x_e = d_{\text{cut}}$. With the aperture width given by $h + d_{\text{cut}} - R$, a simple subtraction gives us $x_w = R - h$. Also, as the aperture can be considered unbounded along the y-direction, y_s and y_n are set to the minimum and maximum values, $y_s = -d_{\text{cut}}$ and $y_n = d_{\text{cut}}$.

5 Wave Optics Noise Functions

In this section, we hope to expand on Section 6.2 in our main paper, providing more insights into our wave optics based noise functions for improving realism in appearance.

5.1 Designing 2D Noise Functions

As mentioned in our paper, we need to build upon our 1D noise functions and generalize our noise model to 2D angular domains—over outgoing directions on the unit sphere—so it can be applied to our full ply scattering model. This is achieved by reusing the 1D power spectral density (PSD) function discussed in the main paper, which we denote as $P_{1D}(\lambda; f_\phi)$ here. Specifically, under an assumption of isotropy, we construct our 2D PSD function $P_{2D}(\lambda; f_\theta, f_\phi)$ from $P_{1D}(\lambda; f_\phi)$ by setting $P_{2D}(\lambda; f_\theta, f_\phi) =$

$P_{1D}(\lambda; \sqrt{f_\theta^2 + f_\phi^2})$ and then renormalizing. We can then generate $r_{2D}^{(n)}(\lambda, \theta_o, \phi_o), i_{2D}^{(n)}(\lambda, \theta_o, \phi_o)$ that satisfy

$$r_{2D}^{(n)}(\lambda, \theta_o, \phi_o),\ i_{2D}^{(n)}(\lambda, \theta_o, \phi_o) \sim \mathcal{N}(0, 1) \tag{S.37}$$

The 2D noise function still takes the form

$$N_{2D}^{(n)}(\lambda, \theta_o, \phi_o) = \frac{1}{2}[r_{2D}^{(n)}(\lambda, \theta_o, \phi_o)]^2 + \frac{1}{2}[i_{2D}^{(n)}(\lambda, \theta_o, \phi_o)]^2 \tag{S.38}$$

and therefore satisfies $\mathbb{E}[N_{2D}^{(n)}(\lambda, \theta_o, \phi_o)] = 1$ at each point.

Our raw 2D noise function is isotropic with respect to θ_o and ϕ_o; however, as $\omega_o \sim (\theta_o, \phi_o)$ represents a direction on the unit sphere, the range of ϕ_o, for larger $|\theta_o|$, corresponds to a smaller cone of directions, causing our isotropic (θ_o, ϕ_o) pattern to be compressed along the ϕ direction near the poles. To adjust for this, we pre-filter our 2D noise tiles before using them in rendering, adopting a filter size along ϕ_o that is proportional to $\frac{1}{\cos\theta_o}$. Fig. 7 compares our unfiltered and filtered 2D noise, with the filtered noise function becoming significantly smoother for large $|\theta_o|$ values (top and bottom rows).

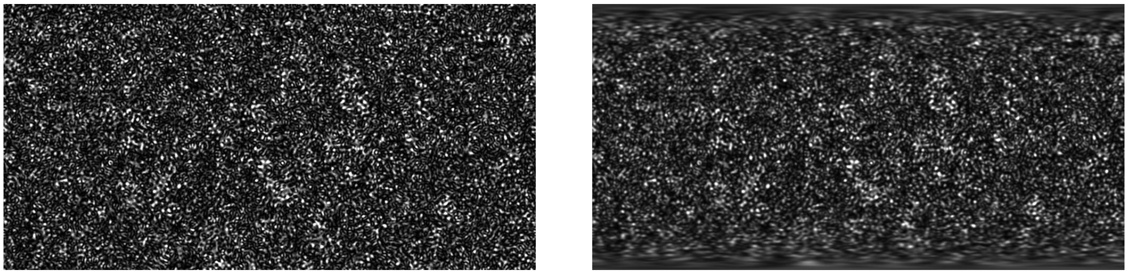

Figure 7: Example unfiltered vs. filtered 2D noise tiles generated using our method.

For rendering, we generate two sets of noise tiles $N_{2D,r}^{(n)}(\lambda, \omega_o)$, $N_{2D,h}^{(n)}(\lambda, \omega_o)$ for separately modeling reflection and higher-order scattering, since these distributions can have very different statistics.

5.2 Noise Application and Memory Effect

To apply our noise model in shading, we need a scheme for evaluating appropriate noise values, for a beam hitting the fabric at any point (spatial variation) and for any incident and outgoing directions (angular variation). As mentioned in the main paper, we handle spatial variation simply by using different noise instances for different positions on the fabric, through a hash function that depends on the ply ID and the shading point's position on the ply.

To augment a ply scattering function with our noise model, we need to properly query noise values in $N_{2D,r}^{(n)}$ and $N_{2D,h}^{(n)}$ and multiply them to our ply scattering functions $f^{\text{refl}}(\theta_i, h; \omega_o)$ and $f^{\text{high}}(\theta_i, h, \sigma; \omega_o)$. for arbitrary $(\theta_i, h, \sigma, \omega_o)$. Note that the absorption coefficient σ, unrelated to spatial or angular variations, has no influence on noise function querying, and the influence of h is handled through using different noise function instances. The influence of θ_i is different for surface reflection and higher-order scattering.

As shown in Fig. 8, the reflection distributions from the same fiber under different incident angles (0° and 15°) are very strongly correlated, while the high-order scattering distributions do not indicate clear correlations—this is because high-order scattering can involve complex light paths within fibers, and even gently changing the incident angle could significantly modify the sequences of scattering events. Thus, we only model this correlation between speckle patterns under different illumination directions, known as the memory effect Xia et al. (2023), for the surface reflection component.

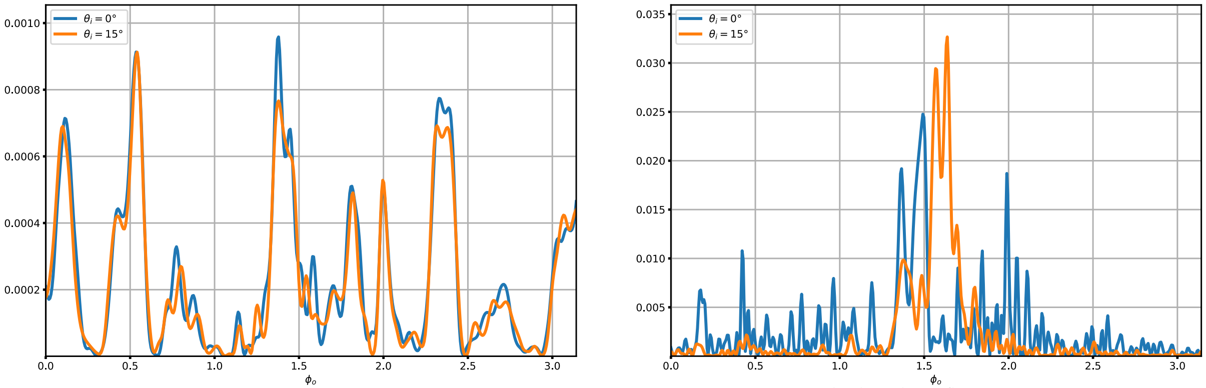

surface reflection

high-order reflection

Figure 8: Memory effect (only) in surface reflection distributions.

Specifically, our noise application is realized by multiplying the average scattering function values with queried noise values:

$$\begin{aligned} f^{(n)}_{\text{refl}}(\lambda, \theta_i, h; \omega_o) &= f^{\text{refl}}(\theta_i, h; \omega_o) \cdot N^{(n)}_{2D,r}(\lambda, \omega'_o) \\ f^{(n)}_{\text{high}}(\lambda, \theta_i, h, \sigma; \omega_o) &= f^{\text{high}}(\theta_i, h, \sigma; \omega_o) \cdot N^{(n)}_{2D,h}(\lambda, \omega_o) \end{aligned} \tag{S.39}$$

We first note that for higher-order scattering, we query the noise function value only based on ω_o, ignoring the impact from θ_i. While this querying approach does not model the precise correlations in high-order scattering distributions along the incident dimension, it is the simplest method that ensures continuity in $f^{(n)}_{\text{high}}$ across incoming and outgoing directions.

For the surface reflection distribution, our ω'_o in Eq. S.39 depends on both θ_i and ω_o via a half vector calculation that preserves the memory effect. To define ω'_o, we assume that the speckle pattern in the noise tile $N^{(n)}_{2D,r}$ shows the details in a full reflection pattern from some ply, when it is illuminated from an nominal incident direction in the ply space, $\omega^n_i = \langle 0, 1, 0\rangle$—corresponding to $\theta_i = 0°$. When we query a noise function value for an incident angle θ_i, our ply-space incident direction is $\omega_i = \langle 0, \cos\theta_i, \sin\theta_i\rangle$, as given in the main paper. To preserve the shift-invariant property of the details in the reflection pattern, we look for a direction ω'_o such that

$$\frac{\omega^n_i + \omega'_o}{||\omega^n_i + \omega'_o||} = \frac{\omega_i + \omega_o}{||\omega_i + \omega_o||} \tag{S.40}$$

In other words, we pick an outgoing direction ω'_o, whose half vector formed with the nominal incident direction equals the half vector between our queried incident and outgoing direction (ω_i, ω_o), and query the reflection-related noise function at ω'_o.

5.3 Designed Noise vs. Random Noise

Finally, we examine the visual effects of augmenting our ply scattering functions with our wave optics based noise model versus simpler, pseudo-random noise functions. Renderings produced using our wave-based noise model are shown in the main paper and reproduced at the end of this document. For comparison, here we present images of the same orange polyester satin fabric generated using less carefully designed noise functions.

We first emphasize that these “less controlled” noise models cannot be simple pseudo-random sequences. Consider the following basic pipeline for generating noise values of unit mean:

- Generate two uniform random numbers $r_1, r_2 \in [0, 1)$.
- Map them to Gaussian random variables $s_1, s_2 \sim \mathcal{N}(0, 1)$.
- Compute $t = (s_1^2 + s_2^2)/2$, mimicking the quadratic form in our wave optics noise construction.

If r_1, r_2 are drawn from standard pseudo-random sequences, then t behaves as an uncontrolled positive-valued random variable with a mean value of 1. However, such randomness averages out across samples in Monte Carlo rendering. With sufficiently many samples per pixel, the resulting images converge to smooth results indistinguishable from those produced without any noise modulation.

To obtain more persistent, deterministic variations, we instead define r_1, r_2 as rapidly oscillating functions of the shaded ply index (ID), the incident angle θ_i, and the outgoing direction (θ_o, ϕ_o). As an example, one possible formulation can be:

$$\begin{aligned} r_1 &= a_1 \cdot \mathrm{ID} + b_1 \cdot \theta_i + c_1 \cdot \theta_o + d_1 \cdot \phi_o \\ r_1 &= \mathrm{fract}(e_1 \cdot \sin r_1) \\ r_2 &= a_2 \cdot \mathrm{ID} + b_2 \cdot \theta_i + c_2 \cdot \theta_o + d_2 \cdot \phi_o \\ r_2 &= \mathrm{fract}(e_2 \cdot \sin r_2) \end{aligned} \tag{S.41}$$

where a_i, b_i, c_i, d_i, e_i are reasonably large constants.

Even with this construction, we find that the rendered fabric often appears overly uniform, since pixel intensities integrate over ranges of directions and the noise values quickly average out. Reducing the constants a_i, b_i, c_i, d_i, e_i sometimes introduces visible banding artifacts, as shown in Fig. 9. The parameters used for each rendered example are listed below.

	a_1	b_1	c_1	d_1	e_1	a_2	b_2	c_2	d_2	e_2
1	505.11	495.22	503.33	497.44	501.55	496.66	504.77	498.88	502.99	499.00
2	55.11	45.22	53.33	47.44	501.55	46.66	54.77	48.88	52.99	499.00
3	5.11	5.22	3.33	7.44	501.55	6.66	4.77	8.88	2.99	499.00
4	505.11	495.22	503.33	497.44	51.55	496.66	504.77	498.88	502.99	49.00
5	505.11	495.22	503.33	497.44	6.55	496.66	504.77	498.88	502.99	4.00
6	55.11	45.22	53.33	47.44	6.55	46.66	54.77	48.88	52.99	4.00

Table 1: Parameters used in our less controlled noise model.

None of the renderings in Fig. 9 exhibit the desirable spatially varying highlights and organic irregularities observed with our wave-optics noise model. Instead, they appear either overly uniform or contain artificial banding. We therefore conclude that casually designed noise functions are inadequate for producing realistic fabric appearance. A wave optics based noise construction is necessary to generate visually pleasing imperfections that contribute to the natural appearance of woven cloths.

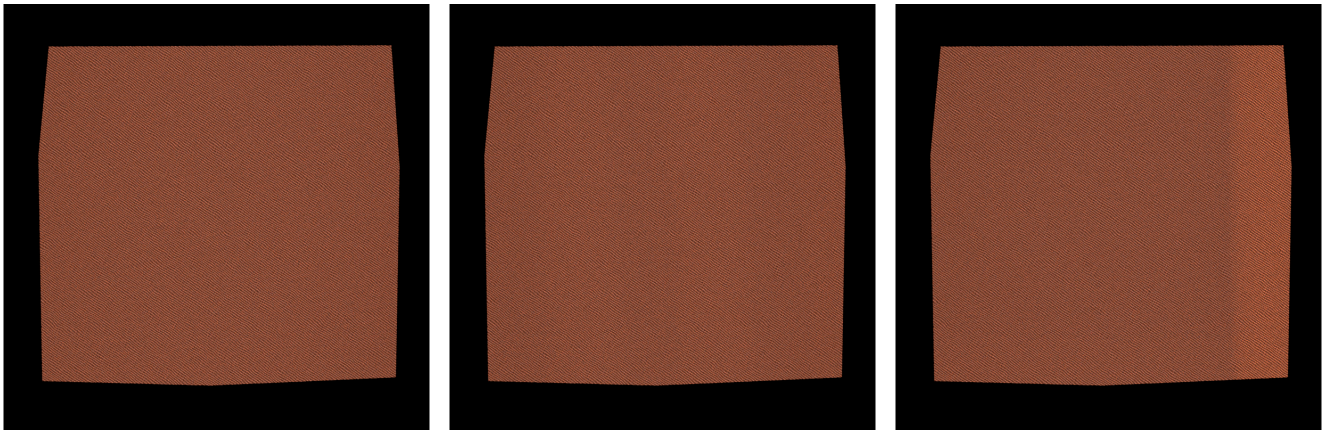

Version 1 Version 2 Version 3

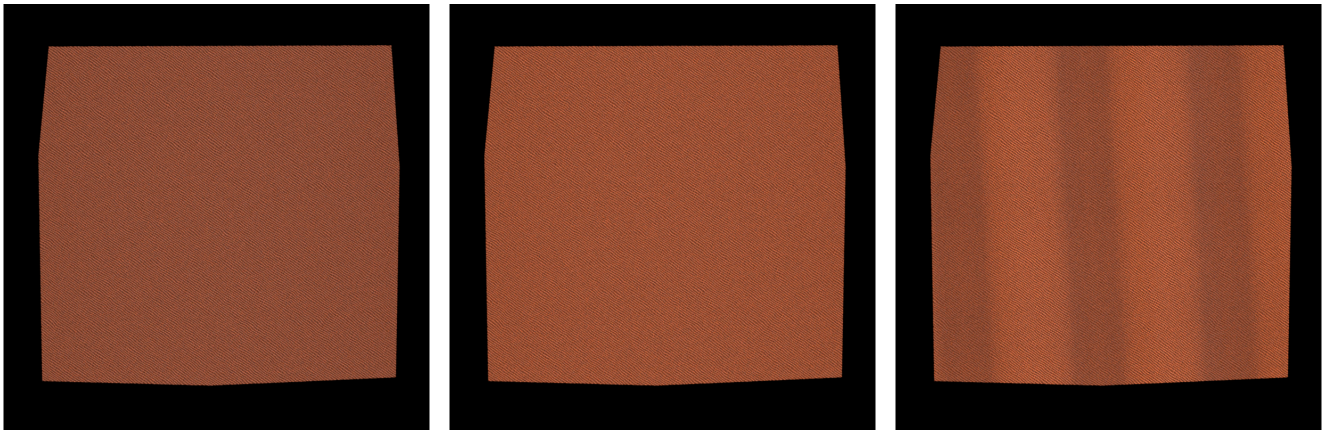

Version 4 Version 5 Version 6

Figure 9: Images of our orange polyester satin sample, rendered with less controlled, pseudo-random noise functions in a few variants.

6 Fabric Color and Absorption

As briefly mentioned in our main paper, in our model, a fabric color is not specified with an albedo, and is instead realized with an absorption spectrum, which is a function of the wavelength (recall that we use spectral rendering). Here we elaborate on the relationships between fabric colors and absorption spectra.

For each type of ply described by its scattering model, given a desired color that is initially described by an albedo in RGB, we first perform spectral upsampling following Jakob and Hanika (2019) (implemented in Mitsuba3), which gives us a per-wavelength albedo spetrum $c(\lambda)$.

We find that the underlying fibers' absorption coefficient $\sigma(\lambda)$ can be approximated using a simple analytical function in the form of

$$\sigma(\lambda) = \log \frac{[c(\lambda)]^{\frac{1}{a}}}{b} \tag{S.42}$$

More specifically, we experiment with integrating the ply's scattering distribution when it is illuminated at $\theta_i = 0$ by a wide beam (average scattering for a range of incident positions), for a dense collection of absorption coefficients. This gives us a relationship between absorption coefficient and total scattered energy. We then fit this relationship to an exponential function and invert the fitted function, which gives us Eq. S.42 parameterized by the coefficients a, b.

Due to different scattering behaviors of fibers with different sizes, shapes, and twists, the coefficients a, b are unique to each ply model. Using Eq. S.42, we can assign different absorption spectra to plies and render fabrics in arbitrary, vivid colors.

Sample 1: Orange Polyester Satin

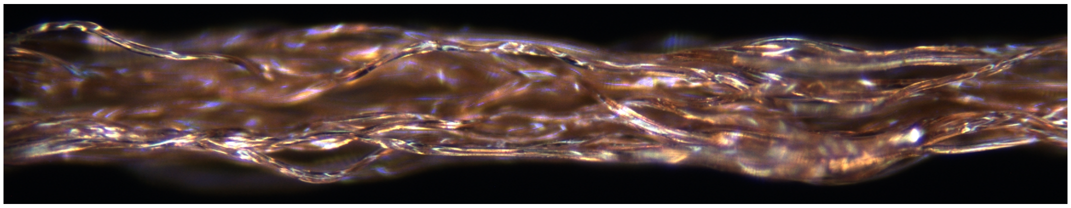

Figure 10: Closeup image of a yarn detached from an orange polyester satin fabric sample.

Our polyester satin fabric features single-ply yarns, as confirmed by manually examining detached yarns. The fiber count within each ply is estimated to be between 45 and 50. The yarn (or ply) diameter is approximately $220\mu m$, with individual fibers measuring about $18\mu m$ in diameter. While the fibers do not exhibit clear helical twisting, we still estimated a twist parameter for modeling purposes. Given the $220\mu m$ ply diameter, the fiber twist pitch was estimated to be $10000\mu m$—this twist is extremely low, as even identifying the twist direction proved difficult. The index of refraction is set to 1.57.

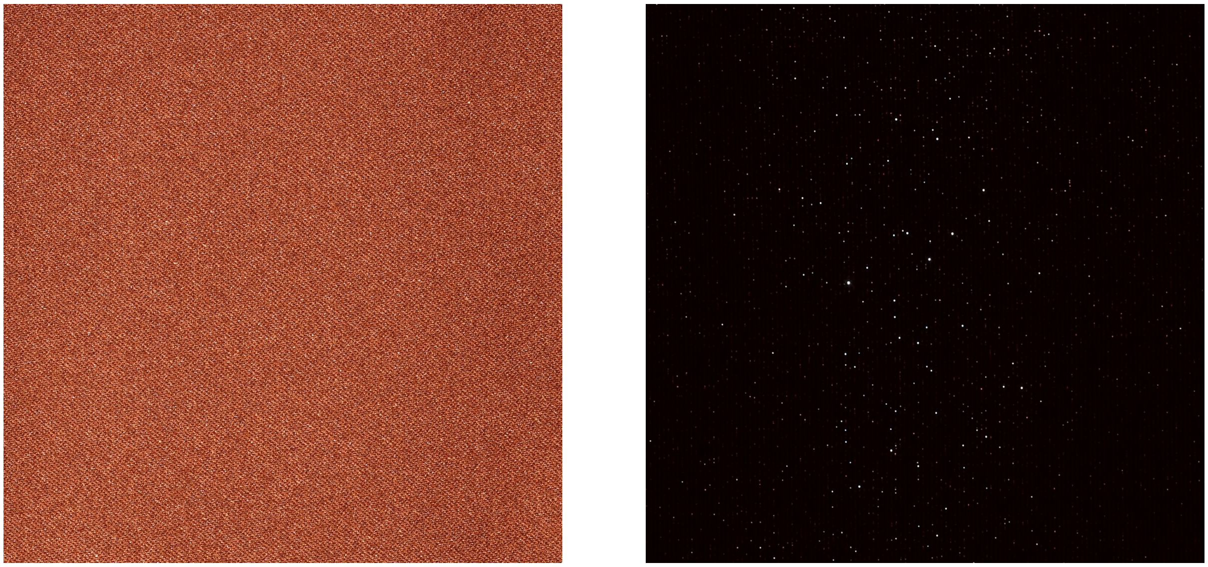

Figure 11: Captured photos of an orange polyester satin sample, front-lit (left) and back-lit (right). As the fabric is opaque, the back-lit photo is dark.

Sample 1: Orange Polyester Satin

We present two back-lit and four front-lit images of the sample, rendered with our shading model. For our transmission images, the important wave optics feature is diffraction, and we compare images rendered with and without modeling diffraction, replacing diffraction with delta transmission in the non-wave variant (matching previous works).

Moreover, textures underlying our surface-based cloth rendering, which describe the weave patterns as well as ply-level geometry, often include some gentle Perlin noise perturbation to avoid overly periodic looking patterns. This Perlin noise, as well as our wave optics noise for infusing irregularities into appearance, can be freely turned on and off when rendering. To demonstrate different combinations of perturbation and noise usage, we rendered four different variants of front-lit images. As we did not make our Perlin noise perturbation aggressive (otherwise it may introduce unnatural looking weave patterns), our wave optics noise makes a much more significant difference to the resulting images.

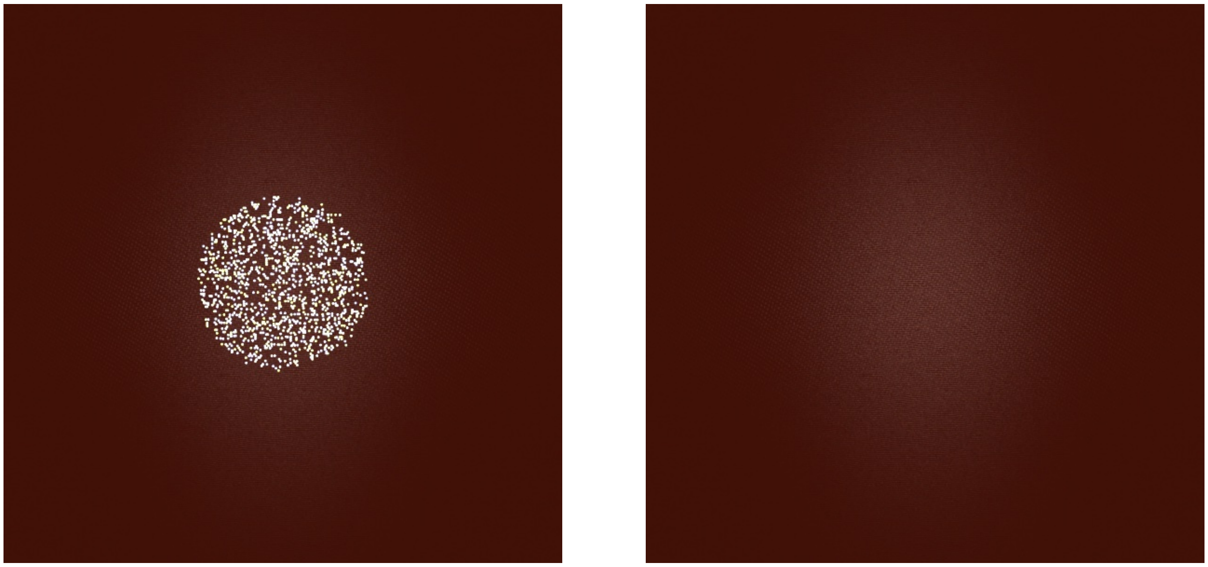

Delta Transmission

Diffraction

Figure 12: Two back-lit images of an orange polyester satin fabric sample rendered with different variants of our method. We can easily replace our diffraction modeling with simple delta transmission, thereby turning off this wave optics feature.

Sample 1: Orange Polyester Satin

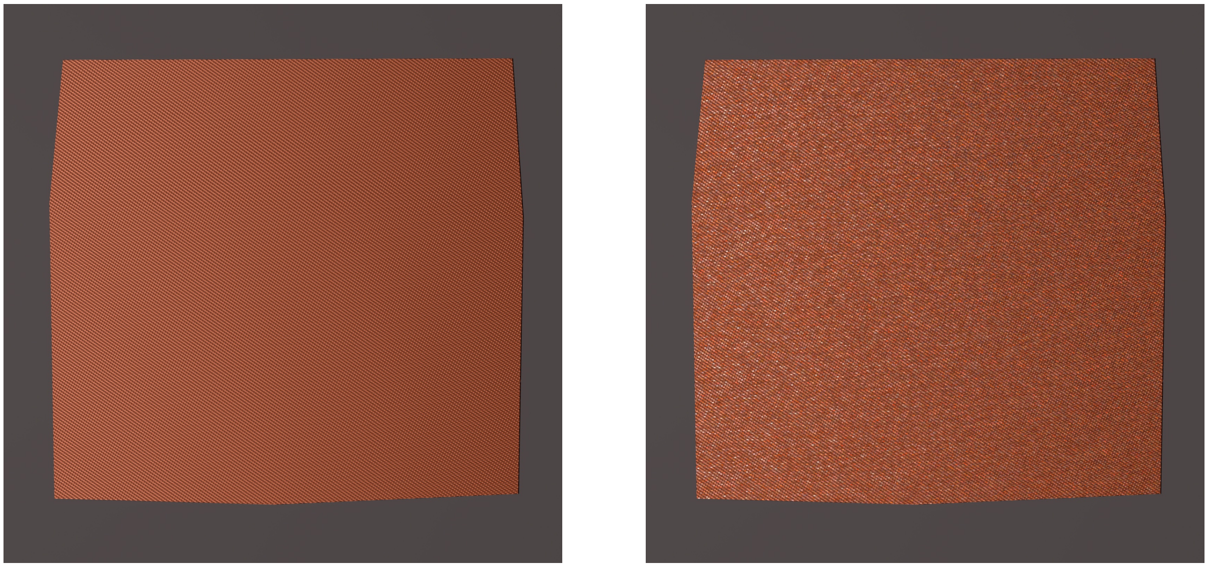

Perlin Off Wave Off

Perlin Off Wave On

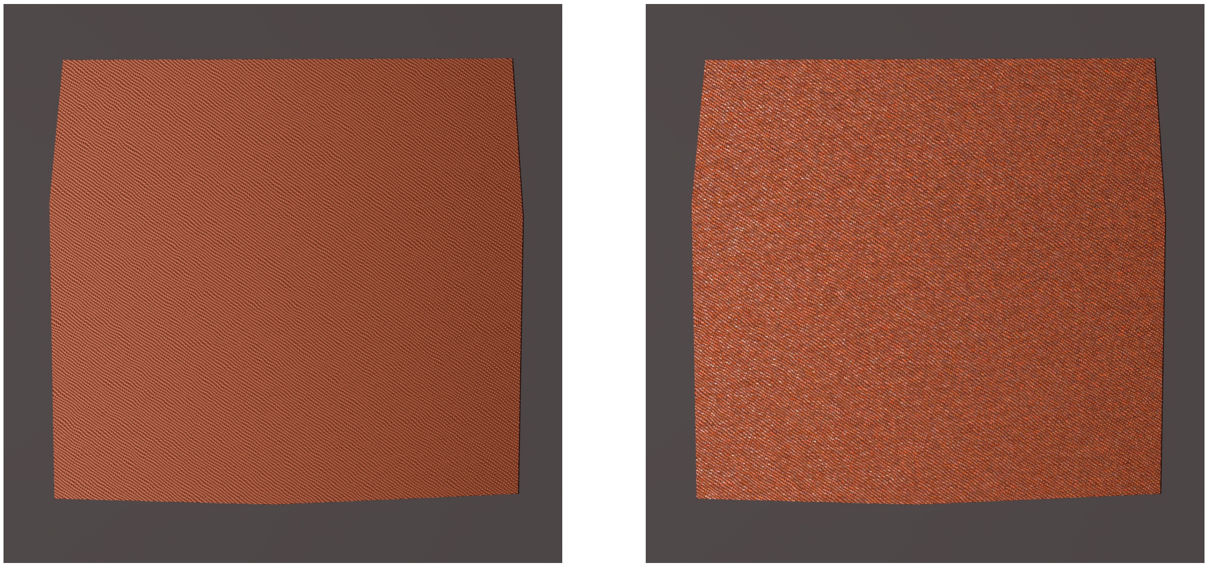

Perlin On Wave Off

Perlin On Wave On

Figure 13: Four front-lit images of an orange polyester satin fabric sample rendered with different variants of our method. The differences between the Perlin on/off images can be observed after zooming in.

Sample 2: White Silk Organza

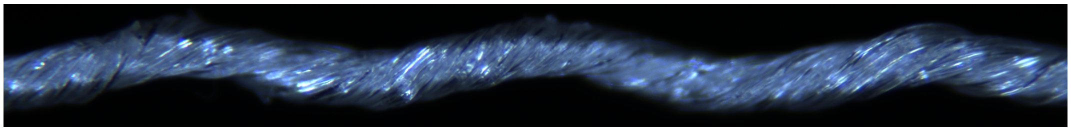

Figure 14: Closeup image of a yarn detached from a white silk organza fabric sample. Fiber color appears to be blue-shifted due to the lighting.

We believe the yarns from the silk organza sample consists of 3 plies, and each ply contains around 20–25 fibers, giving rise to a total of 60–75 fibers in the yarn. The yarn diameter is around $100\mu m$ and ply diameter was set to $45\mu m$. The pitch for fiber twist was estimated to be $1000\mu m$. The fiber diameter was set to $10\mu m$, and fiber index of refraction is 1.54.

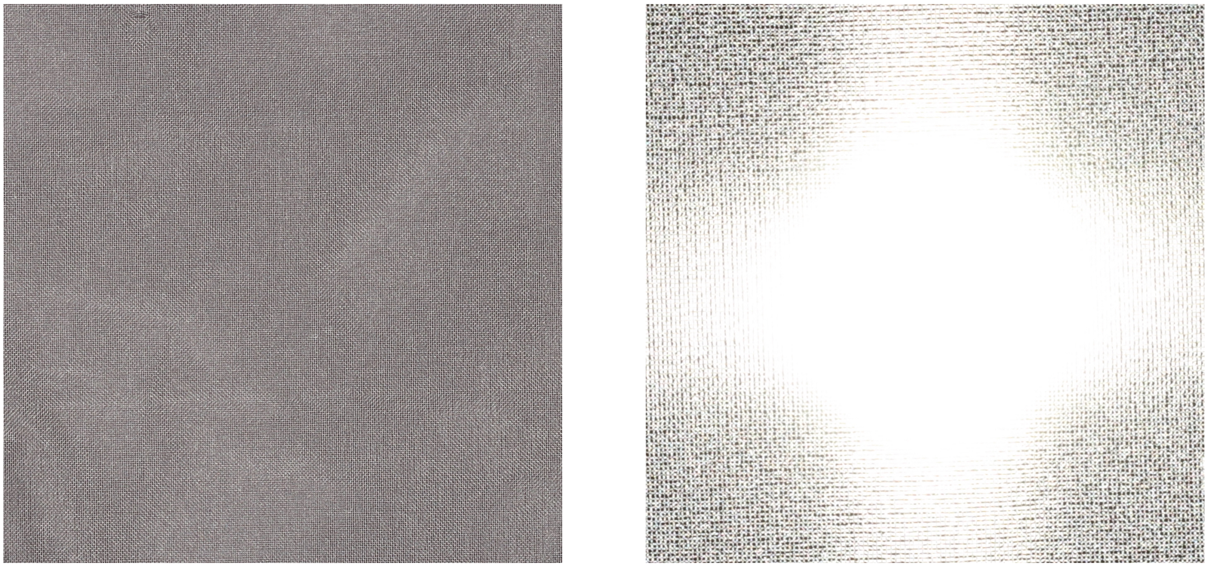

Figure 15: Captured photos of a white silk organza sample, front-lit (left) and back-lit (right). The back-lit photo demonstrates a signature cross shape.

Sample 2: White Silk Organza

We present two back-lit and four front-lit images of the sample, rendered with our shading model. For our transmission images, the important wave optics feature is diffraction, and we compare images rendered with and without modeling diffraction, replacing diffraction with delta transmission in the non-wave variant (matching previous works).

Moreover, textures underlying our surface-based cloth rendering, which describe the weave patterns as well as ply-level geometry, often include some gentle Perlin noise perturbation to avoid overly periodic looking patterns. This Perlin noise, as well as our wave optics noise for infusing irregularities into appearance, can be freely turned on and off when rendering. To demonstrate different combinations of perturbation and noise usage, we rendered four different variants of front-lit images. As we did not make our Perlin noise perturbation aggressive (otherwise it may introduce unnatural looking weave patterns), our wave optics noise makes a much more significant difference to the resulting images.

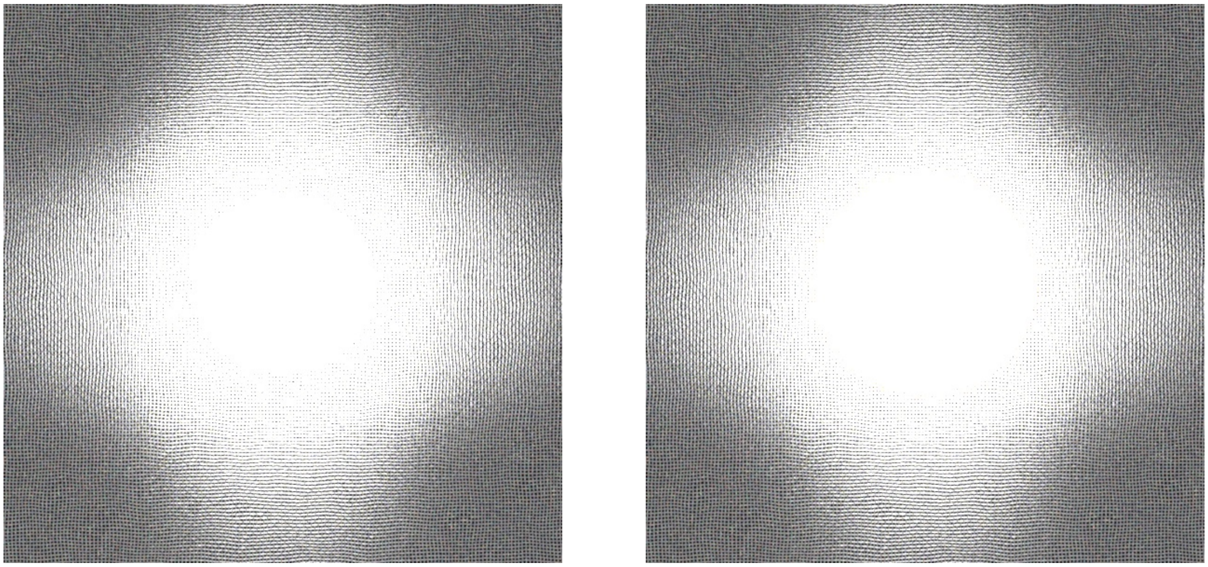

Delta Transmission

Diffraction

Figure 16: Two back-lit images of a white silk organza sample rendered with different variants of our method. We can easily replace our diffraction modeling with simple delta transmission, thereby turning off this wave optics feature.

Sample 2: White Silk Organza

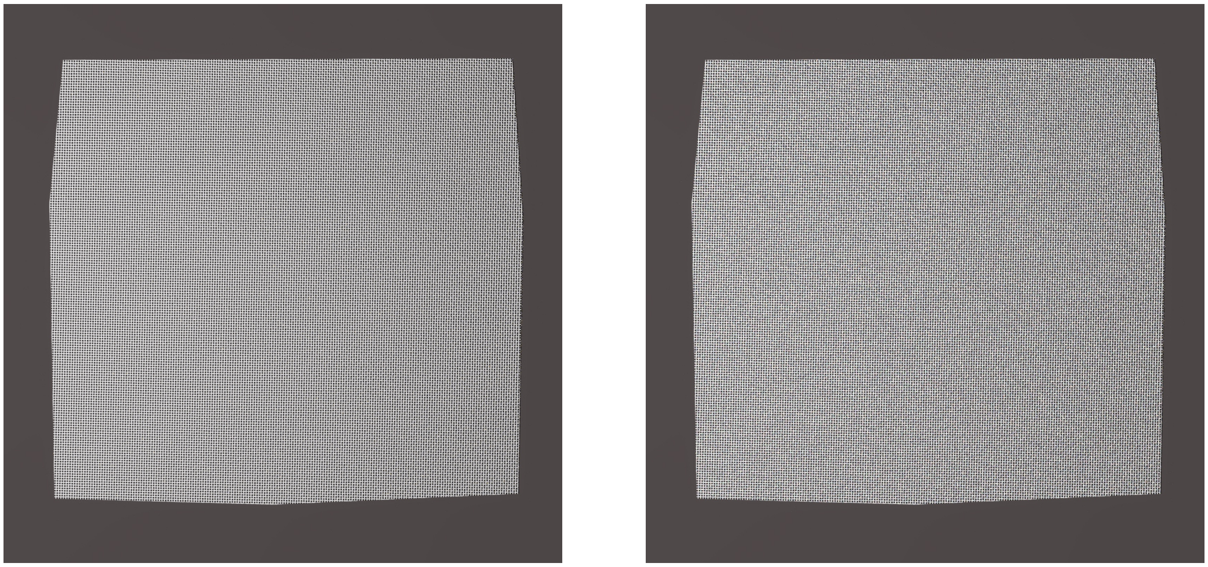

Perlin Off Wave Off

Perlin Off Wave On

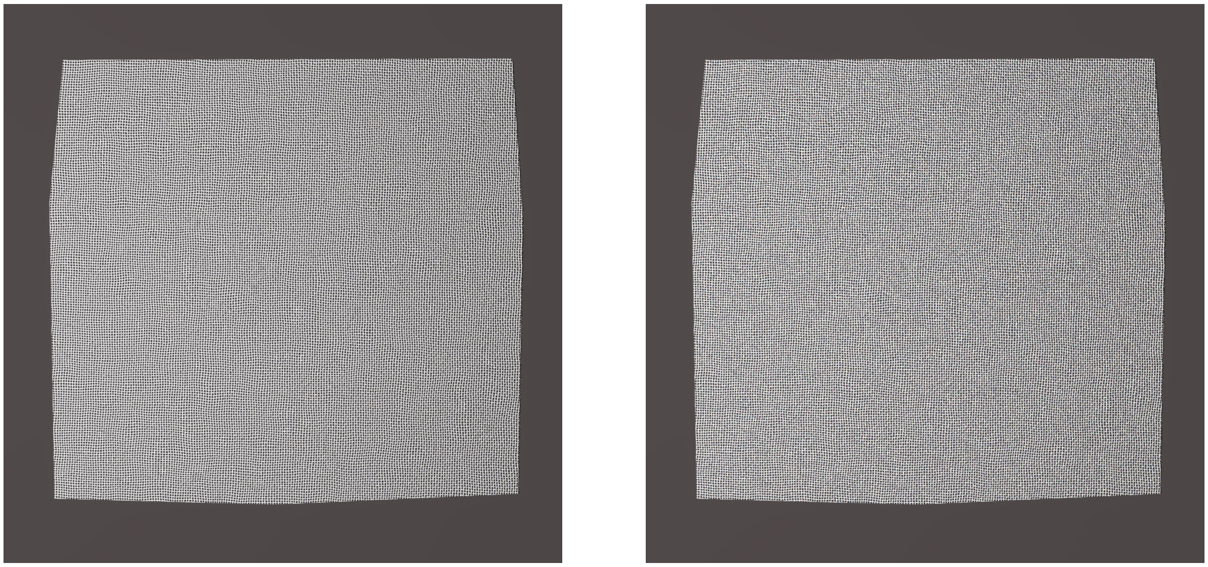

Perlin On Wave Off

Perlin On Wave On

Figure 17: Four front-lit images of a white silk organza fabric sample rendered with different variants of our method. The differences between the Perlin on/off images can be observed after zooming in.

Sample 3: Cream Cotton Twill

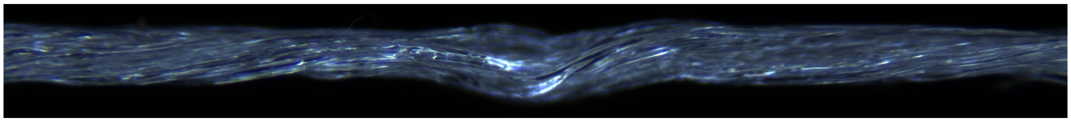

Figure 18: Closeup image of a yarn detached from a cream cotton twill fabric sample. Fiber color appears to be blue-shifted due to the lighting.

The cotton yarns in this sample were also considered to be triple-ply yarn, with each ply containing around 20–25 fibers. As can be seen from the photos, the scales and sizes of the cotton yarn and the silk yarn from the previous sample are rather similar. The yarn diameter is slightly larger, estimated to be $120\mu m$, and ply diameter was set to around $55\mu m$. Interestingly, fiber twist seems to be lower for the cotton yarn than for the silk yarn, with pitch estimated to be around $1800\mu m$. The fiber diameter was also set to $10\mu m$, and fiber index of refraction is 1.53.

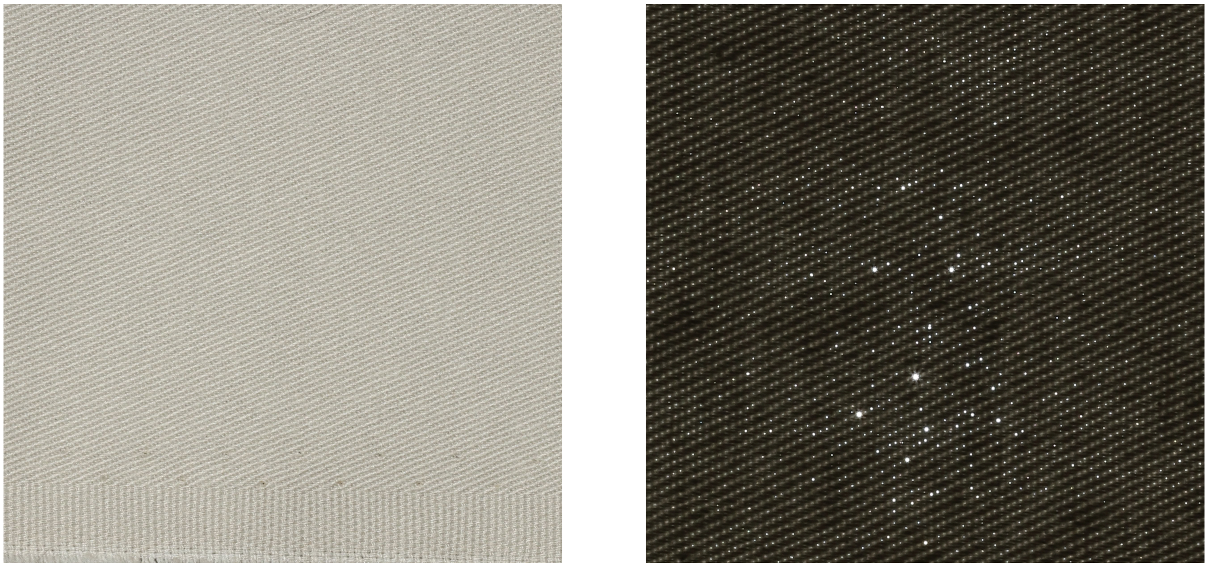

Figure 19: Captured photos of a cream cotton twill sample, front-lit (left) and back-lit (right). As the fabric is opaque, the back-lit photo is rather dark.

Sample 3: Cream Cotton Twill

We present two back-lit and four front-lit images of the sample, rendered with our shading model. For our transmission images, the important wave optics feature is diffraction, and we compare images rendered with and without modeling diffraction, replacing diffraction with delta transmission in the non-wave variant (matching previous works).

Moreover, textures underlying our surface-based cloth rendering, which describe the weave patterns as well as ply-level geometry, often include some gentle Perlin noise perturbation to avoid overly periodic looking patterns. This Perlin noise, as well as our wave optics noise for infusing irregularities into appearance, can be freely turned on and off when rendering. To demonstrate different combinations of perturbation and noise usage, we rendered four different variants of front-lit images. As we did not make our Perlin noise perturbation aggressive (otherwise it may introduce unnatural looking weave patterns), our wave optics noise makes a much more significant difference to the resulting images.

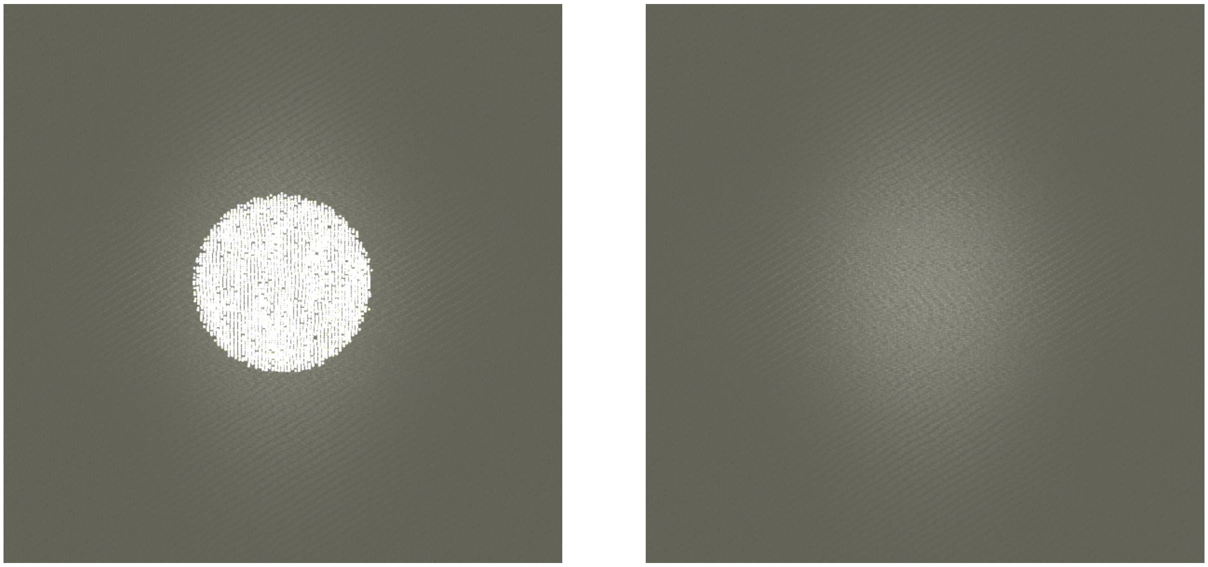

Delta Transmission

Diffraction

Figure 20: Two back-lit images of a cream cotton twill fabric sample rendered with different variants of our method. We can easily replace our diffraction modeling with simple delta transmission, thereby turning off this wave optics feature.

Sample 3: Cream Cotton Twill

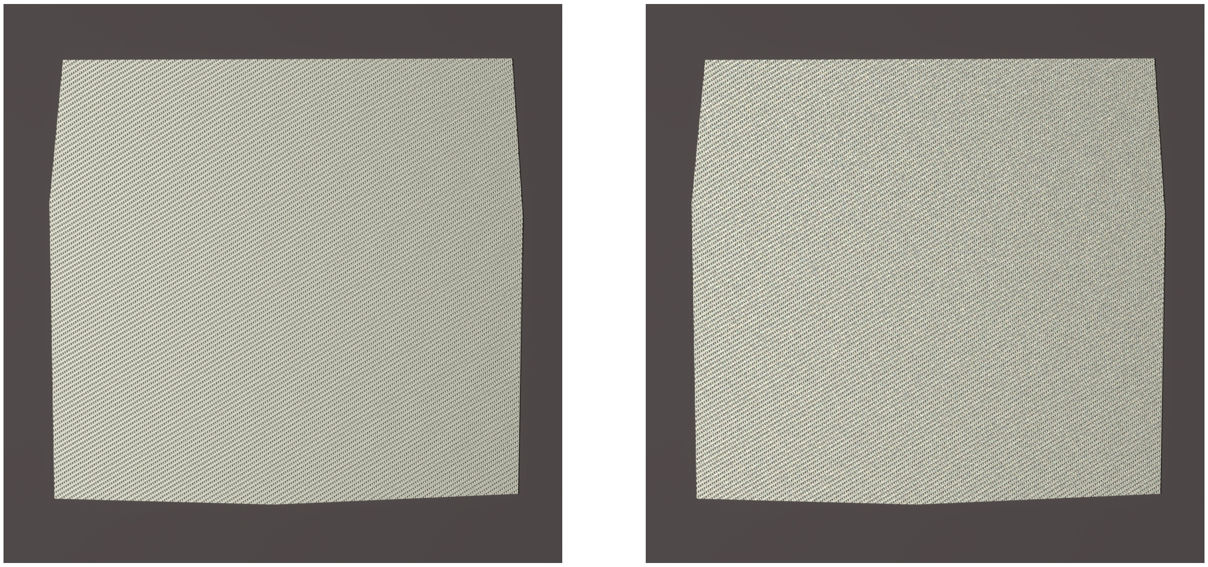

Perlin Off Wave Off

Perlin Off Wave On

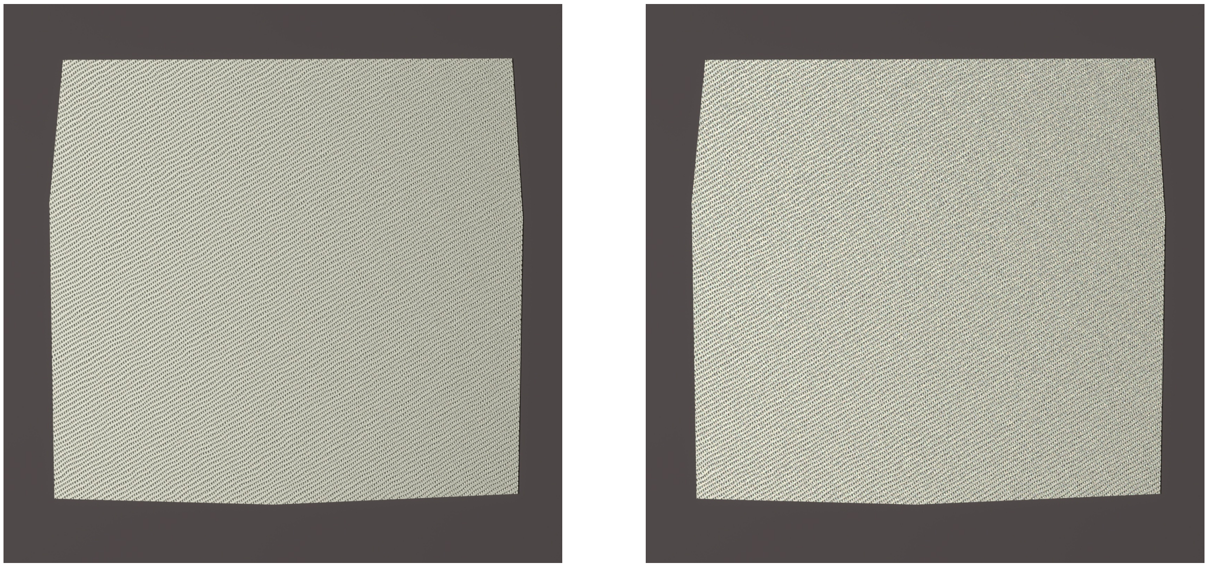

Perlin On Wave Off

Perlin On Wave On

Figure 21: Four front-lit images of a cream cotton twill fabric sample rendered with different variants of our method. The differences between the Perlin on/off images can be observed after zooming in.

Sample 4: Gray Wool Twill

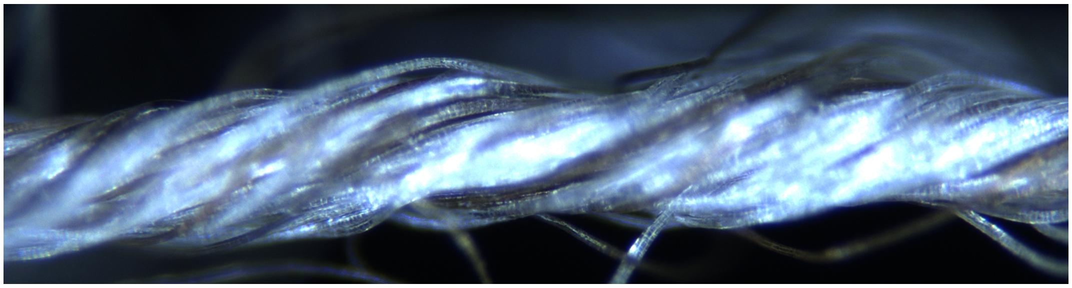

Figure 22: Closeup image of a yarn detached from a gray wool twill fabric sample. Fiber color appears to be blue-shifted due to the lighting.

The wool yarn appears to be messy, and we also consider it a 3-ply yarn, with each ply containing around 38–42 fibers. The yarn diameter is around $300\mu m$ and the ply diameter around $140\mu m$. Fiber diameter is set to $15\mu m$, and the pitch for fiber twist is set to $3000\mu m$. The fiber index of refraction is 1.55.

Figure 23: Captured photos of a gray wool twill sample, front-lit (left) and back-lit (right). As the fabric is very thick and opaque, the back-lit photo is dark.

Sample 4: Gray Wool Twill

We present two back-lit and four front-lit images of the sample, rendered with our shading model. For our transmission images, the important wave optics feature is diffraction, and we compare images rendered with and without modeling diffraction, replacing diffraction with delta transmission in the non-wave variant (matching previous works).

Moreover, textures underlying our surface-based cloth rendering, which describe the weave patterns as well as ply-level geometry, often include some gentle Perlin noise perturbation to avoid overly periodic looking patterns. This Perlin noise, as well as our wave optics noise for infusing irregularities into appearance, can be freely turned on and off when rendering. To demonstrate different combinations of perturbation and noise usage, we rendered four different variants of front-lit images. As we did not make our Perlin noise perturbation aggressive (otherwise it may introduce unnatural looking weave patterns), our wave optics noise makes a much more significant difference to the resulting images.

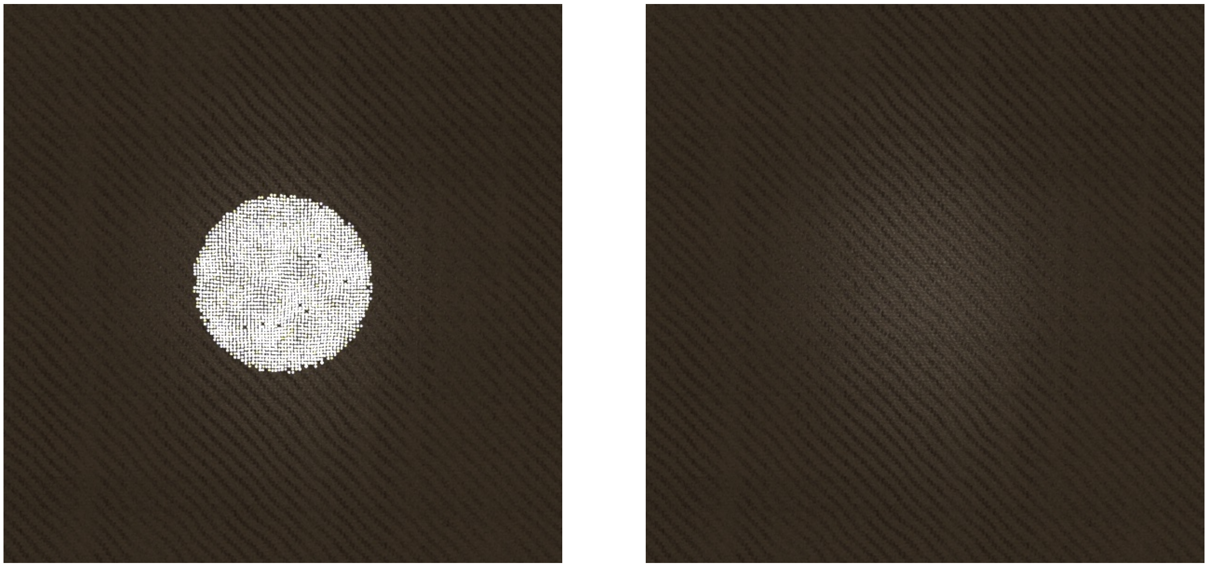

Delta Transmission Diffraction

Figure 24: Two back-lit images of a gray wool twill fabric sample rendered with different variants of our method. We can easily replace our diffraction modeling with simple delta transmission, thereby turning off this wave optics feature.

Sample 4: Gray Wool Twill

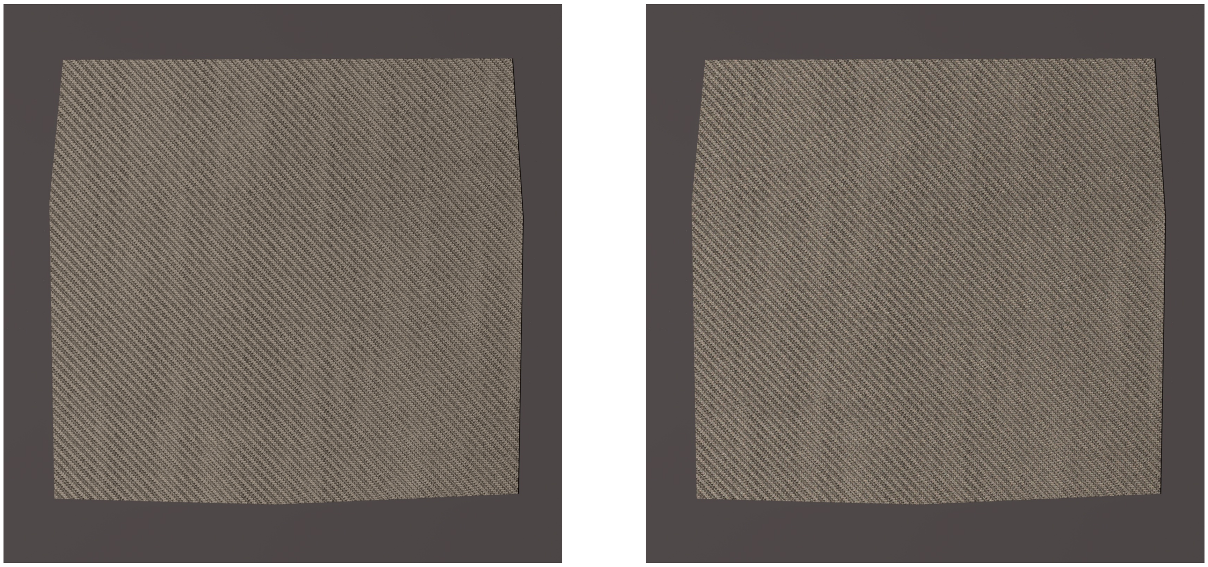

Perlin Off Wave Off

Perlin Off Wave On

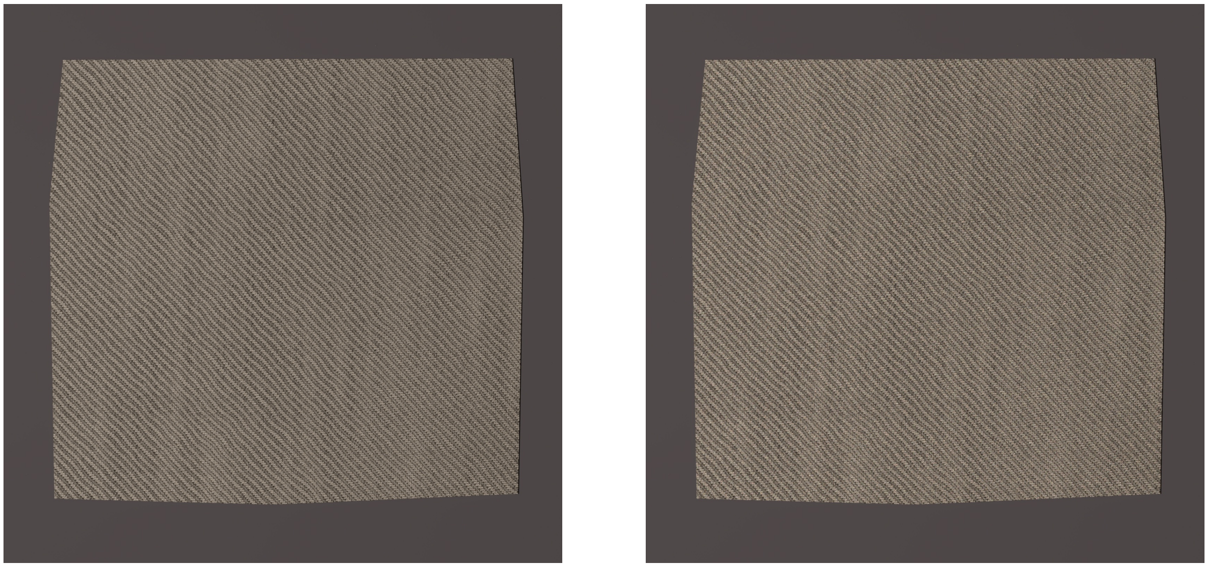

Perlin On Wave Off

Perlin On Wave On

Figure 25: Four front-lit images of a gray wool twill fabric sample rendered with different variants of our method. The differences between the Perlin on/off images can be observed after zooming in.

Sample 5: Black Polyester Plain Weave

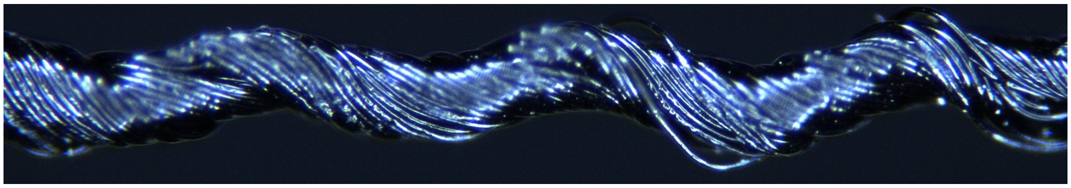

Figure 26: Closeup image of a yarn detached from a black polyester plain weave fabric sample.

We believe the yarns from this polyester plain weave sample consists of 3 plies, and each ply contains around 25–30 fibers, giving rise to a total of 75–90 fibers in the yarn. The yarn diameter is around $160\mu m$ and ply diameter was set to $75\mu m$. The pitch for fiber twist was estimated to be $1200\mu m$. The fiber diameter was set to $12\mu m$, and fiber index of refraction is 1.57.

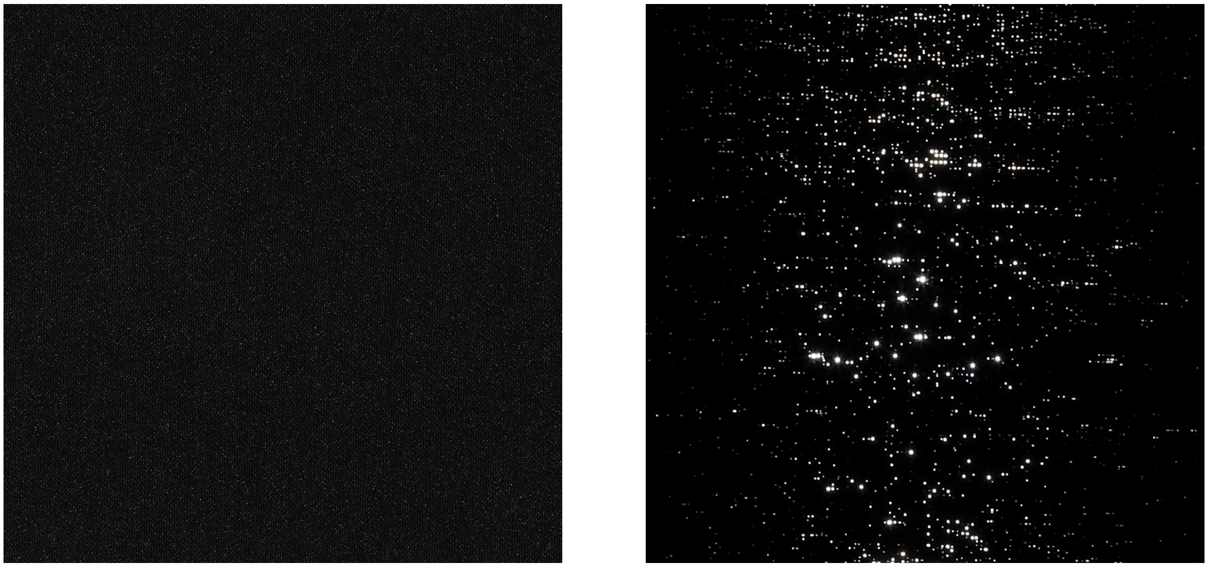

Figure 27: Captured photos of a black polyester plain weave sample, front-lit (left) and back-lit (right). While both photos appear very dark, the back-lit photo reveals lots of pin holes on the fabric.

Sample 5: Black Polyester Plain Weave

We present two back-lit and four front-lit images of the sample, rendered with our shading model. For our transmission images, the important wave optics feature is diffraction, and we compare images rendered with and without modeling diffraction, replacing diffraction with delta transmission in the non-wave variant (matching previous works).

Moreover, textures underlying our surface-based cloth rendering, which describe the weave patterns as well as ply-level geometry, often include some gentle Perlin noise perturbation to avoid overly periodic looking patterns. This Perlin noise, as well as our wave optics noise for infusing irregularities into appearance, can be freely turned on and off when rendering. To demonstrate different combinations of perturbation and noise usage, we rendered four different variants of front-lit images. As we did not make our Perlin noise perturbation aggressive (otherwise it may introduce unnatural looking weave patterns), our wave optics noise makes a much more significant difference to the resulting images.

Delta Transmission

Diffraction

Figure 28: Two back-lit images of a black polyester plain weave fabric sample rendered with different variants of our method. We can easily replace our diffraction modeling with simple delta transmission, thereby turning off this wave optics feature.

Sample 5: Black Polyester Plain Weave

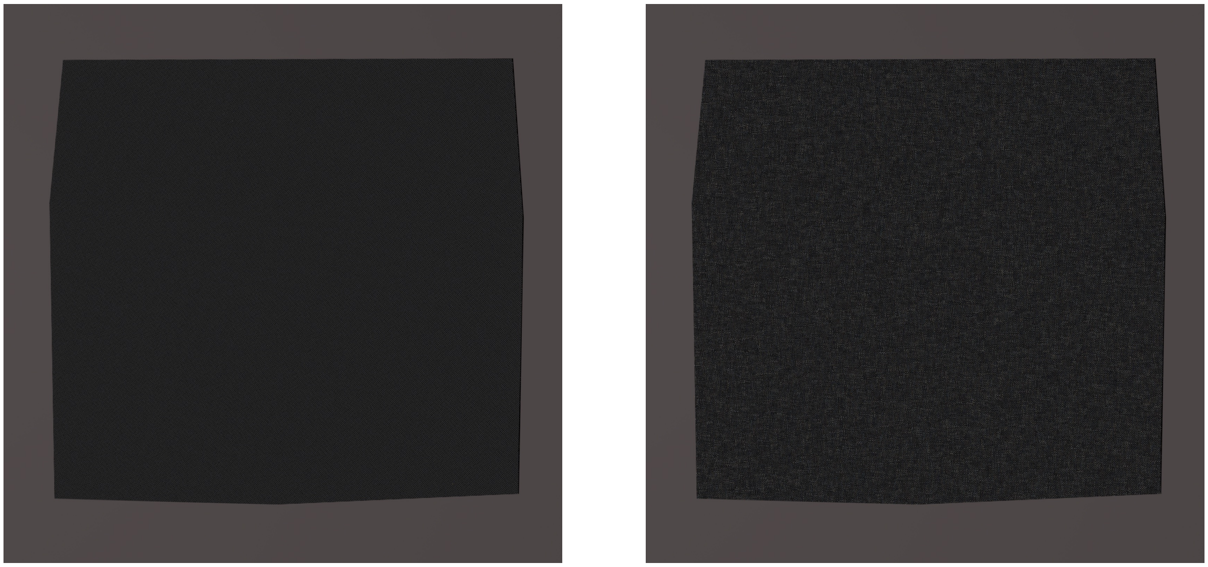

Perlin Off Wave Off

Perlin Off Wave On

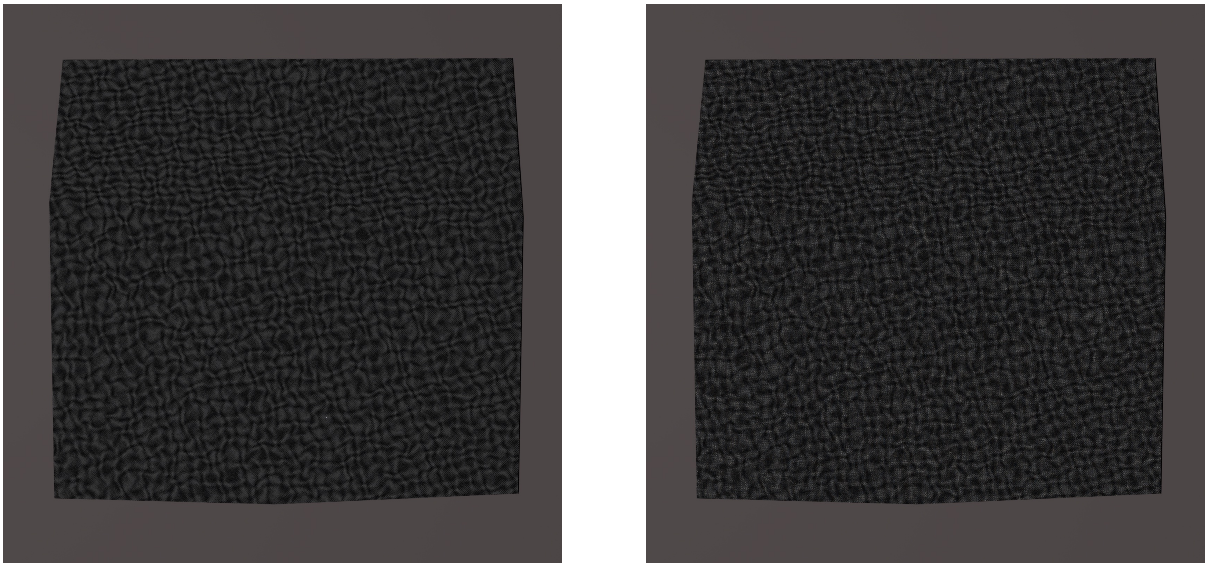

Perlin On Wave Off

Perlin On Wave On

Figure 29: Four front-lit images of a black polyester plain weave sample rendered with different variants of our method. The differences between the Perlin on/off images can be observed after zooming in.

Sample 6: Beige Silk Satin

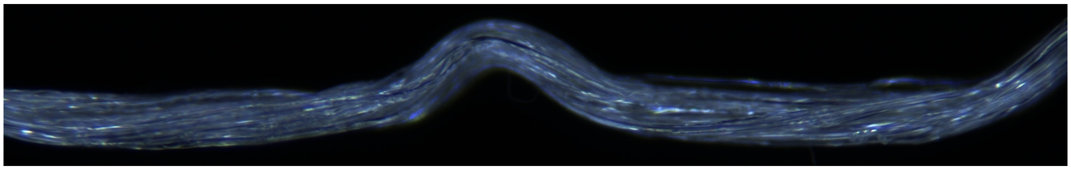

Figure 30: Closeup image of a yarn detached from a beige silk satin fabric sample. Fiber color appears to be blue-shifted due to the lighting.

We think that the yarns from this silk satin sample consists of a single ply, which contains 35–40 fibers. The yarn diameter, or ply diameter is around $110\mu m$. The pitch for fiber twist was estimated to be $2500\mu m$. The fiber diameter was set to $12\mu m$, and fiber index of refraction is 1.54.

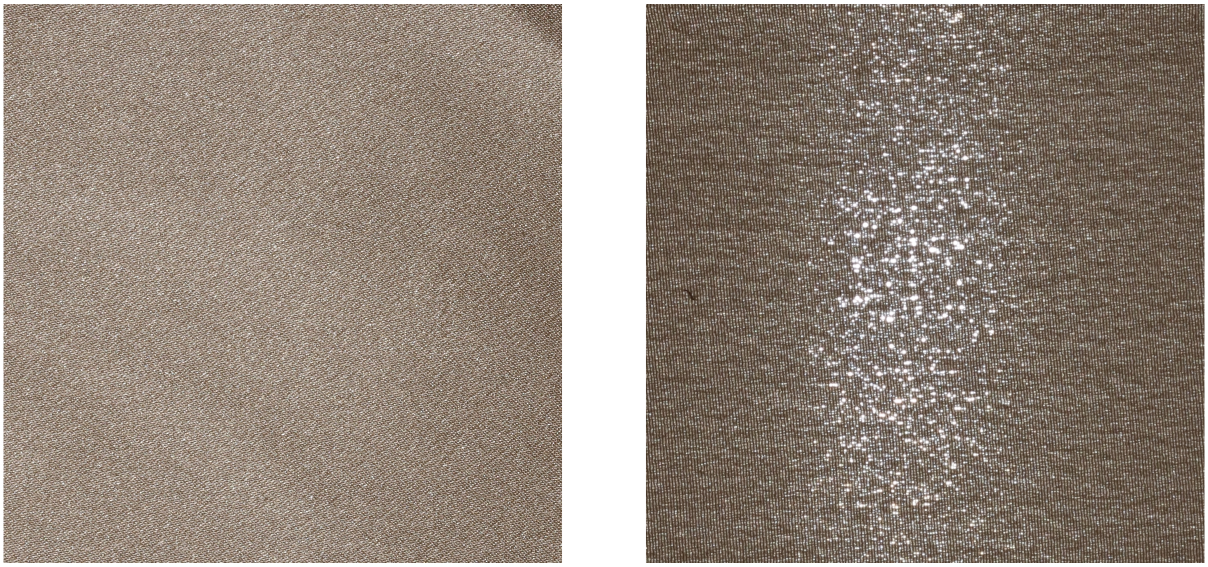

Figure 31: Captured photos of a beige silk satin sample, front-lit (left) and back-lit (right). Note that the back-lit photo features a highly anisotropic transmission pattern.

Sample 6: Beige Silk Satin

We present two back-lit and four front-lit images of the sample, rendered with our shading model. For our transmission images, the important wave optics feature is diffraction, and we compare images rendered with and without modeling diffraction, replacing diffraction with delta transmission in the non-wave variant (matching previous works).

Moreover, textures underlying our surface-based cloth rendering, which describe the weave patterns as well as ply-level geometry, often include some gentle Perlin noise perturbation to avoid overly periodic looking patterns. This Perlin noise, as well as our wave optics noise for infusing irregularities into appearance, can be freely turned on and off when rendering. To demonstrate different combinations of perturbation and noise usage, we rendered four different variants of front-lit images. As we did not make our Perlin noise perturbation aggressive (otherwise it may introduce unnatural looking weave patterns), our wave optics noise makes a much more significant difference to the resulting images.

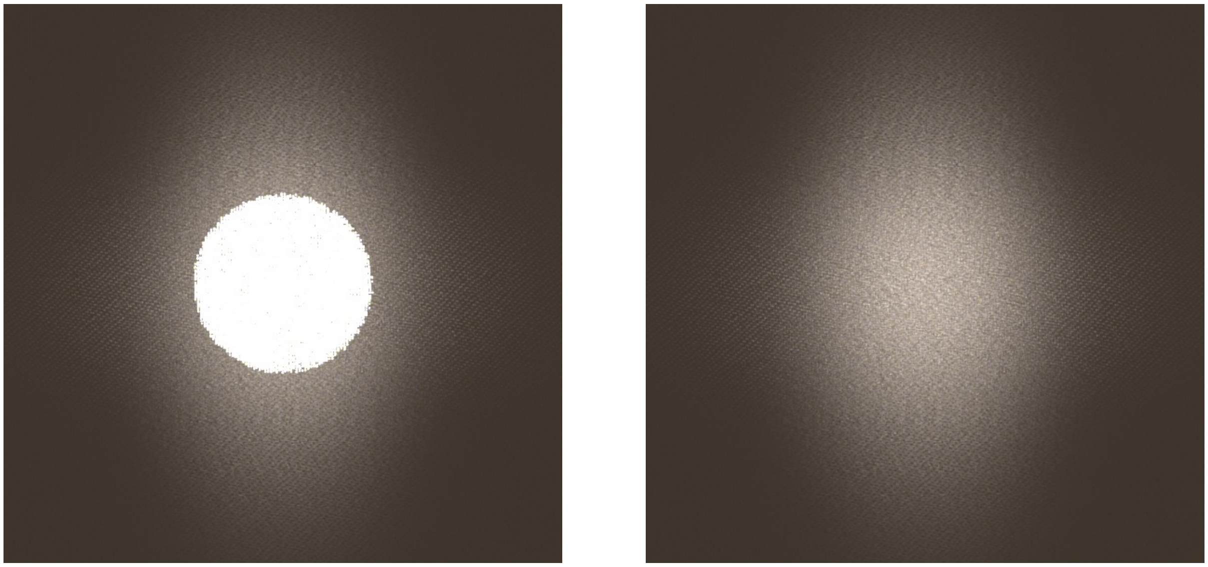

Delta Transmission

Diffraction

Figure 32: Two back-lit images of a beige silk satin fabric sample rendered with different variants of our method. We can easily replace our diffraction modeling with simple delta transmission, thereby turning off this wave optics feature.

Sample 6: Beige Silk Satin

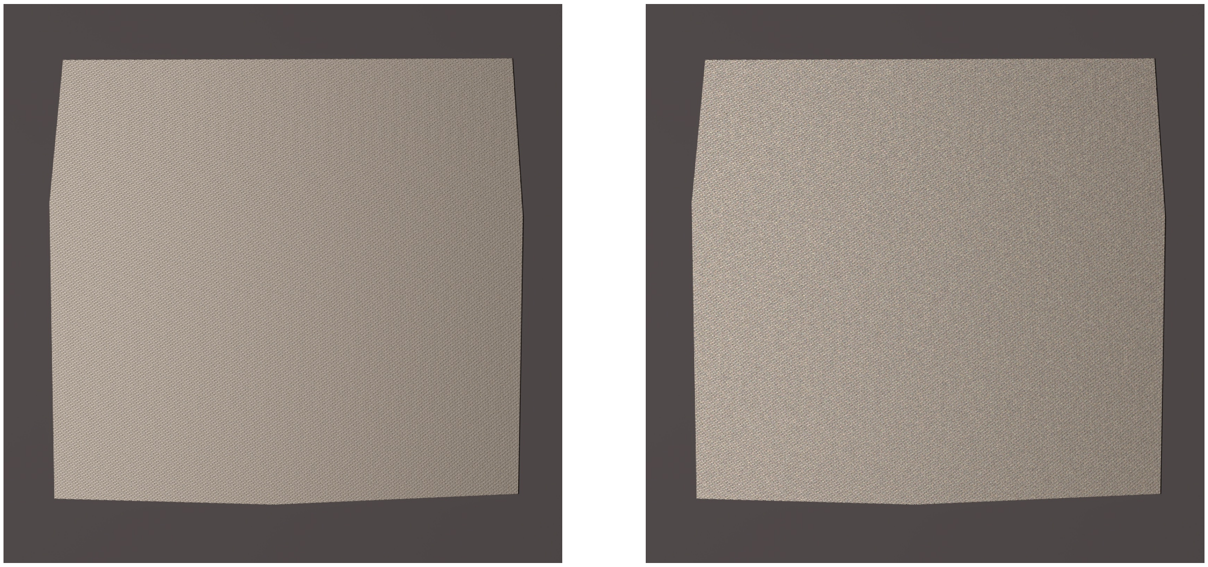

Perlin Off Wave Off

Perlin Off Wave On

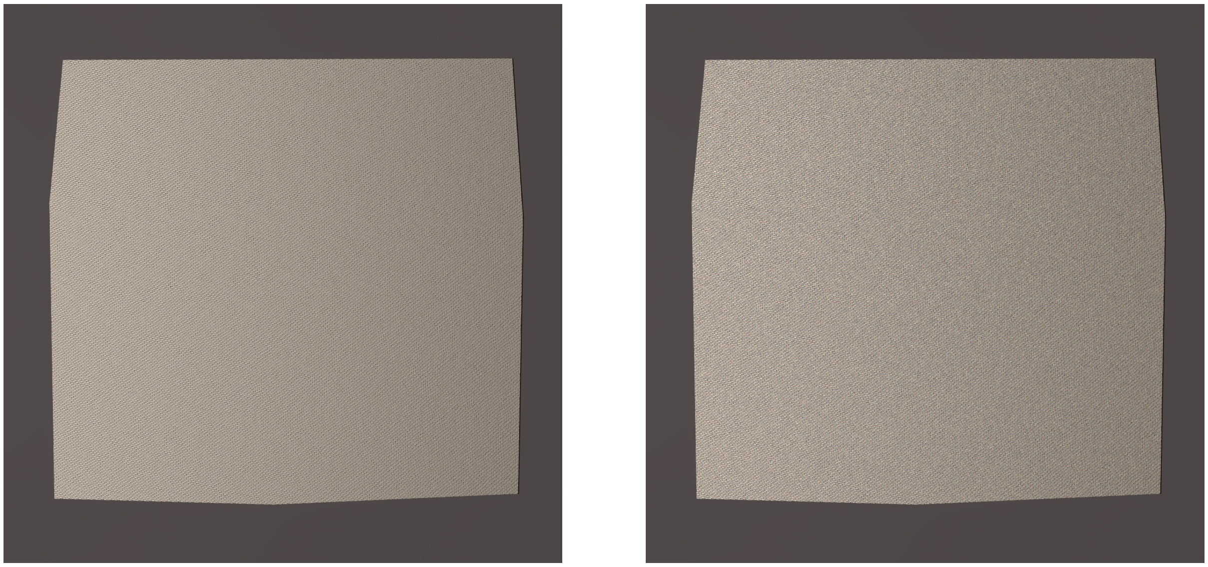

Perlin On Wave Off

Perlin On Wave On

Figure 33: Four front-lit images of a beige silk satin fabric sample rendered with different variants of our method. The differences between the Perlin on/off images can be observed after zooming in.

Sample 7: Blue Cotton Satin

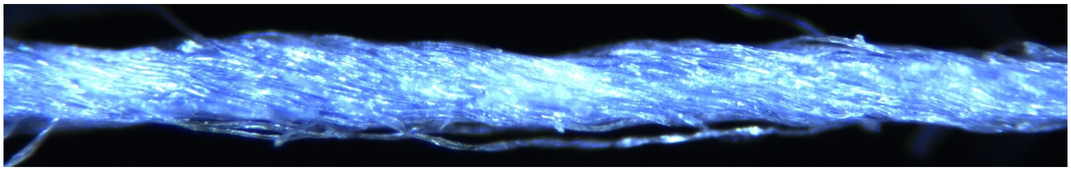

Figure 34: Closeup image of a yarn detached from a blue cotton satin fabric sample.

We believe the yarns from this cotton sample consists of 3 plies, and each ply contains around 30–35 fibers, giving rise to a total of 90–105 fibers in the yarn. The yarn diameter is around $170\mu m$ and ply diameter was set to $80\mu m$. The pitch for fiber twist was estimated to be $2000\mu m$. The fiber diameter was set to $12\mu m$, and fiber index of refraction is 1.53.

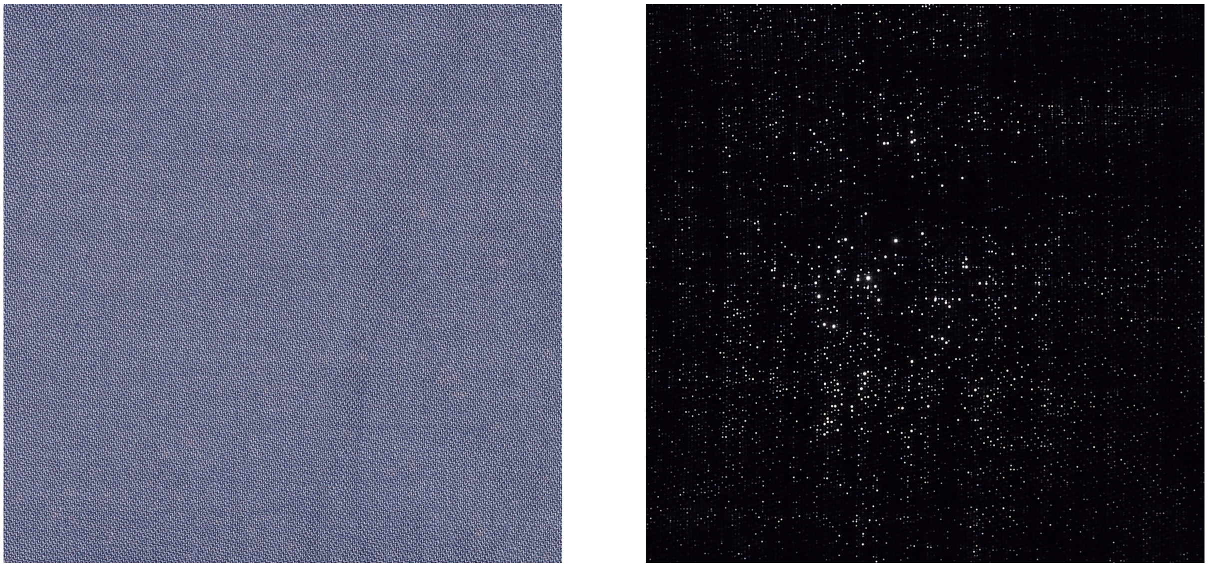

Figure 35: Captured photos of a blue cotton satin sample, front-lit (left) and back-lit (right). As the fabric is opaque, the back-lit photo is dark.

Sample 7: Blue Cotton Satin

We present two back-lit and four front-lit images of the sample, rendered with our shading model. For our transmission images, the important wave optics feature is diffraction, and we compare images rendered with and without modeling diffraction, replacing diffraction with delta transmission in the non-wave variant (matching previous works).

Moreover, textures underlying our surface-based cloth rendering, which describe the weave patterns as well as ply-level geometry, often include some gentle Perlin noise perturbation to avoid overly periodic looking patterns. This Perlin noise, as well as our wave optics noise for infusing irregularities into appearance, can be freely turned on and off when rendering. To demonstrate different combinations of perturbation and noise usage, we rendered four different variants of front-lit images. As we did not make our Perlin noise perturbation aggressive (otherwise it may introduce unnatural looking weave patterns), our wave optics noise makes a much more significant difference to the resulting images.

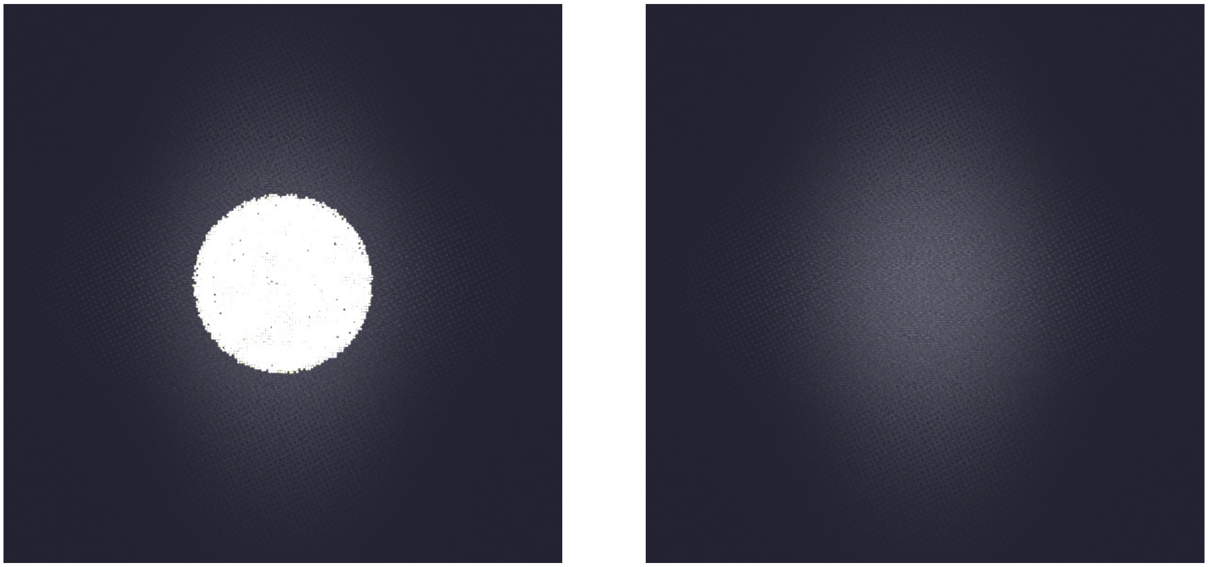

Delta Transmission

Diffraction

Figure 36: Two back-lit images of a blue cotton satin fabric sample rendered with different variants of our method. We can easily replace our diffraction modeling with simple delta transmission, thereby turning off this wave optics feature.

Sample 7: Blue Cotton Satin

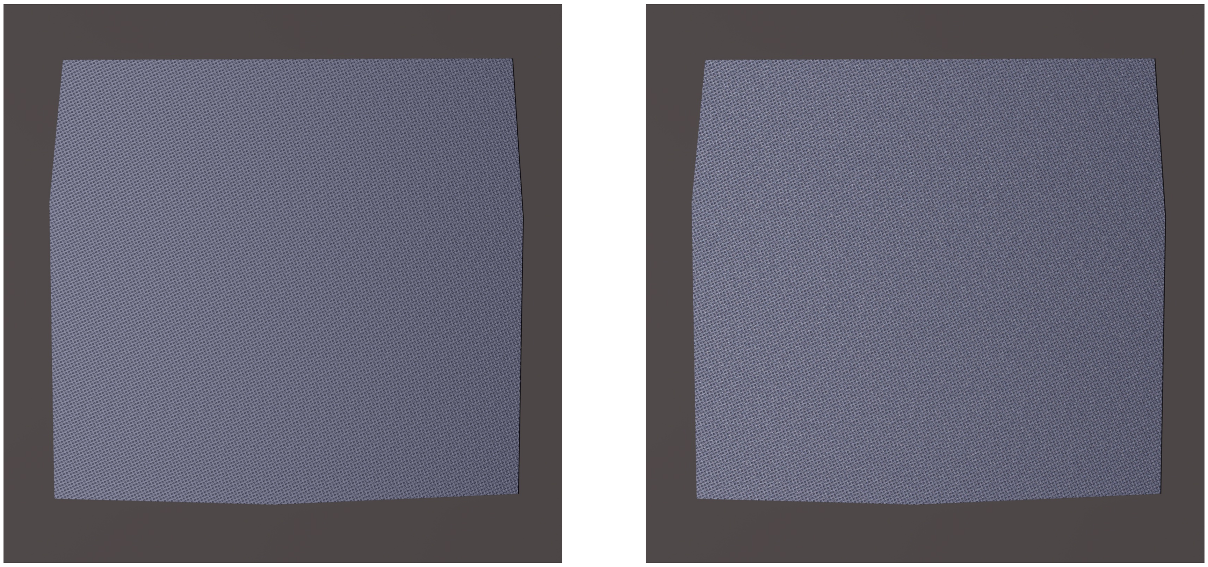

Perlin Off Wave Off

Perlin Off Wave On

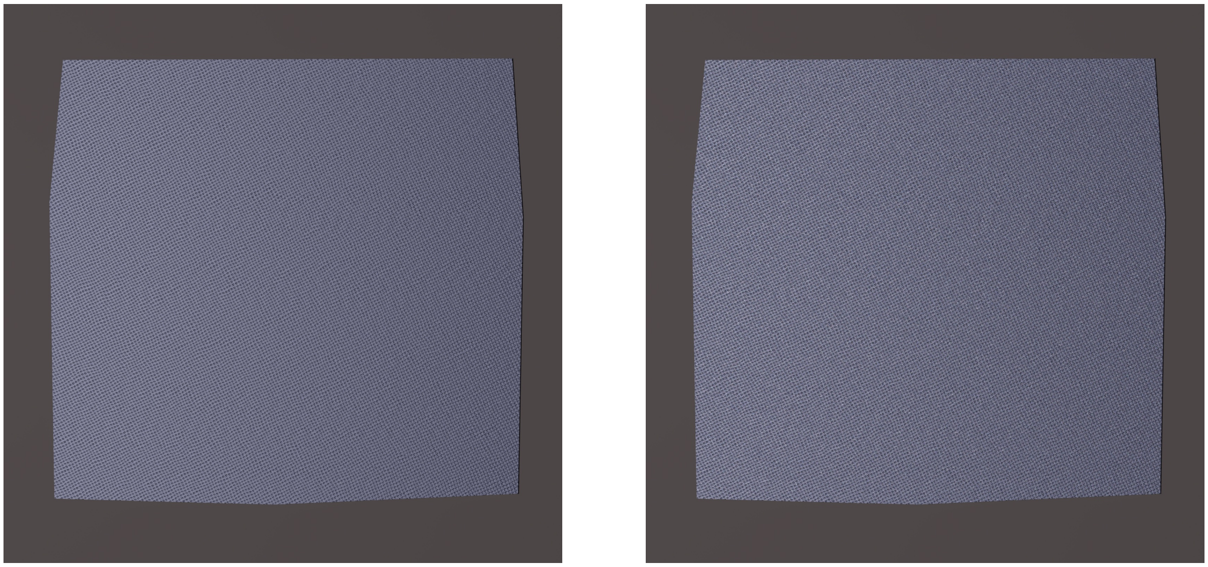

Perlin On Wave Off

Perlin On Wave On

Figure 37: Four front-lit images of a blue cotton satin fabric sample rendered with different variants of our method. The differences between the Perlin on/off images can be observed after zooming in.

References

Walton C Gibson. 2021. *The method of moments in electromagnetics.* Chapman and Hall/CRC.

John C Heurtley. 1973. Scalar Rayleigh–Sommerfeld and Kirchhoff diffraction integrals: a comparison of exact evaluations for axial points. *Journal of the Optical Society of America* 63, 8 (1973), 1003–1008.

Wenzel Jakob and Johannes Hanika. 2019. A low-dimensional function space for efficient spectral upsampling. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 147–155.

Mengqi Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. 2023. A practical wave optics reflection model for hair and fur. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–15.

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 (TOG)* 32, 6 (2013), 1–11.