Predicting Fabric Appearance Through Thread Scattering and Inversion

MENGQI (MANDY) XIA, Yale University, USA
ZHAOYANG ZHANG, Yale University, USA
SUMIT CHATURVEDI, Yale University, USA
YUTONG YI, Yale University, USA
RUNDONG WU, Bytedance, USA
HOLLY RUSHMEIER, Yale University, USA
JULIE DORSEY, Yale University, USA

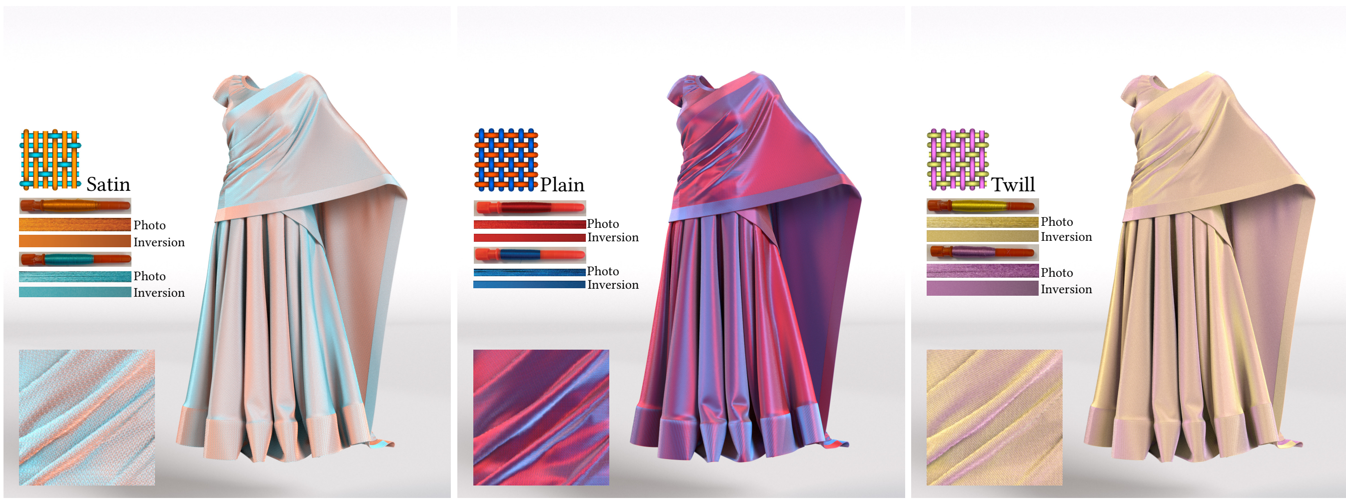

Fig. 1. We predict the appearance of a sari in three different configurations, each with satin, plain, and twill weave patterns, respectively. In each configuration, the warp and weft yarn scattering models are inverted from physical threads. We show the photos of the thread spools, the aligned threads, and their corresponding inversion results. We simulate woven yarn geometry and illustrate the initial configuration of the tiling unit in the figure. We implement a two-scale rendering method to achieve both an accurate and rapid preview of the fabric appearance.

The fashion industry has a real need to preview fabric designs using the actual threads they intend to use, ensuring that the designs they envisage can be physically realized. Unfortunately, today's fabric rendering relies on either hand-tuned parameters or parameters acquired from already fabricated cloth. Furthermore, existing curve-based scattering models are not suitable for this problem: they are either not naturally differentiable due to discrete fiber count parameters, or require a more detailed geometry representation, introducing extra complexity.

In this work, we bridge this gap by presenting a novel pipeline that captures and digitizes physical threads and predicts the appearance of the fabric based on the weaving pattern. We develop a practical thread scattering model based on simulations of multiple fiber scattering within a thread. Using a cost-efficient multi-view setup, we capture threads of diverse colors and materials. We apply differentiable rendering to digitize threads, demonstrating that our model significantly improves the reconstruction accuracy compared to existing models, matching both reflection and transmission. We leverage a two-scale rendering technique to efficiently render woven cloth. We validate that our digital threads, combined with simulated woven yarn geometry, can accurately predict the fabric appearance by comparing to real samples. We show how our work can aid designs using diverse thread profiles, woven patterns, and textured design patterns.

Authors' addresses: Mengqi (Mandy) Xia, Yale University, USA, mengqi.xia@yale.edu; Zhaoyang Zhang, Yale University, New Haven, USA, zhaoyang.zhang@yale.edu; Sumit Chaturvedi, Yale University, New Haven, USA, sumit.chaturvedi@yale.edu; Yutong Yi, Yale University, New Haven, USA, andrew.yi@yale.edu; Rundong Wu, Bytedance, San Jose, USA, wrd920@gmail.com; Holly Rushmeier, Yale University, New Haven, USA, holly.rushmeier@yale.edu; Julie Dorsey, Yale University, New Haven, USA, julie.dorsey@yale.edu.

ACM 0730-0301/2025/8-ART
https://doi.org/10.1145/3731178

CCS Concepts: • **Computing methodologies → Reflectance modeling**.

ACM Reference Format:
Mengqi (Mandy) Xia, Zhaoyang Zhang, Sumit Chaturvedi, Yutong Yi, Rundong Wu, Holly Rushmeier, and Julie Dorsey. 2025. Predicting Fabric Appearance Through Thread Scattering and Inversion. *ACM Trans. Graph.* 44, 4 (August 2025), 12 pages. https://doi.org/10.1145/3731178

1 INTRODUCTION

Fabric rendering plays a crucial role in a wide range of applications, such as design, visual effects, augmented and virtual reality (AR/VR), and online retail. However, one issue that has not been adequately addressed is the prediction of fabric appearance from actual physical threads. This capability is needed for fashion designers to explore the appearance of fabrics that could be physically produced. By accurately modeling and simulating real threads, designers can minimize the reliance on physical prototypes, thus reducing production costs and shortening the time to market.

Modeling fabric appearance is an active research topic in graphics. Based on the underlying geometry representation, fabric rendering works can be roughly categorized into three classes: 1) Curve-based models [Chiang et al. 2015; d'Eon et al. 2011; Marschner et al. 2003; Montazeri et al. 2020; Zhu et al. 2023b], 2) Surface-based models [Irawan and Marschner 2012; Sadeghi et al. 2013; Wang et al. 2022; Zhu et al. 2024, 2023a], and 3) Volumetric models [Jakob et al. 2010; Zhao et al. 2011, 2012]. For the purpose of digitizing physical threads, a curve-based model is desirable. However, existing curve-based scattering models are not suitable for this task. Fiber-based scattering models require detailed geometry representation within a single thread, which is difficult to acquire ground truth data for and adds complexity to the fitting process. The other curve-based scattering models are either not naturally differentiable due to discrete fiber count parameters [Zhu et al. 2023b] or contain sophisticated modes with various parameters and are less desirable for inversion [Montazeri et al. 2020]. Moreover, ply and yarn-based models have mostly been used to perform forward rendering instead of inversion tasks. The closest prior work for achieving the goal is [Sadeghi et al. 2013]. However, their method requires an elaborate capture setup and illumination of a single thread, which is extremely challenging for thin fibers like silk. Additionally, manual fitting is needed to match thread scattering to the capture, further complicating the process. The final rendering is surface-based, losing yarn-level details.

To address the challenge of predicting the appearance of cloth from physical threads, or yarns, we have developed a novel end-to-end pipeline, illustrated in Figure 2. First, we utilize a cost-efficient multiview setup to capture both the reflection and transmission properties of the threads. We collect a thread dataset with diverse colors captured from multiple views. We introduce a novel and compact thread scattering model based on Monte Carlo simulation, which considers light interaction with multiple fibers within a thread. This model focuses on the statistical average of multiple scattering events, rather than specific multi-fiber configurations. We then apply differentiable rendering to digitize these threads. Next, we conduct physical simulations to generate a tileable patch with yarn curves. Finally, we implement a two-scale rendering method that first traces the macro surface and then simulates light interaction within the patch. Our pipeline is easily reproducible and enables efficient and realistic previews of fabric appearance.

We demonstrate that our model significantly improves thread reconstruction accuracy compared to existing models, effectively matching both reflection and transmission views. We acquire a diverse collection of fabric samples featuring different weave patterns and colors, along with the corresponding threads woven into those fabrics. Validation confirms that our approach accurately predicts fabric appearance, whereas existing methods produce inaccurate colors and intensities.

We show how our work aids designs using diverse thread profiles and weaving patterns. We also demonstrate potential usage in embroidery design and cloth design with texture patterns. This work closely connects rendering research and practical demands of fashion design, and opens new possibilities for innovation in the design and fabrication.

Concretely, our main contributions include:

- A novel end-to-end pipeline that is easily reproducible, enabling efficient and accurate predictions of fabric appearance from physical threads.
- A diverse thread dataset captured from multiple reflection and transmission views.
- A compact and differentiable thread scattering model based on Monte Carlo simulation of multiple fiber scattering.

2 RELATED WORK

Surface-based cloth models. Surface-based cloth models treat cloth as 2D thin sheets and model the appearance using Bidirectional Reflectance Distribution Functions (BRDFs) [Adabala et al. 2003; Irawan and Marschner 2012; Sadeghi et al. 2013; Wu et al. 2011]. Wang et al. [2022] approximates layers of an SGGX microflake medium [Heitz et al. 2015] and applies it to render cloth. Subsequent work [Zhu et al. 2023a] extends this approach to add ply-level details, accounting for shadowing and masking effects, while [Zhu et al. 2024] further handles parallax effects and improves efficiency. Although surface-based cloth models are fast to render, they often lack detailed appearance features. They are not suitable to achieve our goal as they do not offer yarn-level control.

Curve-based scattering models. Curve-based scattering models represent fibers, plies, or yarns in fabric as curves and more accurately reproduce detailed appearance effects [Chiang et al. 2015; d'Eon et al. 2011; Luan et al. 2017; Marschner et al. 2003; Xia et al. 2020; Zhao et al. 2016]. Curve-based scattering models have also been used to render hair and fur [Xia et al. 2023; Yan et al. 2017, 2015; Zhu et al. 2022]. Representing threads as curves best suits our goal; however, existing works are not suitable. We demonstrate that the single fiber scattering model cannot adequately represent threads due to the lack of multiple scattering. Zhu et al. [2023b] introduces an aggregated yarn-based scattering model, which is not naturally differentiable due to the discrete fiber count parameter. We show that, given a number of fibers, it produces less accurate thread inversion results and fabric appearance predictions.

Micro-appearance cloth models. Besides curve-based scatteirng models, there are other micro-appearance cloth models such as volumetric models [Jakob et al. 2010; Zhao et al. 2011, 2012] and fiber mesh models [Khungurn et al. 2015; Schroder et al. 2011]. These models are highly versatile and can model detailed appearance effects. However, they are not suitable for predicting fabric appearance from physical threads.

Inverse problems for fabric. Various methods have been proposed for fabric geometry capture, such as those by [Guarnera et al. 2017;

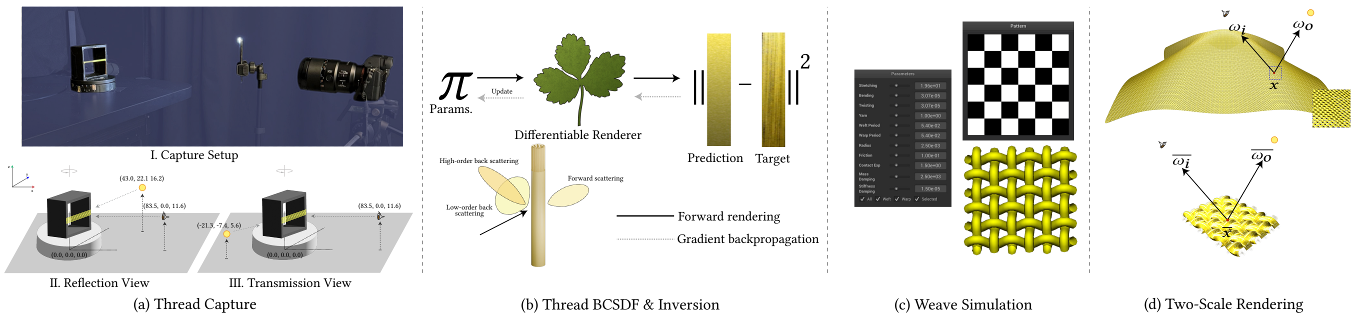

Fig. 2. Overview of our pipeline. (a) **Thread Capture:** We place threads on a hollow box in a dark room and use a commodity camera to capture their reflectance and transmission, illuminated by a cell phone flash light from either the front or back. We acquire 12 reflection and 5 transmission views. A schematic (in cm) illustrates the two capture settings. (b) **Thread BCSDF & Inversion:** We optimize the thread BCSDF parameters using a differentiable renderer (Mitsuba 3). (c) **Weave Simulation:** We perform a physics-based simulation to obtain yarn curves in a woven patch. (d) **Two-Scale Rendering:** To efficiently render woven fabric, we apply a two-scale approach. Only a small, tileable patch needs to be stored in memory. Rays are first traced in the macro-geometry of the surface, followed by a secondary ray-tracing step in the woven micro-geometry.

Schröder et al. 2014], which recover weave patterns and thread parameters from a single photo to instantiate a 3D fabric model. Zhao et al. [2016] fit CT data to procedural models to reconstruct highly detailed yarns. However, CT scans are very expensive to acquire. Recently, Jin et al. [2022] used a simple setup to capture and invert a surface-based cloth model, with [Tang et al. 2024] further extending this work by adding a back-view photo to recover transmission properties. Deng et al. [2024] use PBR and parametric modeling to construct woven textured materials with centimeter and millimeter level 3D structures. Sadeghi et al. [2013] measures the 4D scattering distribution function of a single thread using a *gonioreflectometer*, and renders cloth based on the thread model acquired from the capture. Although this is close to our goal, this work is not easy to apply in practice. It is extremely challenging to focus on thin threads. Additionally, their method relies on manual fitting, which complicates the process. This work renders cloth using a surface based model, losing yarn-level appearance details.

3 THREAD SCATTERING MODEL

3.1 Fiber scattering model

Fiber scattering models represent individual fibers as cylinders and use the Bidirectional Curve Scattering Distribution Function (BCSDF) to characterize the scattering properties of a fiber. Similar to the BSDF, it describes outgoing radiance L_r as an integration of incident radiance L_i multiplied by the BCSDF S:

$$L_r\left(\omega_r\right) = \int L_i\left(\omega_i\right) S\left(\omega_i, \omega_r\right) \cos\theta_i \mathrm{d}\omega_i. \quad (1)$$

We will express the BCSDF in spherical coordinates as $S(\theta_i, \theta_r, \phi_i, \phi_r)$, with θ being longitudinal angles and ϕ being azimuthal angles. Most existing models represent the BCSDF as a sum of reflective and transmissive modes S_p, with each mode S_p factored into a longitudinal function M_p, an azimuthal function N_p, and an attenuation A_p:

$$S\left(\theta_i, \theta_r, \phi_i, \phi_r\right) = \sum_{p=0}^{\infty} M_p(\theta_i, \theta_r) N_p(\theta_i, \phi_i, \phi_r) A_p(\theta_i, \phi_i). \quad (2)$$

$p = 0$ is the first reflection mode (R), $p = 1$ is the two transmission mode (TT), and $p = 2$ corresponds to the transmission, reflection, and transmission mode (TRT). Our model is based on [Chiang et al. 2015], where the longitudinal function is from [d'Eon et al. 2011]

$$M_p\left(\theta_\mathrm{i}, \theta_\mathrm{r}\right) = \frac{1}{2v\sinh(1/v)} \mathrm{e}^{-\frac{\sin\theta_\mathrm{i}\sin\theta_\mathrm{r}}{v}} I_0\left(\frac{\cos\theta_\mathrm{i}\cos\theta_\mathrm{r}}{v}\right), \quad (3)$$

with transformed longitudinal roughness v. $N_p(\theta_i, \phi_i, \phi_r)$ evaluates a trimmed logistic function around the perfect reflection / transmission direction. $A_p(\theta_i, \phi_i)$ computes the fresnel reflection and transmission contribution. The model relies on the Monte Carlo integration inherent in path tracing to integrate across the fiber width. While light can bounce an arbitrary number of times within a fiber, hence the upper bound of infinity in Equation 2, BCSDFs typically contain a finite number of modes.

3.2 Multiple fiber scattering

We assume that azimuthal and longitudinal scattering are separable. To understand the azimuthal scattering behavior, we simulate rays interacting with a single circle and multiple circles in the cross-sectional plane respectively. To simulate light interacting with multiple fibers, we create small circles that represent cross sections of fibers randomly inside a circular boundary that represents the thread. We assign the radius (R) of the thread, the radius (r) of the fibers, the minimum gap (d) between the fibers and the maximum number of attempts (M) as parameters for the simulation. In each attempt, we try adding a new circle within the boundary, which will be successful if it is at least d distance away from all existing ones. Otherwise, we skip to the next step until we have exhausted all attempts.

For each cross-sectional configuration, one million parallel rays pointing to the circle distributed uniformly across the thread diameter are generated to simulate specular reflection, specular transmission, and absorption. Inside the circle, absorption is modeled according to Beer's law. To eliminate the randomness inherent in any specific configuration, we generate 500 random circle configurations for each simulation and average these distributions to produce the final distribution. This can be seen as averaging the random configurations across different thread instances in a real fabric.

The left side of Figure 3 shows one multiple-circle configuration. On the right, we compare the azimuthal scattering function from a single circle with that from multiple circles with different small circle radii, averaged across various configurations. We observe that multi-circle scattering distribution exhibits weaker forward scattering and stronger backscattering. This is expected because, with more fibers present, it is more likely that light will encounter a backscattering event and exit in the reflected directions. We observe that the multiple-circle scattering distribution shares similar characteristics with R, TT, and TRT lobes of the single-circle scattering (Figure 4, left): each lobe appears to be modified independently, and they combine to form the multiple-circle scattering distribution.

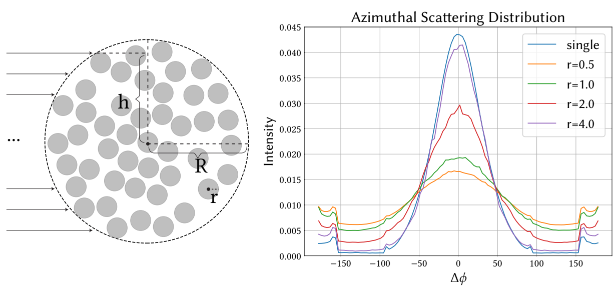

Fig. 3. The left figure illustrates a configuration of multiple circles. The right figure compares the azimuthal scattering functions of single and multiple fiber scattering, with varying small fiber radii, averaged over random configurations.

3.3 Thread BCSDF

Based on the simulation results, we propose a new model S_t for threads containing multiple fibers, comprising three lobes: a low-order backward scattering mode adapted from the R mode in single-fiber scattering, a forward scattering mode modifying the TT mode in single-fiber scattering, and a high-order backward scattering mode modifying the TRT mode in single-fiber scattering.

$$S_t\,(\theta_i,\theta_r,\phi_i,\phi_r) = \sum_{p=0}^{2} M_{tp}(\alpha_p,\theta_i,\theta_r) N_{tp}(\beta_p,\theta_i,\phi_i,\phi_r) A_{tp}(\theta_i,\phi_i),$$
$$A_{t0} = \lambda_0(\theta_i), \quad A_{t1} = \lambda_1(\theta_i) T_r, \quad A_{t2} = \lambda_2(\theta_i) T_r \tag{4}$$

where α_p and β_p are the longitudinal and azimuthal roughness for each lobe p, providing independent control of the width of the distribution; λ_0, λ_1, and λ_2 are scale factors with longitudinal incident angle dependence. $T_r = e^{-\sigma_a l}$ is the transmittance through a single fiber, where σ_a is the absorption coefficient and l is the distance traveled within the fiber. When interacting with multiple fibers, the path length accumulates, but it is impossible to analytically compute the exact path length for each direction, averaged over random fiber configurations. Our representation essentially enables us to find the equivalent scaling and width changes that account for the path length variations in multiple fiber scattering compared to single fiber scattering. We evaluate S_t in a non-integrated manner as proposed in [Chiang et al. 2015]. However, we assume that the thread is subpixel and that the scattering function, integrated over the width of the thread, exhibits the desired scattering behavior of multiple fiber scattering.

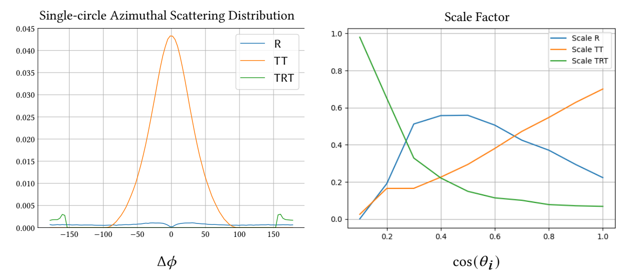

Fig. 4. The left figure shows R, TT, and TRT lobes of the single-circle scattering. The right figure shows the scale factors as functions of $\cos(\theta_i)$.

To verify our model, we first perform azimuthal function fitting tests for different longitudinal incident angles separately (top row of Figure 5). In each fitting, the target is the angular distribution obtained from the simulation. The parameters are the absorption coefficient (σ_a), the individual scale factors (λ_p), and the roughness parameters (α_p). The fitted curve is computed by evaluating for 1000 different height (h) values uniformly spaced across the width. At each height, we evaluate the angular distribution of each lobe using the logistic distribution, compute the weighted sum using the scale factors, and then average the distributions over h. We use the Adam optimizer and iteratively update the parameters until convergence. Our proposed model fits multiple scattering distribution very well, despite a small discrepancy towards the ends of the curves.

We observe that the scale factors $\lambda_p(\theta_i)$ follow specific trends for three lobes (Figure 4, right), leading to the following parametrization of the scale factors

$$\begin{aligned} \lambda_0(\theta_i) &= r_{0a}\cos(\theta_i)^2 + r_{0b}\cos(\theta_i) + r_{0c}, \\ \lambda_1(\theta_i) &= r_{1a}\cos(\theta_i), \\ \lambda_2(\theta_i) &= \frac{1}{r_{2a}\cos(\theta_i) + r_{2b}} + r_{2c}. \end{aligned} \tag{5}$$

With parameterized scale factors, we jointly fit different θ_i and present the results in the bottom row of Figure 5. Our model demonstrates significantly more accurate fitting results compared to the ply-level scattering model in [Zhu et al. 2023b]. As the incident angle becomes more oblique, we observe a large backscattering peak that eventually dominates the distribution. However, in [Zhu et al. 2023b], the azimuthal R lobe at the fiber level and the aggregated ply-level backscattering azimuthal lobe are parameterized using a uniform distribution. This leads to a flat fitting result, compromising accuracy across different incident angles. The PDF and sampling function of the thread BSDF are straightforward to implement, similar to the single-fiber scattering model but with distinct normalized attenuation PDFs.

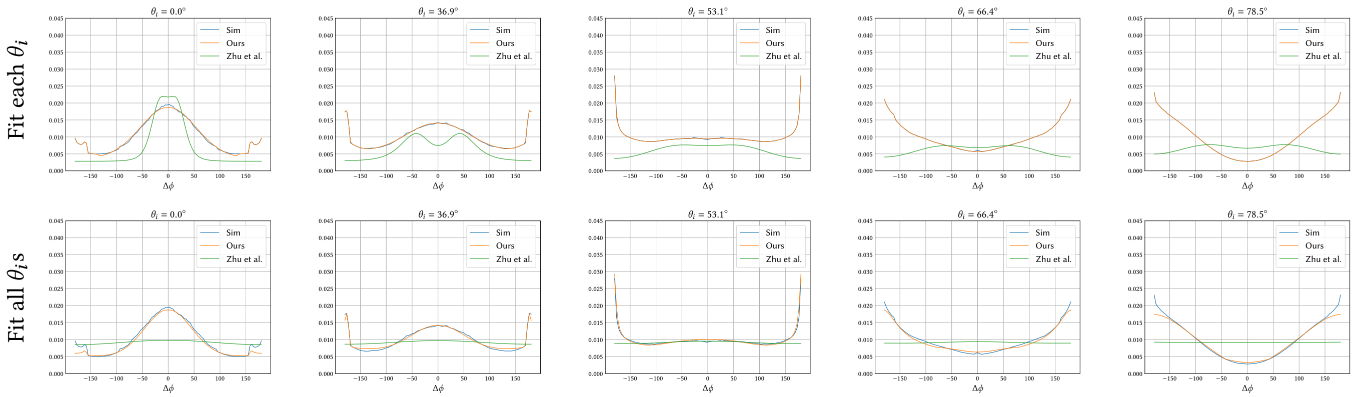

Fig. 5. We compare the fitting results of our model and those from [Zhu et al. 2023b] to the simulated azimuthal functions. In the top row, we fit each incident longitudinal angle separately and in the bottom row we fit different incident angles jointly. This comparison shows that our model produces much more accurate fitting than previous work.

4 DIGITIZE THREADS

Capture. As illustrated in Figure 2, our thread capture system comprises a Sony Alpha 7 II camera, a cell phone flashlight, a rotating platform, and a black hollow box around which the threads are wrapped. The measurement is done in a dark room where the cell phone flashlight is the only light source. The cellphone is placed in front of the rotating platform when capturing reflection and placed behind the rotating platform when capturing transmission. In both cases, we measure the relative positions of the camera and the flashlight with respect to the projection of the center of the rotating platform onto the table, where the origin of our coordinate system is located. As thread scattering exhibits angular variations, especially for materials like silk, we capture multiple views and configurations:

- To capture the reflection from the threads, we first position them in horizontal orientations and rotate the platform to capture six angles. Then, we position the threads vertically and capture the same set of angles.
- To capture the transmission of the threads, we position them horizontally only, as we observe that transmission changes are negligible when the threads are placed vertically and viewed from different angles. We capture five angles.

We collect a diverse dataset with 28 threads captured from 17 views. A selection of captured images are presented in Figures 6 and 7. The threads are tightly wrapped 50 times around a black hollow box, with care taken that there are no gaps between the threads.

Thread reconstruction. We leverage differentiable rendering and use Mitsuba 3 [Jakob et al. 2022] to conduct thread reconstruction. We reproduce the capture set up, modeling the flashlight as a single spherical emitter and representing the thread geometry using aligned B-spline curves. We first estimate the light intensity using a calibrated spectralon. Then we mask the reference photos using the non-zero regions from the initial rendering of the threads. This ensures that the optimization gradient focuses solely on the visual differences between the rendered output and the masked photo captures. We use the Adam optimizer with a learning rate of $1e^{-2}$ and 100 iterations. The optimization loss is computed as

$$\mathcal{L} = \sum_{i=0}^{16} w_i \|I_i - I_i^{\text{ref}}\|^2 + \gamma_0 \max(0, r_{0a} + r_{0b}) + \gamma_1 \max(0, -r_{2b}) + \gamma_2 \max(0, r_{0a}) \quad (6)$$

where $\sum_{i=0}^{16} w_i \|I_i - I_i^{\text{ref}}\|^2$ is the weighted sum of MSE loss from each view; $\gamma_0 \max(0, r_{0a}+r_{0b})$ and $\gamma_1 \max(0, -r_{2b})$ are regularization terms to avoid negative scale factors; $\gamma_2 \max(0, r_{0a})$ is the regularization term to favor a concave down shape for λ_0, following the trend from the simulation. We clamp r_{0c}, r_{1a} to be non-negative. In practice, we use $\gamma_0 = \gamma_1 = \gamma_2 = 1e^{-4}$.

5 SIMULATION AND RENDERING

Yarn geometry simulation. Yarn-level geometry is critical in reproducing the visual appearance of woven fabrics. [Li et al. 2024] represented yarn centerlines using parametric curves, while [Montazeri et al. 2020] employed a reduced simulation model where yarn deformation was restricted to the height direction. In contrast, our method utilizes a full physically-based simulation, enabling a more accurate representation of yarn geometry and its deformations.

We adopt the simulation method in [Leaf et al. 2018] to generate the geometric structure of yarns in woven fabrics. For each cloth sample, a repeating patch is extracted based on the woven pattern. The initial yarn shapes are defined as parameterized curves derived from this pattern. These initial shapes are then relaxed into their deformed configurations using physically-based simulation. The simulation incorporates stretching and bending energy terms, while yarn contraction is adjusted to account for tension introduced during the weaving process. Periodic boundary conditions are applied to ensure that the patch is tileable seamlessly. Once relaxed, the patch is tiled to construct and render the complete fabric. We optionally simulate twisted plies around yarns and explain this in detail in the supplemental.

Two-scale rendering. Rendering human-scale cloth is challenging, and the brute-force method of tiling the entire fabric is impractical. We implement the two-scale rendering method proposed in [Li et al.

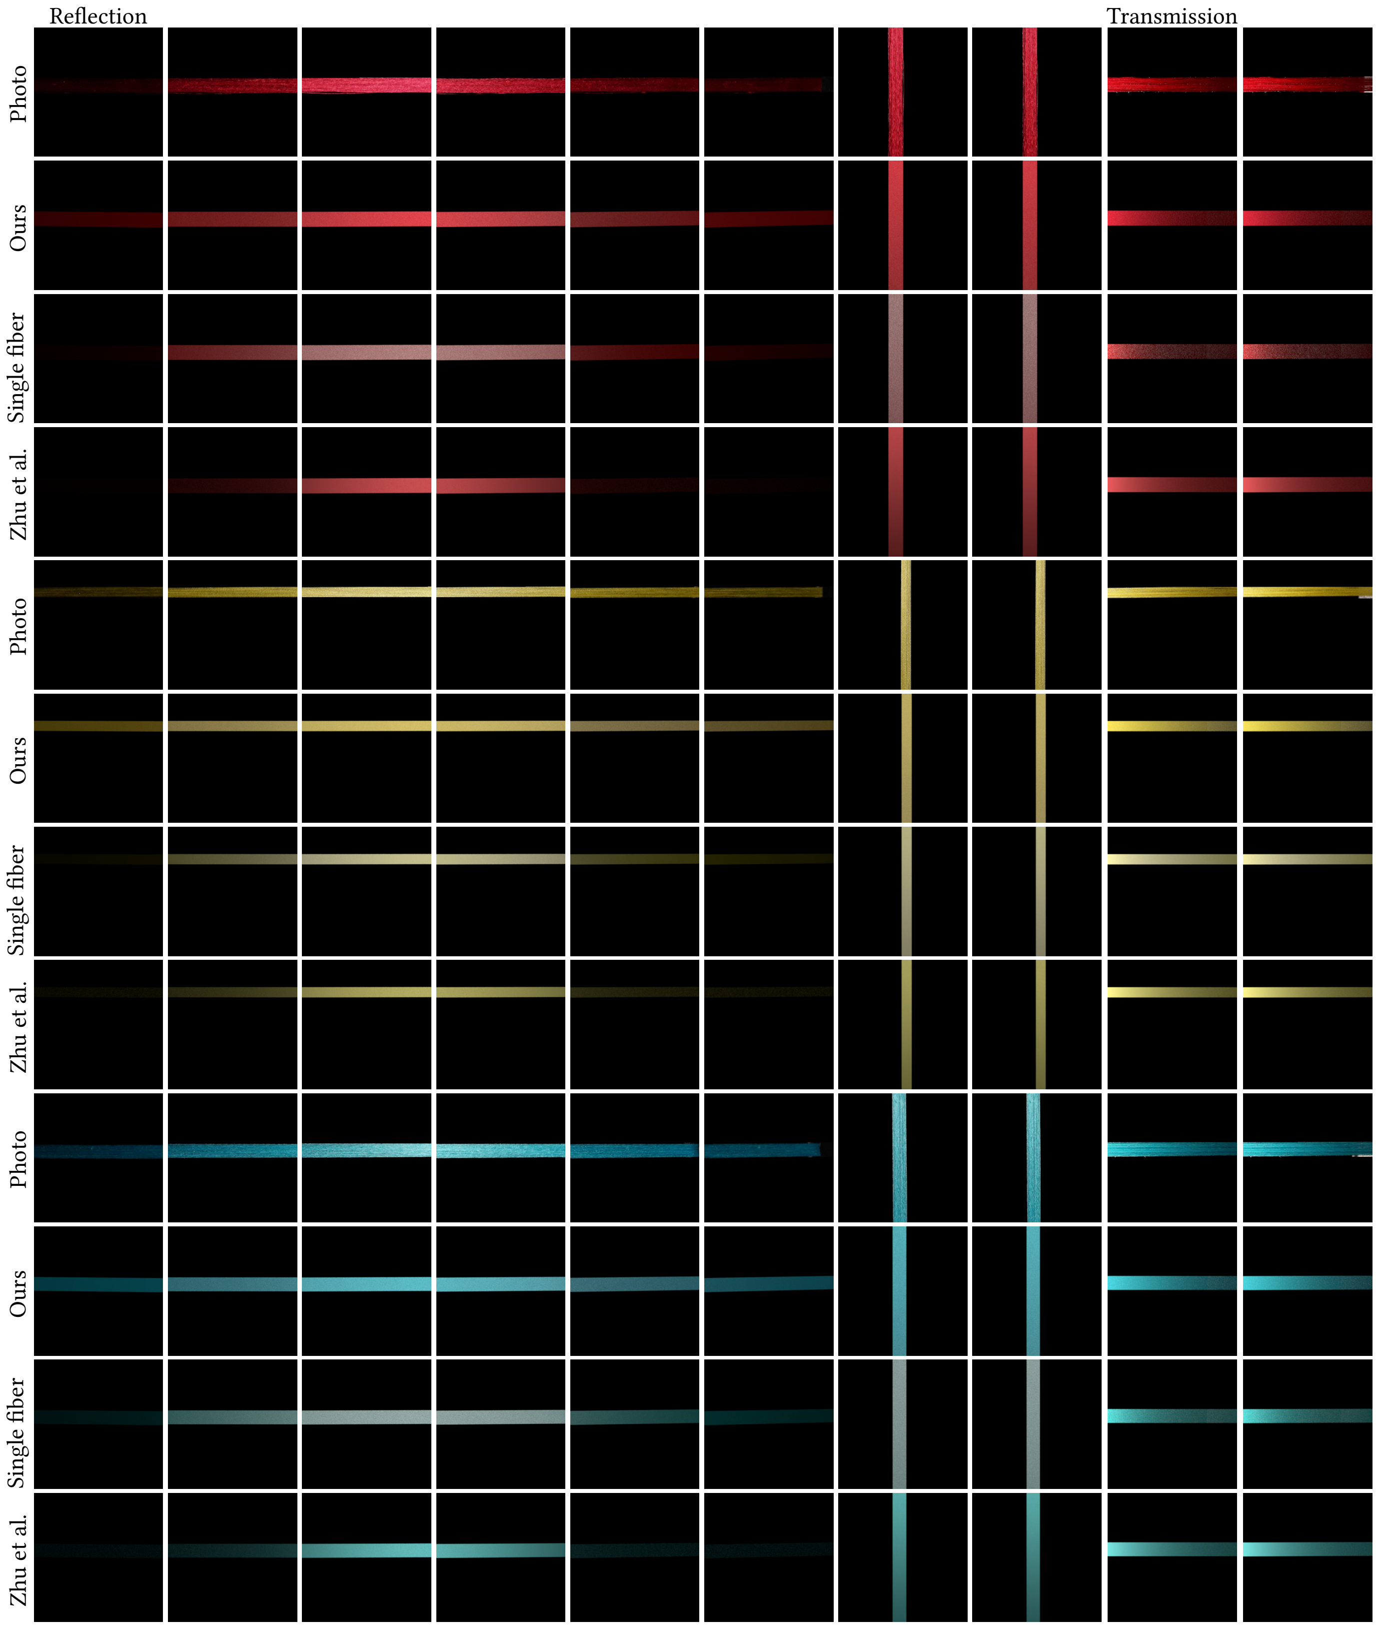

Fig. 6. We compare our model, the single fiber scattering model, and the model by Zhu et al. [2023b] on silk threads with various color. Our model produces significantly better reconstruction in terms of thread color and brightness variation. Please refer to the supplementary material for the complete set of results.

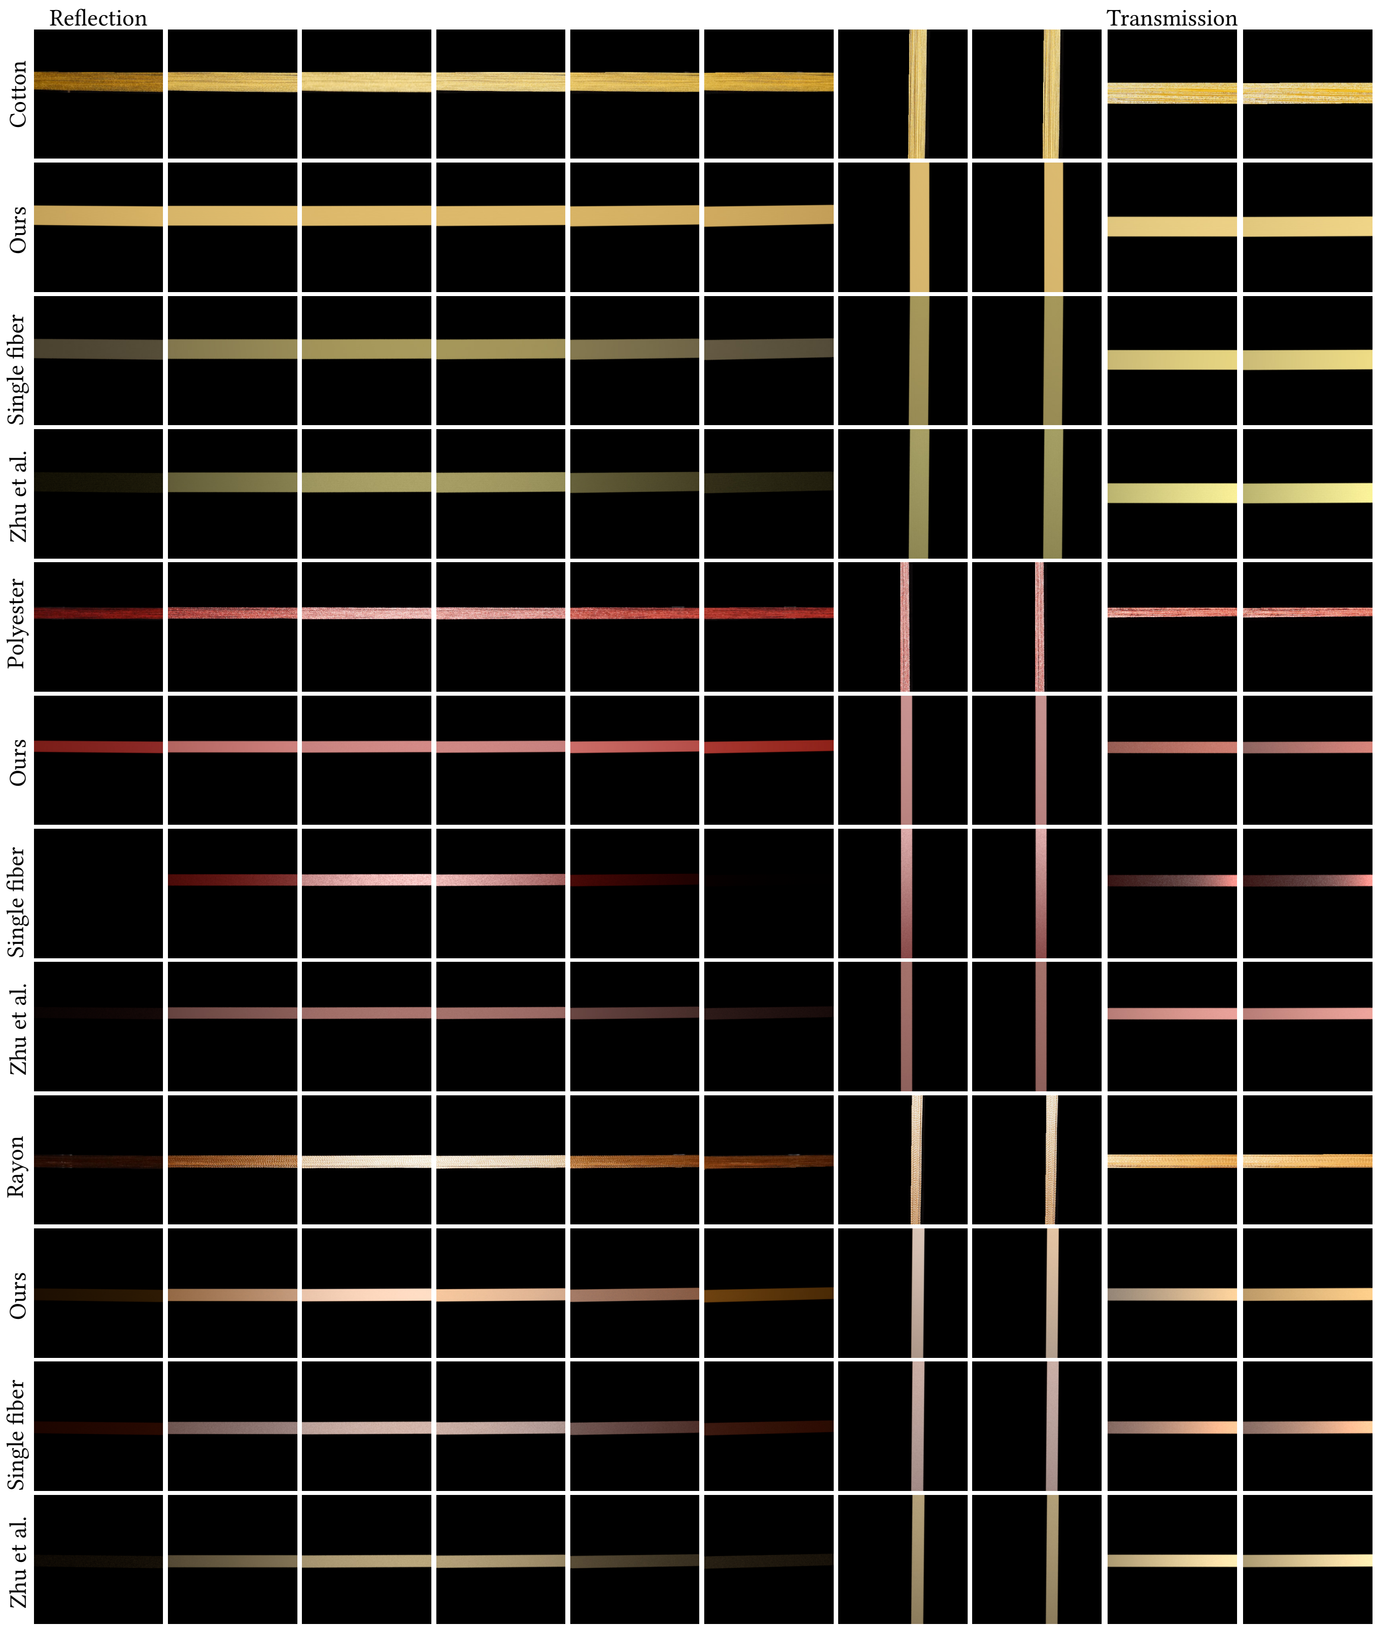

Fig. 7. We compare our model, the single-fiber scattering model, and the model by Zhu et al. [2023b] on cotton, polyester, and rayon threads. Our model consistently achieves better reconstruction across various views.

Fig. 8. This figure shows how our pipeline can aid design by displaying previews of skirts woven from various patterns and thread colors. The last row demonstrates that by arranging four differently colored threads in various ways, we can achieve a similar overall color but with distinct highlights.

2024] to efficiently render the fabric. Rendering a $1m^2$ large silk fabric with a structure similar to that in Figure 10 requires 6.6 GB of memory using the brute-force method. We assume that the woven fabric is tiled from a patch with 6 warp yarns and 6 weft yarns, and only one copy needs to be stored, which requires 6.6 KB of memory. The two-scale rendering approach considers a macro scale for the fabric surface, which is represented using a triangle mesh in the scene, and a micro scale for the yarns. We first generate rays from the camera and compute their intersection with the macro surface. Then, we sample a light direction and transform the intersection point and the directions into the local patch space, computing the intersection in that space. Next, we path tracing at the micro-scale until the ray exits the micro surface. We then transform the ray back into world space and continue the path tracing on the global level. This method allows us to achieve fast and accurate rendering of woven fabric. Each of the three renders in the teaser image (2000 x 2200 resolution, 256 spp) takes 7.8 seconds on an RTX 4090; Figure 12, with the same resolution and spp, takes 9.0 seconds.

6 RESULTS

Thread inversion. We perform thread reconstruction on threads of various colors and materials, including silk, rayon, cotton, polyester, and metallic fibers. A selection of thread inversion results is shown in Figures 6 and 7, with the full set of results provided in the supplemental material. We compare with inversion results using the single fiber scattering model, essentially treating each thread as a single fiber, and the results generated using [Zhu et al. 2023b]. Our method produces much more accurate reconstruction results because it is driven by simulation and able to fit multiple scattering behaviors within a thread. In contrast, the single fiber scattering model lacks backscattering, and [Zhu et al. 2023b] consists of various approximations of the actual multiple scattering behavior.

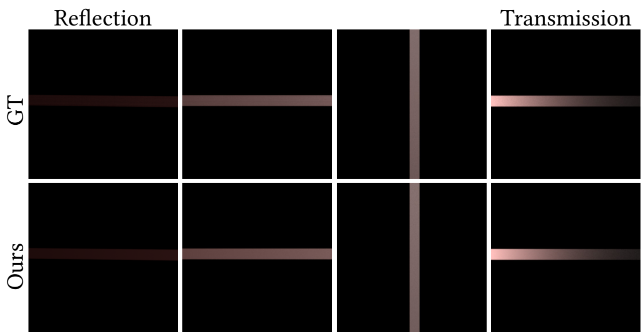

Fig. 9. We conduct thread inversion on synthetic data, where the ground truth is rendered by instantiating multiple fibers within a thread and path tracing them.

Synthetic validation. We validate thread inversion using synthetic data (selected views shown in Figure 9). Specifically, we create 10 fibers randomly distributed within each thread and render the same 17 views as in the thread inversion setup for real data. Each bundle of fibers is reconstructed as a single thread using the thread scattering function. We achieve high accuracy across all views.

Fabric appearance prediction. We acquire 16 fabric samples along with the corresponding threads used to fabricate those samples, and we validate that our pipeline can successfully predict cloth appearance, with two sets of results shown in Figures 10 and 11. Specifically, we gather four weaving patterns: plain, twill, satin, and basket weaves. For each pattern, there are four different color combinations. In all of them, the weft yarns are yellow, and the warp yarns are respectively blue, yellow, purple, and orange for the four samples. We wrap each fabric sample around a black cylinder with a radius of 4.1 cm and a height of 17.1 cm. We capture images using the same viewing configuration as in the reflection setting illustrated in Figure 2. The illumination conditions are the same as those in Figure 10, and the light source shifts from right to left in Figure 11.

To generate appearance predictions, we first capture and invert the corresponding threads, then run yarn simulations to generate tileable patches for each weave pattern. We adjust the warp and weft yarn periods according to their physical lengths in the fabric. We observe a highlight band with a yellow hue in the middle. As previously discussed in [Sadeghi et al. 2013], the highlight width depends on the slope of the underlying yarns, and we achieve the desired highlight width by adjusting the height variation across the patch. Our model reproduces the cloth appearance much more faithfully compared to the results generated using the single fiber scattering model and those of [Zhu et al. 2023b].

In Figures 1 and 8, we demonstrate how our work can aid in cloth design. We predict the appearance of a sari in three different configurations, each with satin, plain, and twill weave patterns in Figure 1. In each configuration, the warp and weft yarn scattering models are inverted from physical threads. We show the photos of the thread spools, the aligned threads, and their corresponding inversion results. In Figure 8, each column represents one weave pattern. In the first row, we show renders of skirts using four differently colored threads. In the second row, we weave the skirts using the threads of the four colors from the first row but with different warp and weft yarn ordering. We observe that these skirts display similar overall colors but with distinct highlight colors.

Textured patterns. We demonstrate that our model can also be used to preview cloth with textured patterns. Instead of using the color reparameterization proposed in [Chiang et al. 2015], we leverage differentiable rendering to determine the absorption coefficients for generating desired patterns. Given a texture pattern, we first render it as a diffuse plane under a constant environment. Under the assumption that clothes are made from threads with material properties similar to those of the captured thread but in different colors, we use the parameters of one acquired silk thread sample to initialize the optimization and invert only the absorption coefficient σ_a. Figure 12 showcases a traditional Chinese wedding dress rendered using this method.

Embroidery. Our method can help preview embroidery designs. We use previous work [Zhenyuan et al. 2023] to convert an embroidery motif (Figure 13 left) into a set of polylines, which are then converted into layered groups of linear curves to represent the embroidery geometry. Similar to the texture application, our method can preview the appearance of embroidery under the assumption that it uses the same materials as the captured threads but in different colors. To achieve this, we reproduce the reference motif in Mitsuba 3 using the converted embroidery geometry and constant lighting. We initialize the materials using the parameters of the captured thread and conduct inversion only on the absorption coefficient (σ_a) of each thread group while keeping other material parameters unchanged.

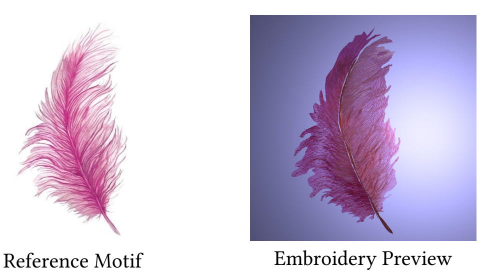

Fig. 13. An embroidery preview example using the parameters of a captured rayon thread sample.

7 DISCUSSION AND CONCLUSION

In this work, we present a novel pipeline that enables fabric appearance prediction based on actual physical threads. We develop a thread scattering model based on Monte Carlo simulations of multiple fiber scattering within a thread. We use a cost-efficient setup to capture threads and apply differentiable rendering to invert thread models. We conduct physical simulations to generate yarn geometry patches according to weaving patterns. We implement a two-scale rendering technique that enables fast and accurate rendering of fabric appearance. Our work tackles the challenge of predicting fabric appearance based on actual physical threads, which is greatly needed in the fashion industry.

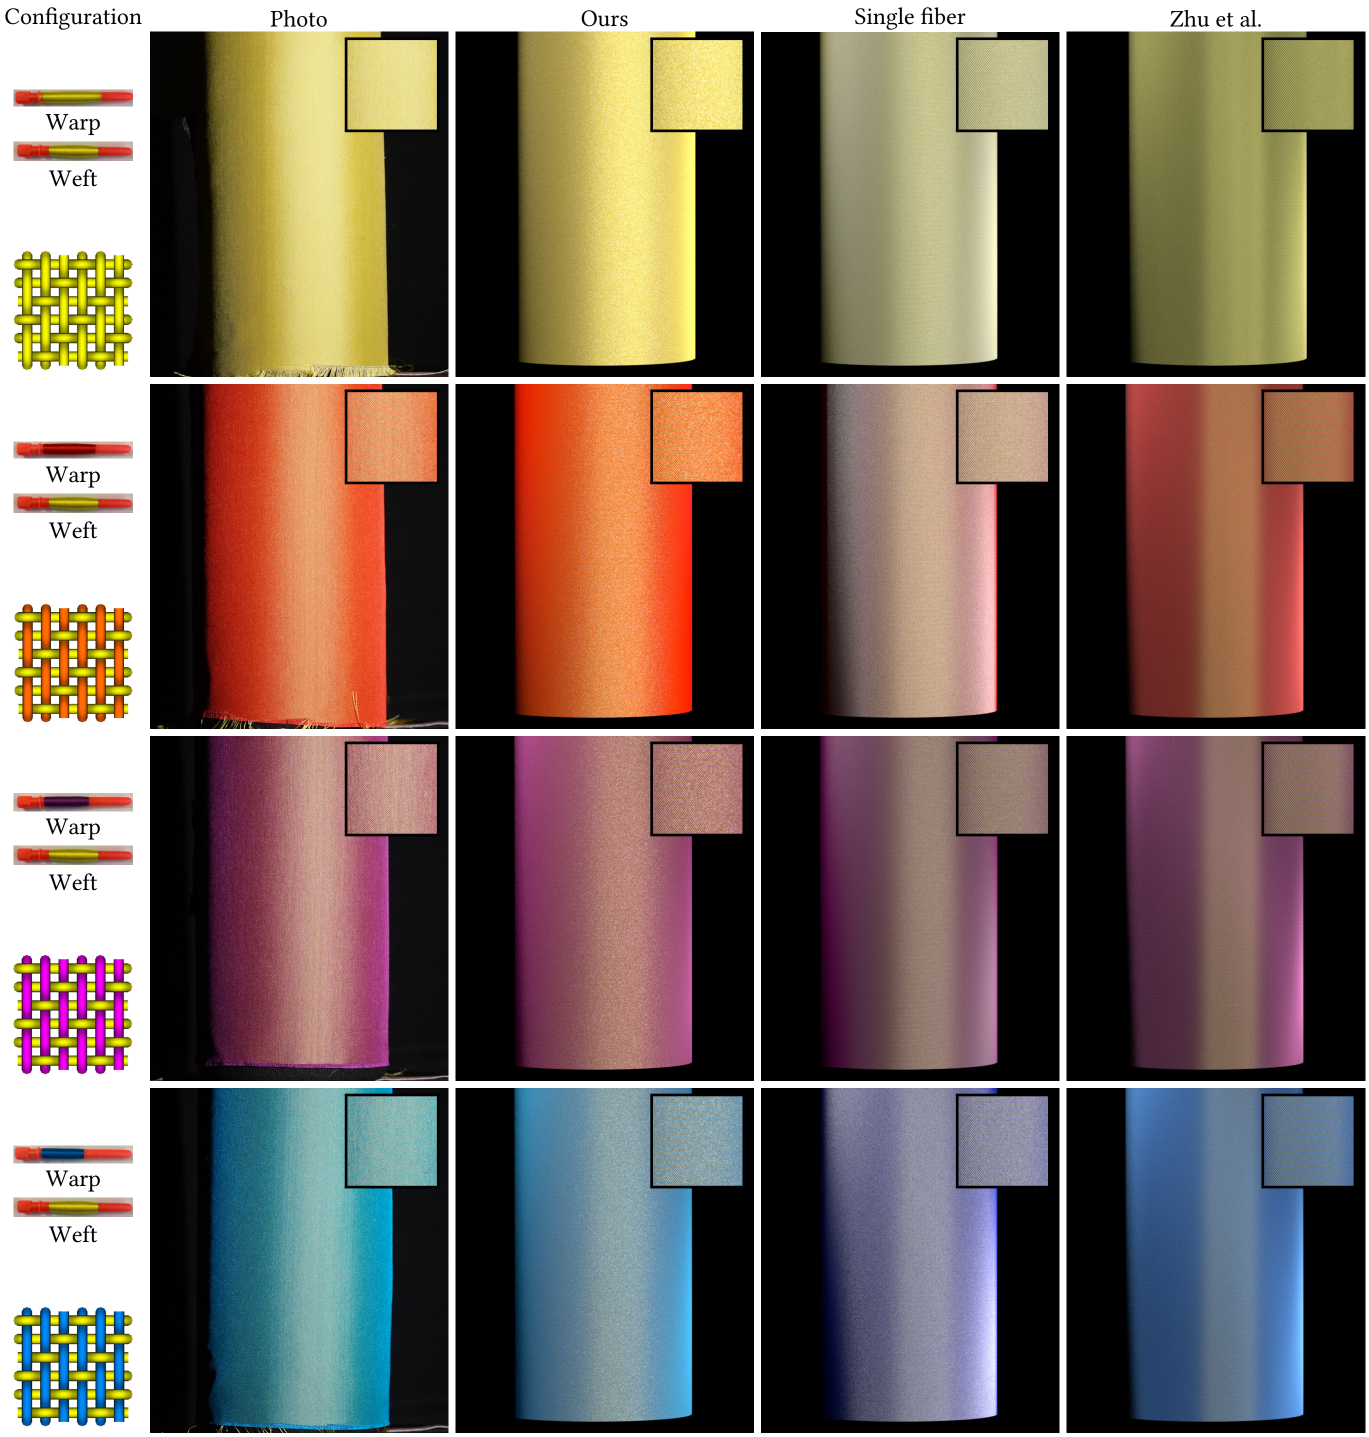

Fig. 10. We compare the appearance predictions of twill weave fabrics wrapped around a cylinder, illuminated from the right by a cell phone flash light. Our model produces significantly more accurate results than previous methods. Please refer to the supplementary material for the full set of results, including three additional weaving patterns.

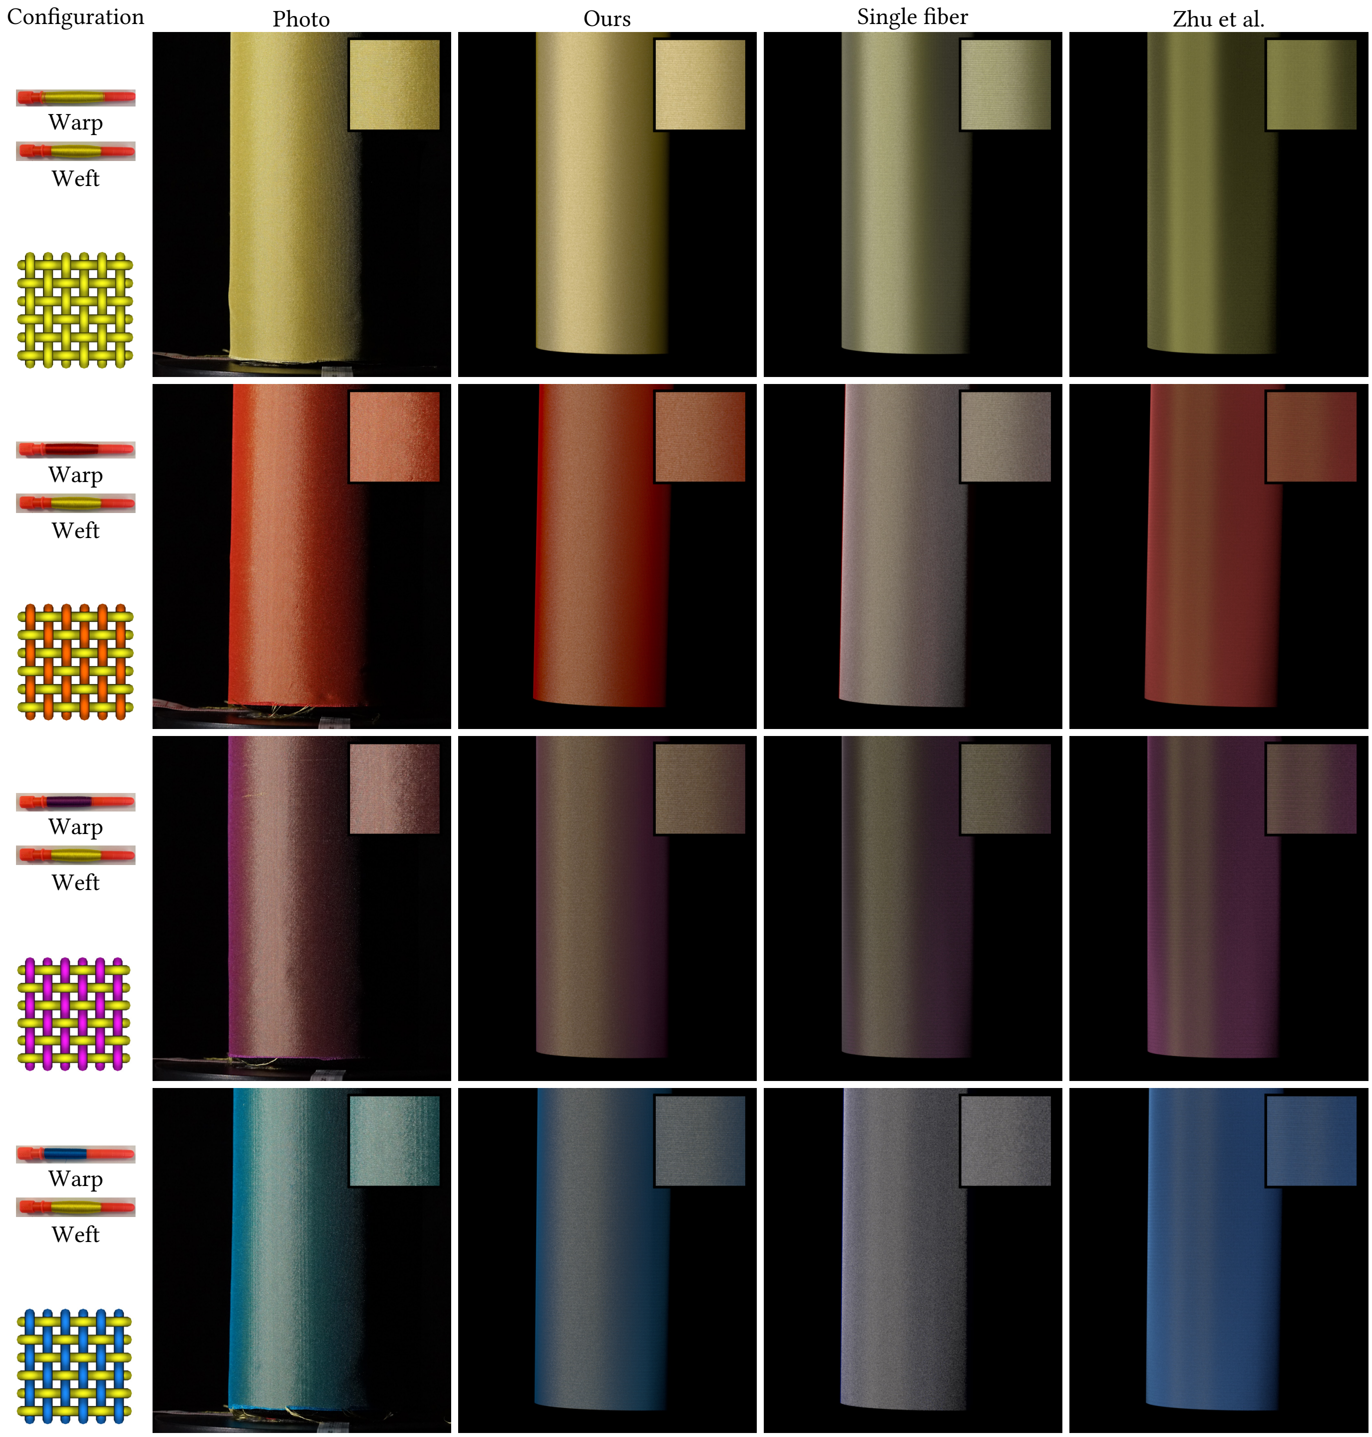

Fig. 11. This figure compares the appearance predictions of plain weave fabrics wrapped around a cylinder, illuminated from the left by a cell phone flashlight. Our model consistently achieves more accurate results than previous methods.


(a) (b) (c)

Fig. 12. Rendering textured patterns: (a) optimized woven plane with thread BSDF; (b) textured absorption coefficient after inversion; (c) rendering using the acquired thread scattering model with optimized absorption coefficient.

Limitations and future work. The current yarn simulation does not model the compression of the cross-section when weaving into the fabric, and a more accurate yarn simulation could lead to even more faithful fabric appearance predictions. Currently, we only consider RGB representation, and upgrading to spectral rendering could potentially produce more accurate color predictions, for example the blue fabric example in Figure 10. In future work, we would also like to explore representing irregularities in the appearance and modeling wave optical effects for both forward and inverse processes.

ACKNOWLEDGMENTS

We would like to thank Zhenyuan (Desmond) Liu for providing the embroidery example, and Xiaoji (Joy) Zhang for guidance on the simulation software. This work is supported by research grants from Tata Sons Private Limited, Tata Consultancy Services Limited, and Titan.

REFERENCES

Neeharika Adabala, Nadia Magnenat-Thalmann, and Guangzheng Fei. 2003. Real-time rendering of woven clothes. In *Proceedings of the ACM symposium on Virtual reality software and technology*. 41–47.

Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. 2015. A practical and controllable hair and fur model for production path tracing. In *ACM SIGGRAPH 2015 Talks*. 1–1.

Wentao Deng, Wei Ke, Zhongmin Deng, and Xungai Wang. 2024. Virtual design of woven fabrics based on parametric modeling and physically based rendering. *Computer-Aided Design* 173 (2024), 103717.

Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. 2011. An energy-conserving hair reflectance model. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 1181–1187.

Giuseppe Claudio Guarnera, Peter Hall, Alain Chesnais, and Mashhuda Glencross. 2017. Woven Fabric Model Creation from a Single Image. *ACM Trans. Graph.* 36, 5, Article 165 (Oct. 2017), 13 pages. https://doi.org/10.1145/3132187

Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX microflake distribution. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–11.

Piti Irawan and Steve Marschner. 2012. Specular reflection from woven cloth. *ACM Transactions on Graphics (TOG)* 31, 1 (2012), 1–20.

Wenzel Jakob, Adam Arbree, Jonathan T Moon, Kavita Bala, and Steve Marschner. 2010. A radiative transfer framework for rendering materials with anisotropic structure. In *ACM SIGGRAPH 2010 papers*. 1–13.

Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. https://mitsuba-renderer.org.

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* (Daegu, Republic of Korea) *(SA '22)*. Association for Computing Machinery, New York, NY, USA, Article 33, 8 pages. https://doi.org/10.1145/3550469.3555380

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.

Jonathan Leaf, Rundong Wu, Eston Schweickart, Doug L. James, and Steve Marschner. 2018. Interactive Design of Yarn-Level Cloth Patterns. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2018)* 37, 6 (11 2018). https://doi.org/10.1145/3272127.3275105

Zixuan Li, Pengfei Shen, Hanxiao Sun, Zibo Zhang, Yu Guo, Ligang Liu, Ling-Qi Yan, Steve Marschner, Milos Hasan, and Beibei Wang. 2024. Fiber-level Woven Fabric Capture from a Single Photo. arXiv:2409.06368 [cs.GR] https://arxiv.org/abs/2409.06368

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.

Stephen R Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light scattering from human hair fibers. *ACM Transactions on Graphics (TOG)* 22, 3 (2003), 780–791.

Zahra Montazeri, Søren B Gammelmark, Shuang Zhao, and Henrik Wann Jensen. 2020. A practical ply-based appearance model of woven fabrics. *ACM Transactions on Graphics (TOG)* 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 Transactions on Graphics (TOG)* 32, 2 (2013), 1–12.

Kai Schroder, Reinhard Klein, and Arno Zinke. 2011. A volumetric approach to predictive rendering of fabrics. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 1277–1286.

Kai Schröder, Arno Zinke, and Reinhard Klein. 2014. Image-based reverse engineering and visual prototyping of woven cloth. *IEEE transactions on visualization and computer graphics* 21, 2 (2014), 188–200.

Yingjie Tang, Zixuan Li, Milos Hasan, Jian Yang, and Beibei Wang. 2024. Woven Fabric Capture with a Reflection-Transmission Photo Pair. In *ACM SIGGRAPH 2024 Conference Papers*. 1–10.

Beibei Wang, Wenhua Jin, Miloš Hašan, and Ling-Qi Yan. 2022. SpongeCake: a layered microflake surface appearance model. *ACM Transactions on Graphics (TOG)* 42, 1 (2022), 1–16.

Hongzhi Wu, Julie Dorsey, and Holly Rushmeier. 2011. Physically-based interactive bi-scale material design. *ACM Transactions on Graphics (TOG)* 30, 6 (2011), 1–10.

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.

Mengqi Xia, Bruce Walter, Eric Michielssen, David Bindel, and Steve Marschner. 2020. A wave optics based fiber scattering model. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.

Ling-Qi Yan, Henrik Wann Jensen, and Ravi Ramamoorthi. 2017. An efficient and practical near and far field fur reflectance model. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–13.

Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. Physically-accurate fur reflectance: Modeling, measurement and rendering. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–13.

Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building volumetric appearance models of fabric using micro CT imaging. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–10.

Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-aware synthesis for predictive woven fabric appearance. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.

Shuang Zhao, Fujun Luan, and Kavita Bala. 2016. Fitting procedural yarn models for realistic cloth rendering. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–11.

Liu Zhenyuan, Michal Piovarči, Christian Hafner, Raphaël Charrondière, and Bernd Bickel. 2023. Directionality-Aware Design of Embroidery Patterns. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, 397–409.

Junqiu Zhu, Christophe Hery, Lukas Bode, Carlos Aliaga, Adrian Jarabo, Ling-Qi Yan, and Matt Jen-Yuan Chiang. 2024. A Realistic Multi-scale Surface-based Cloth Appearance Model. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) *(SIGGRAPH '24)*. Association for Computing Machinery, New York, NY, USA, Article 89, 10 pages. https://doi.org/10.1145/3641519.3657426

Junqiu Zhu, Adrian Jarabo, Carlos Aliaga, Ling-Qi Yan, and Matt Jen-Yuan Chiang. 2023a. A realistic surface-based cloth rendering model. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–9.

Junqui Zhu, Zahra Montazeri, J Aubry, Lingqi Yan, and Andrea Weidlich. 2023b. A Practical and Hierarchical Yarn-based Shading Model for Cloth. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, e14894.

Junqiu Zhu, Sizhe Zhao, Lu Wang, Yanning Xu, and Ling-Qi Yan. 2022. Practical level-of-detail aggregation of fur appearance. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–17.