Matching Real Fabrics with Micro-Appearance Models

PRAMOOK KHUNGURN, DANIEL SCHROEDER, SHUANG ZHAO, KAVITA BALA, and STEVE MARSCHNER
Cornell University

Micro-appearance models explicitly model the interaction of light with microgeometry at the fiber scale to produce realistic appearance. To effectively match them to real fabrics, we introduce a new appearance matching framework to determine their parameters. Given a micro-appearance model and photographs of the fabric under many different lighting conditions, we optimize for parameters that best match the photographs using a method based on calculating derivatives during rendering. This highly applicable framework, we believe, is a useful research tool because it simplifies development and testing of new models.

Using the framework, we systematically compare several types of micro-appearance models. We acquired computed microtomography (micro CT) scans of several fabrics, photographed the fabrics under many viewing/illumination conditions, and matched several appearance models to this data. We compare a new fiber-based light scattering model to the previously used microflake model. We also compare representing cloth microgeometry using volumes derived directly from the micro CT data to using explicit fibers reconstructed from the volumes. From our comparisons, we make the following conclusions: (1) given a fiber-based scattering model, volume- and fiber-based microgeometry representations are capable of very similar quality, and (2) using a fiber-specific scattering model is crucial to good results as it achieves considerably higher accuracy than prior work.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Color, shading, shadowing, and texture*; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—*Physically based modeling*

General Terms: Experimentation, Measurement

Additional Key Words and Phrases: Appearance modeling, fibers, cloth, parameter fitting

ACM Reference Format:

Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching real fabrics with micro-appearance Models. *ACM Trans. Graph.* 35, 1, Article 1 (December 2015), 26 pages.
DOI: <http://dx.doi.org/10.1145/2818648>

This work was supported in part by the National Science Foundation (grants IIS-1513967, IIS-1161645, and IIS-1011919) and the Intel Science and Technology Center for Visual Computing.

Authors' addresses: P. Khungurn (corresponding author), D. Schroeder, S. Zhao, K. Bala, and S. Marschner, Department of Computer Science, Cornell University, Ithaca, NY; corresponding author's email: pramook@cs.cornell.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

2015 Copyright held by Owner/Author(s)

0730-0301/2015/12-ART1

DOI: <http://dx.doi.org/10.1145/2818648>

1. INTRODUCTION

Appearance models that can reproduce the rich appearance of fabric are important in a wide range of applications including textile design, product visualization, retail, and entertainment. Yet, photorealistic rendering of fabrics remains very challenging. Recent research that models fabrics at the scale of fibers [Zhao et al. 2011; Schröder et al. 2014] has produced the most realistic renderings to date. By directly modeling the geometric arrangement of fibers, these methods can reproduce distinctive specular highlights caused by woven structures, subtle diffuse effects of multiple scattering, and details like fuzz and flyaway fibers.

These appearance models are instances of an approach we call *micro-appearance models*, which combines an explicit model of a material's microgeometry with a simple light scattering model. Previous work has shown that these components together can reproduce the subtle and complex light-scattering behaviors seen at larger scales. The approach is also general enough to encompass all types of textiles, including traditional weaves and knits as well as non-woven fabrics.

Using micro-appearance models entails answering a number of questions, which previous work has left open.

- How to represent microgeometry?* Microgeometry can either be represented by volumes (e.g., [Zhao et al. 2011]), or collections of individual fibers (e.g., [Schröder et al. 2014]). It is unclear which approach is better.
- How to model light scattering?* Zhao et al. [2011] proposed a simple microflake phase function, but it does not capture scattering behavior in grazing configurations correctly. Schröder et al. [2014] used a scattering function derived from that of hair fibers, which may handle grazing behaviors better. However, its effectiveness in this regard has never been assessed against measurements.
- How to compute model parameters?* Zhao et al. [2011] employs a simple binary search to fit their model, but it cannot be generalized to more complicated ones. Most other previous work specifies parameters manually with the exception of Schröder et al. [2014], which automatically derives the diffuse color from cloth photographs. To our knowledge, there has been no general framework for fitting model parameters.

This article attempts to answer these three questions. To do so, we introduce a number of innovations. To identify the best microgeometry model, we develop an algorithm that converts a micro CT scan of cloth fabric to an explicit mesh of the fibers that compose the fabric so that we can compare the effectiveness of the two approaches. To identify the best light scattering model, we use the methodology for developing hair scattering function [Marschner et al. 2003] to develop an improved scattering model for textile fibers that takes into account reflection from and refraction through fiber surfaces. To compute the parameters of the model, we develop an appearance matching framework which takes into account multiple fabric observations under different lighting conditions. It uses stochastic gradient descent to optimize the parameters, so it is general enough to fit any parameter with respect to

which the partial derivative of a single path tracing sample can be computed.

These innovations enabled us to systematically evaluate micro-appearance modeling approaches against one another and against measurements. In general, we gather ground-truth data in the form of gonireflectometric measurements of real cloth samples. The appearance matching framework uses some of the data to optimize for parameter values. Performance can then be assessed by comparing the rest of the ground-truth data to renderings yielded by the models and the fitted parameters. We believe that this evaluation procedure is essential to developing and effectively testing new models. Using the procedure, we provide an extensive evaluation of the effectiveness of our appearance matching pipeline and the relative performance of different microgeometry and light scattering models using six fabrics with very different characteristics.

The conclusions we draw from our investigation are as follows.

- (1) A scattering model based on previous models for hair scattering works much better than the microflake model, especially at reproducing bright grazing highlights.
- (2) In choosing between fiber- and volume-based models for cloth microgeometry, there is no clear winner: both are capable of matching measurements when used with the right scattering model. However, we did find that smooth orientation fields are important to achieving good results with volume models.
- (3) Our system worked well for a number of fabrics, but we also experienced a number of difficult cases, which point out venues for improving both the system itself and the light scattering model.

Our work as a whole comprises a complete and practical appearance modeling system which we believe is an important step forward in achieving predictivity and photographic realism for textiles. It can generate fabric models that capture both far-field reflectance properties and near-field fine textures of different types of textiles, as demonstrated in Figure 1.

The system implemented in this article, however, does not cover all aspects of fabric appearance. The experiments presented here focus on light *reflected* from fabric and do not examine light *transmitted* through the fabric. Moreover, our system currently only handles fabrics with a single yarn type and color. While it can be readily extended to address these aspects, the extension is not in the scope of this article and is left for future work.

2. PREVIOUS WORK

A system for modeling fabric appearance must specify two components: a representation for fabric geometry and a model for how light scatters from that geometry. Additionally, a practical and complete system should offer a way to infer model parameters from observations of real cloth samples.

2.1 Representations of Fabric Geometry

In graphics, three approaches have been proposed to represent fabric geometry. The first approach is to abstract the fabric to a two-dimensional surface, represented by a mesh or curved surface [Sattler et al. 2003; Adabala et al. 2003; Irawan and Marschner 2012]. Surfaces can be quite successful in distant views, but close-up views, especially at edges and silhouettes, look incorrect because the three-dimensional structure of yarns and fibers is missing.

The second approach is to use volumetric data. Kajiya and Kay [1989] pioneered this technique in the context of fur rendering.

Fig. 1. Rendering of a scene with (a) velvet, (b) silk, and (c) twill fabric models produced by our pipeline. *The exposure of the (a2) image was scaled up by 8 times to make the fibers visible. The velvet uses volume-based microgeometry, and the other two use fiber-based microgeometry.

Xu et al. used the approach to represent yarns in knitwear [2001]. Zhao et al. [2011] obtained fabric volumes from micro CT scans of real cloth samples and were able to reproduce detailed and irregular appearance of cloth at an unprecedented level. Schröder et al. [2011] generated cloth fibers procedurally and converts them to a volumetric representation.

The third approach is to represent fabric with a collection of discrete fibers, each represented with explicit geometry. While it has been employed to model materials with much larger fiber-like units such as a head of hair [Ward et al. 2006] and feathers [Chen et al. 2002], only a few previous works have used it to represent woven cloth. Zhang et al. [2013] proposed a new variant of subdivision surfaces that procedurally generates fibers on the fly at the right level of detail. Schröder et al. [2014] also developed another procedural model for fiber generation but generated the complete fiber mesh before rendering.

2.2 Models for Scattering

To model scattering based on a surface model, a popular option is to employ bidirectional texture functions (BTFs) [Sattler et al. 2003].

Table I. Approaches to Fabric Appearance Modeling

Appearance model	Geometry model		
	Fabric mesh	Volume	Fiber mesh
BTF	[Sattler et al. 2003]		
BSDF	[Adabala et al. 2003] [Irawan and Marschner 2012] [Sadeghi et al. 2013]		[Zhang et al. 2013]
Volumetric		[Zhao et al. 2011]	
Fiber-based		[Kajiya and Kay 1989] [Schröder et al. 2011] This work	[Schröder et al. 2011]* [Schröder et al. 2014] This work

*Schröder et al. used the fiber geometry coupled with fiber-based appearance models to produce ground truth for comparison with their volumetric geometry representation.

This approach captures view-dependent appearance by exhaustive sampling, but often suffers from undersampling and shows limited quality at edges and silhouettes and under grazing illumination.

The second approach, also coupled with the surface approach to geometry representation, is to use a surface-based BSDF model coupled with textures. Adabala et al. [2003] generated textures from weave pattern data and modeled scattering behavior with a microfacet BSDF. Irawan and Marschner [2012] proposed a procedural texture model based on weave patterns and a reflectance model based on the analysis of specular reflection from spun fibers. Sadhegi et al. [2013] constructed cloth BSDFs from weave patterns and a yarn scattering model.

The third approach is to specify volumetric appearance based on scattering particles, as dictated by the radiative transfer equation (RTE), by specifying the volumetric absorption coefficient, volumetric scattering coefficient, and a phase function. Zhao et al. [2011] modeled appearance of cloth with an anisotropic microflake phase function, the foundation of which was laid by Jakob et al. [2010].

The last approach employs light scattering models for fibers. The first of such models was proposed by Kajiya and Kay [1989] and is capable of reproducing the characteristic sheen in fur and hair. Most fiber scattering models, including ours, are based on the work of Marschner et al. [2003], which introduced the *bidirectional curve scattering distribution function* (BCSDF) to encapsulate the far-field appearance of hair fibers. Zinke and Weber [2007] generalized the BCSDF to the *bidirectional fiber scattering distribution function* (BFSDF) to make it capable of capturing near-field appearance, which is subsequently used unmodified by Schröder et al. [2014]. Schröder et al. [2011] used the BFSDF with fiber-based geometry and a derived volume-based approximation. However, because the shift to volume rendering caused changes in appearance, they had to include special corrections in their volumetric rendering algorithm.

Table I summarizes the range of approaches to fabric modeling and situates our work in this space. In particular, we generate fibers from micro CT scan data instead of creating them procedurally. Our scattering model is similar to models for yarn scattering used in surface scattering models [Irawan and Marschner 2012; Sadeghi

et al. 2013] but is adapted to the BCSDF framework, making it suitable for rendering with explicit fibers.

2.3 Parameter Estimation

For the volume-based cloth model of Zhao et al. [2011], the authors selected rendering parameters for the scanned cloth volumes by binary searching the values of single parameters in sequence to match statistics of a reference image of the cloth sample. By contrast, our method recovers all parameters of the fiber scattering model simultaneously and compares the rendered cloth to several photographs under different lighting and viewing conditions. As in our approach, Gkioulekas et al. [2013] recover the unknown parameters of their scattering model through stochastic gradient descent. Their method combines gradient descent and Monte Carlo rendering to recover linear combinations of predefined materials from a material dictionary for a wide range of translucent media. Recently, Schröder et al. [2014] introduced an image-based technique to reverse engineer physical fabric samples. Their approach, however, focuses on recovering weave patterns and diffuse yarn colors and requires many other parameters needed for rendering to be specified manually. Sadeghi et al. [2013] estimated parameters in two stages. They first determined their yarn scattering model's parameters to match to a dense sample of yarn BSDF measurements, and then they determined parameters of the yarn curves. Both stages, however, were carried out manually.

We note that there are two approaches to the number of training data used for parameter fitting. On one hand, works such as Zhao et al. and ours use a *sparse* sampling of appearance. Indeed, Zhao et al. use one photograph whose pixel values are then averaged into three numbers. We use 16 (as will be discussed in Section 5.2), and the images are averaged in the same way. On the other hand, Sadeghi et al. and Gkioulekas et al. use a *dense* sampling.¹ One criticism to the sparse sampling approach is that the samples might not be representative enough to capture all the details of the fabric's scattering behaviors. For example, Sadeghi et al. documented that fabrics exhibit multiple highlights due to weave patterns and not including such highlights may lead to incorrect parameter values, especially lobe widths.

It is instructive to compare our approach to that of Sadeghi et al. [2013]. Appearance is produced by a combination of microgeometry and optical properties, and both methods seek to determine geometry and scattering properties separately. Sadeghi et al. consider the yarn as the basic unit, so they begin by measuring the optical properties of a yarn very accurately and then finish by adjusting the geometry to match the overall appearance. We consider the fiber as the basic unit and proceed in the opposite order. We establish geometry first—both the arrangement of fibers within yarns and also the geometry of yarns in the cloth—using micro CT scanning and then adjust the scattering properties to match appearance. In a sense, we rely on detailed microgeometry to give rise to complex structures in fabric's scattering behaviors such as multiple highlights and to compensate for approximated lobe widths values. Encouraged by the success of Zhao et al. and constrained by the need to render images during the fitting process, we have taken the approach of fitting the optical properties to the low-frequency BRDF. As discussed in Section 7, this approach produced excellent results for many fabrics, but there are others where additional training configurations may be helpful.

¹Gkioulekas et al. take only 18 photographs, but each pixel of each photograph is considered a measurement, so the sampling is in fact very dense.

2.4 Derivative Estimation

Our appearance matching framework extends a path tracer to estimate derivatives of images with respect to parameters. Hařan and Ramamoorthi [2013] estimate derivatives with respect to the albedos of volumes. Our method reduces to the same form when applied to estimate derivatives with respect to albedo but supports other parameters as well. The operator-theoretic method of Gkioulekas et al. [2013] considers the full set of rendering parameters of an isotropic volume but only obtains derivatives with respect to the coefficients of a convex linear combination of predefined materials. By comparison, the approach in Section 5 can individually estimate derivatives with respect to any rendering parameter that has a differentiable effect on the samples used to render the image.

2.5 Fiber Generation

Previous works that employ fiber-based geometry representations generate textile fibers procedurally, using mathematical models such as that proposed by Srebrateep and Bohez [2006].

Researchers in the hair modeling community have presented many works that generate hair fibers from different types of data such as a single photograph [Chai et al. 2013] (for 2D manipulation), multiple photographs from different views [Paris et al. 2004; Luo et al. 2013], focus stacks [Jakob et al. 2009], and thermal images [Herrera et al. 2012]. All the cited works generate hair strands by tracing particles along a direction field defined by the captured data. In contrast, we take volumetric data as input and generate fibers by identifying locations of fiber centers and then connecting them globally instead of sequentially tracing particles.

Based on recent advances in textile research, Schröder et al. [2014] proposed a technique to build fiber geometry for woven fabrics. Their approach estimates yarn curves using one photograph taken under uniform lighting and procedurally generates fibers around them.

Constructing fiber geometry from volumetric data has been well studied in the neuroscience community for the purpose of understanding the brain’s networking. Approaches include tracing particles through a direction field [Basser et al. 2000], growing a level set [Parker et al. 2002], and simulating water diffusion in the volume [Kang and Zhang 2005].

While Shinohara et al. used micro CT images to extract yarn positions in fabric [2010], we are not aware of any previous works, especially in the graphics community, that generate cloth fibers from micro CT images.

3. OVERVIEW

Our appearance modeling pipeline is outlined in Figure 2. Our appearance matching process takes as input a set of photographs of a fabric under different lighting and viewing configurations, together with the corresponding scene geometry, and finds fits for the parameters of our new light scattering model. The light scattering model is described in Section 4. The appearance matching process is introduced in Section 5. The appearance matching process takes input as photographs and fiber microgeometry, the latter of which is constructed, as described in Section 6.

We now describe each part of the pipeline in more detail. First, in Section 4, we introduce a new light scattering model for textile fibers. The scattering model has two terms. The first term models light reflected directly off fiber surfaces and contains a Fresnel term that makes the reflection brighter at grazing angles to address the inaccuracy observed in Zhao et al.’s model. The second term models light transmitted forward through the fibers, accounting

Fig. 2. Appearance modeling pipeline.

for the fact that textile fibers are generally translucent. While the scattering model is conceptually simple, it cannot be fitted using the simple iterative binary search method described in Zhao et al. [2011]. Therefore, we introduce a new appearance matching process in Section 5.

The appearance matching is done using gradient descent optimization to find values for the parameters of the scattering model that achieve the best match between the photographs and physically-based renderings of the cloth microgeometry model. Multiple scattering contributes significantly to the renderings (the majority of the reflected light is due to multiple scattering in most cases). Thus, the optimization has to account for it. To do so, we extend our renderer, a Monte Carlo path tracer, to compute derivatives of the output image with respect to the parameters of the scattering model. The derivatives are computed as unbiased estimates of the true derivatives and are used in a *stochastic gradient descent* optimization method, which converges to a minimum despite the uncertainty in the individual estimates of pixel values and their derivatives (under the condition that the objective function is convex). Our process is agnostic both to the microgeometry model and to the scattering model, allowing us to directly compare the abilities of different models to recover cloth appearance.

A goal of this article is to study the effectiveness of two representations of fabric microgeometry. While previous work has been able to create micron-scale volumetric representation from micro CT images of real fabrics, no work has addressed the creation of fiber mesh representations from such data. To enable direct comparison of the two approaches, we develop an algorithm to construct fiber meshes from micro CT images in Section 6. The method is based on identifying fiber centers in slices of micro CT volumes and connecting them.

Finally, in Section 7, we present the results of our investigation into fabric appearance models, carried out using the tools developed in the earlier sections, and draw conclusions about which methods should be used.

4. FIBER SCATTERING MODEL

In this section, we describe a fiber scattering model, developed with the goal of addressing the shortcomings of the microflake scattering function used by Zhao et al. [2011]. Our model builds upon the considerable research in scattering models for rendering hair [Marschner et al. 2003; Zinke and Weber 2007; d’Eon et al. 2011]. Like hair, textile fibers are long cylindrical structures made of dielectric material, so they can be expected to exhibit similar

Fig. 3. The two modes of our scattering function.

specular reflection geometry. However, because textile fibers are smaller and less visible individually, as well as more irregular in cross section, we use a simpler model than the full hair model.

4.1 Preliminaries

We adopt the radiometry and notation of Marschner et al. [2003], with the fiber positioned on the x -axis and a direction ω represented in spherical coordinates θ and ϕ such that

$$\omega = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \sin \theta \\ \cos \theta \cos \phi \\ \cos \theta \sin \phi \end{bmatrix}. \quad (1)$$

The angle θ is called the *longitudinal angle* and ϕ the *azimuthal angle*. A fiber scattering model is a function $S(\omega_i, \omega_o)$ that denotes the ratio of the outgoing *curve radiance* in the direction ω_o to the incoming differential *curve irradiance* from direction ω_i .

Like most hair scattering models, ours is a sum of *separable modes*:

$$S(\omega_i, \omega_o) = \sum_{p=0}^k S_p(\omega_i, \omega_o) = \sum_{p=0}^k M_p(\theta_i, \theta_o) N_p(\theta_i, \phi_i, \phi_o). \quad (2)$$

Here, each of the *per-mode scattering functions* $S_p(\omega_i, \omega_o)$ is a product of a longitudinal function $M_p(\theta_i, \theta_o)$ and an azimuthal function $N_p(\theta_i, \phi_i, \phi_o)$. This separability is a significant convenience for importance sampling. A mode describes a component of scattered light that survives the same sequence of interactions with the fiber surface and is also named after the sequence. For example, the R mode describes light reflected off the fiber surface, the TT mode describes light transmitted into the fiber in and then immediately out, and the TRT mode describes light transmitted into the fiber and reflected internally once before being transmitted out.

In agreement with previous BSDF models for cloth [Irawan and Marschner 2012; Sadeghi et al. 2013], informal measurements of individual fibers suggest an azimuthally uniform R mode and a TT mode with a single forward scattering lobe. Moreover, the more detailed TRT mode that appears in hair fibers does not appear to be important for textile fibers. (See the supplementary material for evidence.) This motivated us to include only the first two modes in our model:

$$S(\omega_i, \omega_o) = \frac{M_R(\theta_i, \theta_o)}{2\pi} + M_{TT}(\theta_i, \theta_o) N_{TT}(\theta_i, \phi_i, \phi_o). \quad (3)$$

The two modes are depicted in Figure 3.

Our model has five parameters that determine the intensities and widths of the two modes.

- C_R is the color of the R mode.
- C_{TT} is the color of the TT mode.
- β_R is the longitudinal width of the R mode.

— β_{TT} is the longitudinal width of the TT mode.

— γ_{TT} is the azimuthal width of the TT mode.

We make use of two differently normalized Gaussian-like functions. One is $\bar{g}$, a renormalized Gaussian in θ :

$$\bar{g}(\theta; \mu, \sigma) = \frac{g(\theta; \mu, \sigma)}{G(\mu, \sigma)}, \quad (4)$$

where $g(\theta; \mu, \sigma)$ denotes the Gaussian distribution with mean μ and standard deviation σ , and $G(\mu, \sigma)$ is a normalization factor (defined in Appendix A) chosen to satisfy energy conservation.

The other is the von Mises distribution f (also used in [Irawan and Marschner 2012]), which is the analog of the Gaussian distribution on the circle:

$$f(\phi; \mu, \sigma) = \frac{\exp(\sigma^{-2} \cos(\phi - \mu))}{2\pi I_0(\sigma^{-2})}, \quad (5)$$

where $I_0(x)$ is the modified Bessel function of order 0.

4.2 The R Mode

The R mode accounts for light that reflects specularly from the surface of the fiber. The total amount scattered into this mode depends on θ due to Fresnel reflection; we model this dependence using a heuristic formula, since the surface is not planar and the actual fraction transmitted depends on the cross section and surface properties, neither of which we wish to model. We introduce a parameter C_R that specifies the reflectance at $\theta_i = 0$, then use Schlick's approximation to let reflectance increase to 1 as the incident direction becomes parallel to the fiber:

$$\mathcal{F}_R(\theta_i) = C_R + (1 - C_R)(1 - \cos \theta_i)^5. \quad (6)$$

In the absence of a particular cross section, we assume that the reflected light is distributed uniformly in ϕ , so that

$$N_R(\theta_i, \phi_i, \phi_o) = \frac{1}{2\pi}.$$

Moreover, it is scattered to a small range of θ that increases with surface roughness, which we model using the normalized lobe $\bar{g}$:

$$M_R(\theta_i, \theta_o) = \mathcal{F}_R(\theta_i) \bar{g}(\theta_o; -\theta_i, \beta_R). \quad (7)$$

Because the values of C_R may be different among the red, green, and blue channels, our model allows reflection from fiber surfaces to be colored. Although Fresnel's equations do not predict this, we deliberately make the reflection colored as we found empirically that this led to the model being better at reproducing the color of some fabrics.

4.3 The TT Mode

The TT mode represents light that transmits into the fiber and then out. It is responsible for the remaining fraction $1 - \mathcal{F}_R(\theta_i)$ of incoming light. It is colored (via the parameter C_{TT}) to account for light absorbed by colorants in the interior of the fiber.

As with the R mode, we model the longitudinal spread using a normalized Gaussian, but since transmitted light is generally focused forward, we model the dependence on ϕ using the von Mises distribution centered at $\phi_i + \pi$:

$$M_{TT}(\theta_i, \theta_o) = C_{TT}(1 - \mathcal{F}_R(\theta_i)) \bar{g}(\theta_o; -\theta_i, \beta_{TT}), \quad (8)$$

$$N_{TT}(\theta_i, \phi_i, \phi_o) = f(\phi_o; \phi_i + \pi, \gamma_{TT}), \quad (9)$$

where γ_{TT} controls the azimuthal width of the forward scattering peak.

Discussion. The two components of the model together define a simple but expressive model for scattering from fibers. It can model rougher fibers like cotton or wool compared to smoother fibers like nylon or silk by adjusting β_R and β_{TT} ; it models the color of fibers primarily using C_{TT} , and the effects of different cross sections that produce more or less strongly forward-directed scattering are modeled by adjusting γ_{TT} .

4.4 Volumetric Appearance Model

The scattering model can be adapted into a volumetric appearance model compatible with the anisotropic RTE [Jakob et al. 2010]. To do so, we need to specify a normalized phase function, an albedo, and a directionally varying coefficient of attenuation.

To specify a phase function, consider a single color channel. For each value of ω_o , we can construct a probability distribution over the sphere of incoming directions that is proportional to $S(\omega_i, \omega_o)$ and use this distribution as the (direction-dependent) phase function. Our implementation uses the probability distribution, computed by tabulation of both ω_i and ω_o , for importance sampling of the fiber scattering model as the phase function. However, because C_R and C_{TT} can have different values in different channels, we construct a separate phase function for each channel. When rendering, we render three monochrome images for each channel before combining them to a single colored image.

We leave the albedo $\alpha = \sigma_s/\sigma_t$ as a parameter to be fitted by our fitting process.

For the extinction coefficient, we choose a function such that σ_t is maximal when the light's direction is perpendicular to the local fiber direction and decreases smoothly to 0 when the light's direction is parallel to the local fiber direction. In particular, we choose

$$\sigma_t(\omega) = \sigma_{t,\max} \sqrt{1 - (\omega \cdot \mathbf{d})^2}, \quad (10)$$

where $\mathbf{d}$ is the local fiber direction, and $\sigma_{t,\max}$ is the maximum coefficient of extinction, which is a parameter to the model. In theory, $\sigma_{t,\max}$ is a parameter our fitting process can fit. However, we set it to constant as will be discussed in the next section.

5. APPEARANCE MATCHING

In this section, we explain our appearance matching method. As outlined in Section 3, the method receives as input photographs of a material and corresponding scenes with the same geometric configuration of camera, cloth sample, and light source. Its goal is to find model parameters that result in renderings that match the photographs. The method quantifies the differences between photographs and renderings with an objective function, and minimizes its value through stochastic gradient descent on the model parameters.

To implement this method, we must obtain derivatives of the objective function with respect to the model parameters, which in turn requires differentiating the rendered images with respect to the parameters. We begin with an introduction to a derivative estimation method from stochastic simulation and its application to path tracing. We then discuss the choice of objective function and the gradient descent optimization.

5.1 Derivative Estimation

In the path integral formulation of physically based rendering, the intensity of a pixel is an integral of a contribution function over path space. Consider a scattering model parameter θ . The intensity

depends on θ as follows:

$$I(\theta) := \int_{\Omega} f_{\theta}(x) dx, \quad (11)$$

where Ω is the space of all paths in the scene, and $f_{\theta}(x)$ is the contribution of the path x .

With some reasonable smoothness conditions, it is possible to compute the derivative of $I(\theta)$ by passing the derivative operator through the integral:

$$I'(\theta) = \frac{d}{d\theta} \int_{\Omega} f_{\theta}(x) dx = \int_{\Omega} \frac{df_{\theta}(x)}{d\theta} dx. \quad (12)$$

This integral, then, can be estimated by Monte Carlo integration:

$$I'(\theta) = \int_{\Omega} \frac{(df_{\theta}/d\theta)(x)}{p(x)} p(x) dx = E_{\bar{X}} \left[\frac{(df_{\theta}/d\theta)(\bar{X})}{p(\bar{X})} \right], \quad (13)$$

where p is *any fixed* probability distribution, and $\bar{X}$ is a path random variable with distribution p .

Now, we want to evaluate the derivative at $\theta = \theta_0$. Since p is an arbitrary fixed probability distribution, we can use importance sampling based on the value θ_0 to pick an efficient probability distribution p_{θ_0} . Then, an unbiased estimator of $I'(\theta_0)$ is given by

$$\frac{(df_{\theta}/d\theta)|_{\theta=\theta_0}(\bar{x})}{p_{\theta_0}(\bar{x})}, \quad (14)$$

where $\bar{x}$ is a path sampled according to p_{θ_0} .

Recall that f consists of a product of terms such as BSDF evaluations, volume transmittances, and emitter intensities. Abstracting the origins of the individual terms², this may be written as

$$f_{\theta}(\bar{x}) = \prod_{i=1}^N f_{i,\theta}(\bar{x}). \quad (15)$$

As a result,

$$\frac{df_{\theta}(\bar{x})}{d\theta} = f_{\theta}(\bar{x}) \sum_{i=1}^N \frac{(df_{i,\theta}/d\theta)(\bar{x})}{f_{i,\theta}(\bar{x})}. \quad (16)$$

Substituting into (14), the expression evaluated by the estimator for $I'(\theta_0)$ is

$$\frac{f_{\theta_0}(\bar{x})}{p_{\theta_0}(\bar{x})} \left[\sum_{i=1}^N \frac{(df_{i,\theta}/d\theta)|_{\theta=\theta_0}(\bar{x})}{f_{i,\theta_0}(\bar{x})} \right]. \quad (17)$$

As a result, instrumenting a path tracer to estimate $I'(\theta_0)$ is relatively straightforward. For each term $f_{i,\theta_0}(\bar{x})$ calculated while tracing a path, the path tracer need only accumulate the associated term $r_{i,\theta_0}(\bar{x})$ into a running sum to obtain from the same path an estimate of I and I' . Moreover, the same path may yield estimates of derivatives of I with respect to an arbitrary number of parameters by accumulating sums of r for each parameter.

In our implementation, how we compute derivatives depends on how different parts of the BCSDf are evaluated. When a part is evaluated directly (e.g., the term for the R mode), we symbolically differentiate the expression for that part and write another piece of

²One may note that we make no assumption regarding path length or termination criteria, such as Russian Roulette. The choice of termination criteria and any terms introduced by the use of Russian Roulette are folded into $p_{\theta_0}(\bar{X})$.

Fig. 4. Renderings and derivatives with respect to fitted parameters of the Fleece model. For two view configurations, we show the rendered Fleece and visualizations of its derivatives with respect to each parameter of the scattering model of Section 4. In the derivative images, gray (shown in the borders) indicates a value of zero, while lighter and darker values in each channel indicate positive and negative values, respectively. For scalar parameters, the image shows how each channel changes with respect to the parameter. For the color parameters, each channel of the parameter affects a single channel of the rendering, so the image visualizes the derivative of each channel of the rendering with respect to the color parameter in that channel. Derivative magnitudes are not to scale across images.

Fig. 5. Measurements used in the fitting process. The icons indicate the orientation of the material (bolded edges), the light source (brown), and the camera (blue); the normal angles to the camera and light are given beneath. Due to limitations of the measurement apparatus, normal angles of 80 degrees in the final row are constrained to 64 degrees.

code to carry out the derived calculation. Other parts, such as the probability distribution in the volumetric model, are tabulated. For these parts, we also tabulate the derivatives with respect to relevant parameters as we tabulate the parts.

Figure 4 shows visualizations of the derivatives of images of the Fleece fabric model under two lighting/viewing configurations.

5.2 Measurements

To derive a model that matches a real piece of fabric, we need measurements of that fabric to match against. For each fabric, we took 16 photographs of a flat sample of the cloth illuminated by a $10\text{cm} \times 10\text{cm}$ square light source located about 61cm from the sample. Between measurements, we move the camera and the light source around hemispheres centered at a point on the fabric. Each photograph is cropped to a square covering roughly $1\text{cm} \times 1\text{cm}$ area of the material around the center point of the camera's orbit. Figure 5 visualizes the 16 measurement configurations we use for appearance matching, with the corresponding cropped photographs of Fleece.

A number of factors influenced our choice of using 16 measurements. First, because there are 12 parameters to fit³, there must be at least 12 observations to have a well-posed problem. Second, we generally would like to use as few measurements for fitting as possible because rendering detailed geometry with full light transport simulation is time consuming. Third, however, we would like enough diversity in lighting and viewing configurations. We settled on 16 measurements, which allows us to include four types of light/camera elevations (the columns of Figure 5), two types of fabric rotations (the odd rows versus the even rows), and whether the light source and the camera are in the same plane or not (Row 1 and 2 versus Row 3 and 4).

5.3 Objective Function

To optimize for parameter values, we need an objective function, which is a scalar valued function that summarizes the difference between the photographs and the corresponding renderings. We denote the objective function by $f(\vec{R}_{\vec{\theta}}, \vec{M})$, where $\vec{\theta}$ is the vector of values of all the rendering parameters, $\vec{R}_{\vec{\theta}}$ is the vector of pixel intensities of the rendered images when the parameters are set to $\vec{\theta}$, and $\vec{M}$ is the same vector of the photographs.

The photographs and the rendered images show views of the material under matching conditions but do not show the same piece of fabric, so the objective function must compare the images without depending on the details being the same. The simplest way to do this is to average the whole image so that differences in the spatial details do not matter. Therefore, we form the measurement vector $\vec{M}$ by concatenating the average intensities of the 16 measured images in each color channel. In effect, we rely on fiber-level microgeometry to ensure a texture that is at least plausible. We optimize only against the average intensities of our measurements, guaranteeing that the base color and highlight of the material, which are visible at near and far scales, are captured accurately.

³The number of parameters comes from the combination of using the volumetric microgeometry model and the phase function derived from the fiber scattering function in Section 4. The parameters are C_R (3), β_R (1), C_{TT} (3), β_{TT} (1), γ_{TT} (1), and α (3), so in total there are $3 + 1 + 3 + 1 + 1 + 3 = 12$ parameters.

We now discuss our choice of the objective function f . Recall that the rendered measurements $\vec{R}_{\vec{\theta}}$ and their derivatives are only available as unbiased estimates from the renderer, so we will not be able to calculate f or its gradient exactly. Moreover, if one desires the theoretical convergence properties of stochastic gradient descent, we must calculate an unbiased estimate of f and its gradient. Thus, f must interchange with expectations as follows:

$$\mathbb{E}[f(\vec{R}_{\vec{\theta}}, \vec{M})] = f(\mathbb{E}[\vec{R}_{\vec{\theta}}], \vec{M}). \quad (18)$$

We observe that all multivariate polynomials in the components of $\vec{R}_{\vec{\theta}}$ satisfy this property, as long as any rendered values $\vec{R}_{\vec{\theta}}(i)$ that occur in the same product are uncorrelated, for example, by being calculated in separate renderings.

Our objective function takes the form of a weighted sum of terms f_{image} calculated per image:

$$f(\vec{R}, \vec{M}) = \sum_{i=1}^N w_i f_{\text{image}}(\vec{R}(i), \vec{M}(i)), \quad (19)$$

where $\vec{R}(i)$ and $\vec{M}(i)$ contain the average intensities in each channel of the render and photograph of configuration i , respectively. The per-image terms are

$$w_i = \frac{1}{\max(a_i, \tau)^2} \quad f_{\text{image}}(\vec{R}(i), \vec{M}(i)) = \sum_{c \in \{r, g, b\}} (\vec{R}(i)_c - \vec{M}(i)_c)^2. \quad (20)$$

The weight w_i is chosen such that each term of the sum in (19) approximates the square of the relative error between f_{image} and $\vec{M}(i)$. For simplicity, we use the average intensity $a_i = (\vec{M}(i)_r + \vec{M}(i)_g + \vec{M}(i)_b)/3$ of photograph i instead of assigning different weights to different channels. We calculate the *relative* error of each measurement to avoid overfitting the intensities of particularly bright images such as when the camera views the specular highlight produced by the light source at grazing angles.

Additionally, we threshold the weights of dark photographs to not be lower than $1/\tau$, where τ is set to 0.02 for all the fittings. Introducing the threshold prevents the optimization from overfitting the intensities of very dark images, where even differences due to the dark current noise produced by the camera may result in a large relative error in intensity.

As mentioned earlier, we want to evaluate an unbiased estimate of f , given the unbiased estimates of average intensities obtained from the renderer. Because f_{image} contains squares of the per-channel average intensities $R_{i,c}$ of the rendered images, each image must be rendered twice independently to make them uncorrelated. In this way, we may compute the squared difference without introducing bias.

5.4 Stochastic Gradient Descent Optimization

We now wish to explore the space of rendering parameter configurations $\vec{\theta}$ to minimize the value of f . For convenience, consider f as a direct function of the parameters $\vec{\theta}$. In stochastic gradient descent, we start with an initial parameter value $\theta^{(0)}$. We then iteratively modify the parameter values in the opposite direction of the (estimated) gradient:

$$\vec{\theta}^{(i+1)} = \vec{\theta}^{(i)} - \alpha_i \nabla f(\vec{\theta}^{(i)}), \quad (21)$$

where α_i is a scaling factor, often called the *learning rate*, that changes as the iteration proceeds. Despite the noise in the gradient

estimate, it can be shown that the iteration converges to the global optimum given that f is convex over the search space and that the learning rates decrease at the appropriate speed: namely, if $\sum_{i=1}^{\infty} \alpha_i^2 < \infty$ and $\sum_{i=1}^{\infty} \alpha_i = \infty$ [Bottou 2010].

In our implementation, we choose the harmonic series $\alpha_i = a/i$ as the learning rates, where a is a constant. While the choice of a has no effect on the convergence guarantee, choosing a value that makes the optimization converge quickly is critical in practice. To make finding one possible, we nonlinearly scale the parameter space using the process described in Section 5.5.

We perform the gradient descent as follows. Initializing $\vec{\theta}$ to a starting value $\vec{\Theta}$ (defined in the next section), we run six different stochastic gradient descent iterations for exponentially bracketed choices of the learning rate multiplier a , running each for 60 iterations. To slow the $1/i$ decay of the step size while preserving the convergence guarantees of SGD, we initialize the iteration number i to 50. We select the optimized parameters with the lowest reported objective function value $f(\vec{R}_{\vec{\theta}}, \vec{M})$ as our result. In some instances, the gradient descent yielding the least residuals continued to oscillate intensely around a local minimum of the objective function after 60 iterations; in such cases, we performed a final 10 iterations at one tenth the last learning rate to descend to the local minimum.

5.5 Parameter Rescaling and Ranges

A suitable learning rate is hard to find when regions in parameter space have large differences in their gradient magnitudes. Such a situation complicates the selection of the learning rate because of the following.

- The process may enter a neighborhood where gradients have small magnitude. If the learning rate is too small, it will spend a long time moving in short increments through this region.
- The region surrounding a local optimum may have gradients with large magnitude. If the learning rate is too large, the optimization process will repeatedly overshoot the local minimum many times before converging.

For example, a rendering of cloth will change much more dramatically with respect to the TT mode color, C_{TT} , when the parameter value is large than when the value is small. As such, a large learning rate is needed when C_{TT} is small, and a small learning rate is needed when C_{TT} is large. A situation might arise where no single learning rate works well on all regions.

We mitigate the effect of disparity in gradient magnitudes by automatically defining a mapping r from the rendering parameters $\vec{\theta}$ to a space of search parameters $r(\vec{\theta})$. The stochastic gradient descent is performed in this rescaled space instead of the space of rendering parameters. Our goal is to find a mapping such that a unit change in the remapped space $r(\vec{\theta})$ corresponds to a constant change in the objective function value. To approximate this, we choose a starting configuration $\vec{\Theta}$ of the rendering parameters and rescale each parameter individually based on its effect on the objective function near this configuration.

For each parameter θ_p , we define a mapping r_p for the parameter as follows. We fix all other parameters to the values specified by $\vec{\Theta}$ and calculate the average intensities $\vec{R}$ obtained by setting parameter θ_p to an ascending sequence of values $c_1, \dots, c_k$ spanning the range we will permit the optimization to explore for this parameter. We then define r_p at the values c_j as shown next, and extend it to a piecewise linear function on the domain $[c_1, c_k]$:

Fig. 6. Calculated rescaling curves for the C_{TT} and γ_{TT} parameters of the scattering model of Section 4 for Fleece. The curves relate the parameters $r(\theta)$ optimized by the gradient descent to the rendering parameters θ . The starting parameter configuration $\bar{\Theta}$ is given in Table II.

$$r_p(c_1) = 0$$

$$r_p(c_{j+1}) - r_p(c_j) = \sum_{i=0}^N w_i \sqrt{f_{\text{image}}(\bar{R}_{\bar{\Theta}, \bar{\theta}_i=c_j(i)}, \bar{R}_{\bar{\Theta}, \bar{\theta}_i=c_{j+1}(i)})}. \quad (22)$$

As $\bar{\theta}_p$ increases from c_j to c_{j+1} , the increase in the function $r_p(\bar{\theta}_p)$ is equal to the change in the averages $\bar{R}$, though we actually compute the square root to counter the nonlinearity of f_{image} . Consequently, regions of the domain of permitted values of $\bar{\theta}_p$ that correspond to a large change in $\bar{R}$ are mapped to larger regions in the space of search parameter values $r_p(\bar{\theta}_p)$. Figure 6 visualizes the calculated rescaling curves r_p for parameters of the scattering model of Section 4 that have large disparity between derivative magnitudes in different regions.

The rescaling function r for all the parameters, a vector function, is formed by assembling the per-parameter functions r_p . However, we treat scalar parameters (β_R , γ_R , and γ_{TT}) differently from color parameters (C_R and C_{TT}). For each scalar parameter p , we compute r_p as detailed previously and use the same function in the assembly of r . On the other hand, we treat a color parameter p as three separate scalar parameters that always have the same value when computing the rescaling curve r_p . Consequently, changing the (scalar) value of p is equivalent to changing the value of three scalar parameters at the same time. Thus, when assembling r , we set the per-channel function r_p^{red} , r_p^{green} , and r_p^{blue} to $r_p/3$ instead of r_p .

Compared to not rescaling the parameters, our approach handles the images' different sensitivity to unit change of different parameters better. For example, we found that the rendered images are much more sensitive to changes in C_{TT} than in γ_{TT} , so at least a component-wise linear rescaling was essential. However, adding nonlinearity was also very important. Before employing the rescaling method we proposed, we found it difficult to fit to a material that was relatively bright in at least one channel in grazing configurations. As a specific example, the rendered images were so sensitive to C_{TT} in the grazing configurations that unless the initial condition was very near to the correct value, the gradients were large enough to push the C_{TT} value to both extremes of the domain unless the step size was very small.

Table II lists the ranges of values we permit each parameter to take during the optimization. The varying ranges permitted for the lobe widths are due to practical considerations. We observed that

Table II. Fitting Domains of Rendering Parameters

Parameter	Lower Bound	Upper Bound	Θ
C_R	0.001	0.999	0.1
C_{TT}	0.001	0.999	0.85
β_R	1.0°	10.0°	5.0°
β_{TT}	10.0°	45.0°	10.0°
γ_{TT}	1.0°	45.0°	10.0°
α	0.001	0.999	0.85
γ	0.005	1.5	N/A
d	N/A	N/A	4000.0

The final column lists the default parameter values comprising the $\bar{\Theta}$ of the rescaling process described in Section 5.5. α and d denote the albedo and density multiplier when rendering volume geometry (the latter is held fixed). Parameter fitting for the microflake model by the method of [Zhao et al. 2011] use the stated ranges for α and γ . For the black Velvet, a lower bound of 0.04 was used for C_R .

Fig. 7. Fitted Fleece results for volumes of different density scales d . When the camera and light are both above the plane of the material, reducing the density from (a) 4,000 to (b) 2,000 to (c) 1,000 has minimal effect on the ability of the optimization to recover the same appearance. In draped configurations, though, a scale of 1,000 leads to an unrealistically thin appearance.

values of β_R greater than 10 degrees led to an implausibly diffuse fiber-level appearance without significantly affecting the average intensities fitted against by the optimization. Moreover, values of β_{TT} less than 10 degrees led to instability in the fitting process, as they greatly increased the sensitivity of the rendered images to the C_{TT} parameter in grazing views with intense highlights. For the black Velvet, we constrained C_R to be at least 0.04.

The rightmost column of Table II lists the default parameter values used when rendering images to define the rescaling functions r_p . We chose values we believed would be typical of an “average” material so that the rescaling is defined relative to a region of the parameter space we expect to be most heavily explored by the optimization. For each parameter, we choose a nine-value sequence $c_1, \dots, c_9$ for use in defining the function r , assigning values c_j more finely towards the upper bound of C_{TT} and towards the lower bounds of C_R , β_R , β_{TT} , and γ_{TT} , capturing the greater sensitivity of the rendered images to values of the parameters in those regions. More details on these values can be found in the supplementary material.

For volume geometry, we include the single-scattering albedo α in the fitting process but leave the density multiplier d fixed. As shown in Figure 7, we observed that when fitting against the photographs, all of which place both the camera and light source

Fig. 8. Fabric geometry creation pipeline.

above the plane of the material, the value of d did not significantly affect the ability of the optimization to match the photographed appearance. Under draped configurations, though, fitted parameters with a low d resulted in an unacceptably “thin” appearance.

6. FABRIC GEOMETRY CONSTRUCTION

In this section, we show how to construct fabric microgeometry. We construct two representations: a surface-based one and a volumetric one. The surface-based representation is a collection of discrete surfaces modeling the surfaces of individual fibers, while the volumetric representation stores density and fiber direction in a high resolution voxel grid.

Both representations are created from micro CT scans of cloth samples according to the pipeline in Figure 8. We use Zhao et al.’s [2011] image processing pipeline to compute a preliminary volumetric representation, a voxel array containing (1) the density of the material in each voxel, and (2) the local direction of the fiber at that voxel. From this volume, we infer locations of textile fibers and then construct cylindrical surfaces to cover them. This gives us the surface-based fiber representation. Lastly, we use this fiber model to improve the consistency of the preliminary fiber direction volume, resulting in the final volumetric representation.

We first focus on the process that converts the volumetric representation to fibers. The input to this process is a voxel array with density and direction, and the output is fiber geometry in the form of a collection of 3D polylines. Each segment of polyline acts as the axis of symmetry of a cylinder with circular cross-section, and consecutive cylinders sharing a vertex are joined with a miter joint, producing a continuous 3D surface for rendering.

We assume that the input volume is roughly axis-aligned, that is, the weft fibers are roughly along the x -axis, the warp fibers along the y -axis, and in case of fabrics such as velvet, vertical fibers along the z -axis.

To produce fiber geometry, the following steps are taken.

- (1) *Volume decomposition.* The input density volume is decomposed into three volumes corresponding to the x -, y -, and z -axis. In this way, the warp, weft, and vertical fibers can be processed separately along their length.

Fig. 9. The density volume obtained from a micro CT scan of Velvet and its three decomposed volumes.

- (2) *Fiber center detection.* The 2D slices of each volume along the corresponding axis are processed independently to detect cross sectional centers for fibers going through the slices.
- (3) *Polyline creation and smoothing.* Nearby fiber centers in adjacent slices are linked together to form chained polylines in 3D. These polylines act as skeletons of the reconstructed fibers. Short polylines are removed, and the union of all polylines from the three volumes become the polylines for the entire volume. These polylines are then smoothed.
- (4) *Radius determination.* The single radius of all the cylinders constituting the fiber surfaces is then determined.

We now discuss each of these steps in more detail.

6.1 Volume Decomposition

The first step is to separate the model into three subvolumes containing only the fibers oriented (approximately) along the warp, the weft, and the vertical directions. The three subvolumes are created simply by associating each voxel with the coordinate axis closest to its direction vector (determined by selecting the component of the direction vector that is largest in absolute value). The three subvolumes are later processed independently to recover fibers that run primarily in each volume’s associated direction.

Figure 9 depicts the result of decomposing a micro CT scan volume of Velvet, which has *pile* fibers sticking up perpendicular to the overall plane of the fabric. We see that the x - and y -dominant volumes primarily contain voxels in the woven part (one for the warp, the other for the weft), and the z -dominant volume mostly contains the pile fibers. Noise voxels can be observed in all the volumes and can give rise to short extraneous fibers, which are removed later in the pipeline.

6.2 Fiber Center Detection

Each subvolume is processed in 2D slices perpendicular to its dominant direction, with the goal of locating where each fiber crosses each slice. Since fibers are roughly perpendicular to the slices, they appear as compact blobs of higher density. To determine the centers of these blobs, we apply the standard blob detector which convolves the slice image with Laplacian-of-Gaussian filters at several scales, and then finds the local minima of the response in both scale and spatial domain [Lindeberg 1998]. Figure 10 shows fiber centers detected in a slice of the z -dominant axis of the Velvet volume in Figure 9. The result is a collection of fiber locations $\{(i_1, j_1), (i_2, j_2), \dots, (i_m, j_m)\}$ for each slice, which must be matched up across slices to produce 3D fibers.

6.3 Fiber Building

In this step, we connect the detected fiber centers to create polylines representing the individual fibers. Contrary to previous approaches to fiber detection that grow fibers one after another [Jakob et al.

Fig. 10. (a) A slice of the z -dominant subvolume of Velvet in Figure 9 in 3D view, (b) the slice viewed as a 2D image, and (c) the result of the convolution of the Laplacian-of-Gaussian filter to the slice image with the detected blob positions depicted as small red squares.

2009; Luo et al. 2013], we view this as a matching task: we first decide which detected centers in all pairs of neighboring slices belong to the same fiber, then extract maximal paths in the resulting graph to determine the polylines to generate.

We connect fibers across slices by solving a series of bipartite graph matching problems, each matching the fiber centers detected in slice k to the centers in slice $k + 1$. These fiber centers become vertices in a weighted bipartite graph. Edges are constructed between vertices on different slices with weights inversely proportional to the in-plane distance between the end points. In particular, between point a in slice k and point b in slice $k + 1$, there is an edge with weight

$$w(a, b) = \exp(-d(a, b)^2 / (2\sigma^2)) \quad (23)$$

$$d(a, b) = \sqrt{(i_a - i_b')^2 + (j_a - j_b')^2}. \quad (24)$$

or zero if $d(a, b) > \sigma$. We used $\sigma = 5$ voxels for all the fabric volumes we processed in this article.

The maximum weighted bipartite matching can be solved to an approximation ratio of $1/2$ using an $O(n^2 \log n)$ greedy algorithm [Drake and Hougardy 2003], which is much faster at our problem size than the optimal but $O(n^3)$ Hungarian algorithm. We found that the greedy algorithm worked well with our data.

Due to noise, this process generates many short polylines in addition to the long polylines corresponding to well-tracked fibers. We retain only polylines with at least 10 vertices when we collect the results from the x , y , and z volumes together. Figure 11 shows the result of this step on the Velvet volume.

6.4 Polyline Smoothing

Because fiber centers are located independently per slice, the raw polylines are noisy (see Figure 11). To reduce noise, we smooth them, as described in Luo et al. [2013]. New vertex positions $p_1, p_2, \dots, p_n$ are computed to minimize the energy:

$$\mathcal{E} = \sum_i \alpha \|p_i - p_i^{(0)}\|^2 + \|p_{i-1} - 2p_i + p_{i+1}\|^2, \quad (25)$$

where $p_i^{(0)}$ is the original 3D position of the i th vertex, and $\alpha = 0.1$.

6.5 Radius Determination

Given the collection of polylines previously determined, we now need to compute a value for the fiber radius to fully define the set of 3D cylinders that represent the fibers. We choose the radius to match the volume covered by the fibers to the voxels that is occupied by fibers in the original volume.

Fig. 11. Polyline generated by the fiber growing process and the effect of smoothing on some generated polylines.

More concretely, we first upsample the original volume by a factor of 4. Then, given a candidate value r of the radius, we rasterize cylinders of radius r around all polylines into a volume of the same resolution. The score for the value r is given by

$$\text{score}(r) = |C(r) \cap V| - |C(r) \otimes V|, \quad (26)$$

where,

— $C(r)$ is the set of voxels contained in one of the cylinders when the radius is set to value r ,

— V is the set of non-empty voxels in the original volume, and

— $\otimes$ is the symmetric difference operator:

$$A \otimes B = (A \cup B) - (A \cap B).$$

The radius values are found by trying out 20 evenly spaced values between w to $4w$, where w is the width of a voxel in the original volume.

6.6 Results and Discussions

Figure 12 compares renderings of the original micro CT scan volumes of six fabrics with the constructed fiber geometry. In general, the fibers agree with the geometry of the volume but look a bit thinner because (1) the pipeline filtered out many short fibers, and (2) the radii were determined in part by trying to minimize the number of voxels covered by the cylinders but not the original volumes.

We note that our pipeline can break highly curved fibers into disconnected pieces when different parts of the fiber belong to

Fig. 12. Visualizations of original micro CT scan volumes of six pieces of fabric and their reconstructed fibers.

different decomposed volumes. We do not attempt to reconnect the fibers, since, as we shall see in the final rendering of the fabric model, some broken fibers are acceptable at typical viewing distances where fibers are not clearly resolved.

Our pipeline requires a scan resolution that resolves individual fibers well. It also requires that the fiber cross sections appear circular or elliptical at the scanned resolution, which might not be true for synthetic fibers such as rayon, nylon, and acetate.

6.7 Improving Volume Direction Fields

In real cloth, fine-scale irregularities in the fibers have a major effect on luster: smooth, well-aligned fibers (as in many silks and synthetics) reflect coherently and produce bright highlights, whereas kinky and irregularly arranged fibers (cotton, wool) reflect to a range of directions, producing a less shiny appearance. In fiber or volume models, noise in the fiber curves or the direction volumes can introduce a similar change in appearance—noise makes it impossible to match the appearance of smooth fibers as seen in silk or velvet.

Smoothing of polylines is very effective in removing this noise without disturbing fabric structure or over-smoothing less organized fabrics. But when smoothing directions in a volume, naïve approaches can easily smooth out important features by mixing the directions belonging to different fibers. Rather than pursuing more complex noise-reduction methods for the volume direction fields and to keep the comparison between fibers and volumes on an equal footing, we smoothed the volume direction fields by simply transferring the directions from the fiber models onto the volume model. This is implemented by setting the direction of any voxel contained in a fiber cylinder to a unit vector parallel to the fiber. To increase coverage, we use cylinders of radius $2r$ instead of r , where r is the radius determined in Section 6.5.

The result of this process can be observed in the second column of Figure 13. We can see that the direction field becomes much less noisy, and the highlight in the rendering becomes much more prominent relative to the base color of the material.

In summary, we have shown how to compute a surface-based fiber representation from the micro CT data, and how to transfer the smoothed directions from this representation to create a volumetric representation. We will now compare these two representations to evaluate their ability to represent real fabrics accurately.

Fig. 13. The effect of fiber direction consistency on appearance. The first two rows depict the local fiber direction at the first interaction between the eye ray and the volume, visualizing the absolute values of the x , y , and z components of direction as an RGB color. The volume in the second row is obtained by tiling the swatch from the first row and then warping the tiled volume with a shell map, as done in Zhao et al. [2011].

Table III. Cloth Samples, Volumetric Model Parameters, and Fiber Radii

Name	Material	Weave	Color	Voxel size	Data size	s	t	h	ε_d	ε_J	Fiber radius
Fleece	N/A	unwoven*	blue	$6.60 \mu\text{m}$	$500 \times 540 \times 586$	2.00	3.00	11	0.23	-14	$8.68 \mu\text{m}$
Gabardine	wool	twill	red	$5.00 \mu\text{m}$	$671 \times 457 \times 233$	1.00	2.00	16	0.45	-10	$16.05 \mu\text{m}$
Silk	silk	satin	red	$1.40 \mu\text{m}$	$780 \times 530 \times 160$	1.00	2.00	11	0.55	-30	$4.72 \mu\text{m}$
Velvet	N/A	unwoven	black	$3.03 \mu\text{m}$	$430 \times 478 \times 524$	0.25	0.75	15	0.58	-155	$9.25 \mu\text{m}$
Twill	N/A	twill	green	$2.51 \mu\text{m}$	$570 \times 715 \times 165$	1.00	2.00	11	0.46	-16	$4.88/7.27 \mu\text{m}$
Cotton	cotton	gauze	white	$5.50 \mu\text{m}$	$461 \times 440 \times 160$	0.25	0.75	13	0.41	-5	$8.97 \mu\text{m}$

Note: The data sizes were computed after the rotation and cropping were performed. The s , t , h , ε_d and ε_J are the parameters for the image processing pipeline described in [Zhao et al. 2011]. Twill fabric has two fiber radii because the warp and the weft fibers are of different sizes. *Fleece is a knit fabric processed so that the visible surface is similar to felt.

7. RESULTS

In this section, we present the results of our appearance matching pipeline and compare various model representations for a range of fabrics. We first detail our use of the geometry processing pipeline. Then, we present the matching results and evaluate our appearance-matched models against the photographed materials in different configurations to validate our approach.

7.1 Data Acquisition and Processing

We processed six cloth samples whose details are given in Table III. All samples, except for Gabardine, were scanned with the XRadia VERSA XRM-500 scanner at the Cornell Imaging Multiscale CT Facility. Gabardine's scan was made at the High-Resolution X-ray Computed Tomography Facility at The University of Texas at Austin. Each volumetric scan, except Fleece, was rotated so that the fibers are aligned with the x -, y -, and z -axis according to the requirement in Section 4. The volumes were then cropped so that they can be easily tiled.

Some volumes and their associated fibers received the following special processing.

- The Silk volume was sheared to align the warp and weft yarns with the x - and y -axes, respectively.
- The warp and weft yarns of Twill are composed of different types of fibers with different radii. To capture this difference, we independently ran the radius determination procedure for fibers generated from the x -dominant and y -dominant subvolumes.
- To reduce seams when tiling fabrics with evident regular structure (i.e., Silk, Gabardine, and Twill), we created Wang tiles [Cohen et al. 2003] of the volumes based on their top-down views and used the Wang tiles as source volumes to generate fibers instead of the original volumes.
- Some x - and y -slices of the Cotton volume were removed to ensure that the spacing between the warp and weft yarns roughly match the photographs taken for parameter fitting. This is done so that the area of the rendered image that is covered by a fiber is roughly the same as that of the photographs. Since the weave pattern of the material is quite loose, the inter-yarn spacing may change locally depending on the handling of the material.

The photographs of cloth samples were taken using the Cornell spherical gantry equipped with a Canon EOS 50D camera and a $10\text{cm} \times 10\text{cm}$ LED area light source. The fabrics were mounted on a turntable whose top is a black metal plate. As a result, the photographs did not capture the appearance of the fabrics alone but the appearance of the fabrics with the metal plate underneath them. We simulate this condition by inserting a black mesh underneath all the fabrics being rendered.

Fig. 14. Insufficient brightness of microflake phase function at grazing angles. The camera and light source are at a normal angle of 80 degrees and in a plane with the imaged point on the Fleece volume. Bracketing the density scale multiplier d and microflake standard deviation γ , no configuration is able to match the intensity of the photograph. The albedo is left fixed at 0.999 as the image brightens monotonically with albedo.

For each of the fabric samples, we created three complete fabric models.

- The first uses the fiber mesh geometry we reconstruct (Section 6) and our fiber scattering model (Section 4). This model allows us to explore the capability of the fiber geometry/fiber appearance approach. We refer to this combination as *Fiber/BCSDF*.
- The second uses the volumetric geometry representation with improved direction field (Section 6.7) and the phase function derived from the fiber scattering function (Section 4.4). This model allows us to assess the volumetric geometry/fiber appearance approach. We refer to this combination as *Volume/BCSDF*.
- The third uses the same volumetric geometry representation as the second approach with the microflake-based phase function as specified in Zhao et al. [2011]. This model allows us to compare the previous two new approaches against Zhao et al.'s work. We refer to this combination as *Volume/Microflake*.

We compute the scattering parameters for the volume/microflake model by adapting the fitting procedure from Zhao et al.'s work to the measurements used for the other models we evaluate. Zhao et al. photograph a curved sample of the material, isolate a region of the image with varied appearance, and assign microflake model parameters to match the mean and variance of the pixels in the region. To approximate this, we concatenate the images used by our appearance matching method and treat the concatenated image as the region to optimize. As shown in Figure 14, the microflake model is incapable of reproducing the extreme highlights seen in two of the fitting images, so we omit these when fitting the model.

Fig. 15. Renderings of Velvet in a draped configuration. Our appearance matching process yielded parameters that resulted in different highlight brightness between the images. However, the highlight consistently emerges as the fabric turns away from the camera in all models.

We performed all appearance matching using PCs with 2.27GHz Intel XEON CPUs and at least 64GB of RAM. A scene configuration is rendered into a 64×64 image at 64 samples per pixel. An iteration with 16 configurations, each rendered two times, took about 14 minutes to complete on a single core. We performed 60 iterations, and the whole process took about 14 core-hours for a fabric sample modeled with either the fiber/BCSDF or volume/BCSDF models. The volume/microflake fitting took 2 core-hours because the microflake model was faster to compute and 27 binary search iterations were performed.

7.2 Validation

Validation in planar configurations. Figure 18 in the Appendix shows the photographs used by the appearance matching process for each material alongside renderings produced by the three models. Figures 19 and 20 in the appendix show the scatter plots of the average pixel values of the photographs versus the renderings. The supplementary material contains 492 more validation configurations that were not used for fitting. The corresponding fitted parameter values are given in Table IV in the Appendix.

Validation in nonplanar configurations. To evaluate the qualitative appearance in more natural configurations, Figure 21 in the Appendix shows the fabrics wrapped around a cylinder of radius 1.5cm and the renderings of the three models in the same configuration. Figures 22 and 23 in the Appendix plot the average pixel values of each column of the images in Figure 21, allowing us to quantitatively compare the methods.⁴ Figures 24 and 25 in the Appendix shows the renderings of the fabrics in a simple draped configuration. In Figure 15, we also show Velvet in a more elaborate draped configuration which better reveals the fabric’s characteristic highlights.

7.3 Discussion

We now discuss the results of each fabric in turn before drawing conclusions.

Fleece. The appearance of Fleece is matched well by the fiber/BCSDF and volume/BCSDF models. The volume/microflake

model also performs well in nongrazing configurations, but it is unable to reproduce the brightness in grazing configurations (i.e., the first two rows of the first column in Figure 18 and the $\phi_i - \phi_o = 126^\circ$ and 162° in Figures 21 and 22). The second row of Figure 24 shows that this discrepancy has a large impact on appearance particularly when the light source and the camera are on the opposite side of the fabric. Because the microflake model has no transmission component, it produces an opaque appearance, so the microflake Fleece looks less soft than our models.

Gabardine. In Figure 18 and Figure 21, our two models are able to match the appearance in all configurations well, while the fiber/BCSDF performs somewhat better in the $(80^\circ, 80^\circ)$ grazing configurations in Figure 18. On the other hand, the volume/microflake still cannot produce bright highlights in these grazing configurations. Figures 21 and 22 also exhibit the same trends. All models are close to the photographs when the fabric is in retroreflection configurations ($\theta_i - \theta_o$ low). The volume/microflake model, however, becomes too dark in grazing configurations ($\theta_i - \theta_o$ high), but the other two models become somewhat too bright. This indicates that the new models are improvements over the previous approach for Gabardine.

Silk. Our model still performs better than the volume/microflake at grazing configurations. However, in Figure 18, the grazing highlights rendered with our models have a slight cyan tint. Additionally, in Figure 22, the BCSDF models performs less well in the green and blue channel in configurations with low $\theta_i - \theta_o$ in contrast to their good performance in the red channel.

These behaviors are because the red channel of C_R in the two BCSDF models is much higher than the blue and green channels. The parameters cause light reflected off fibers to be very bright in the red channel and dim in the others. They also cause the light transmitted through fibers to have a cyan tint. We surmise that the fitting process arrived at these values because it needs to set the red channel of C_R high to match the chromaticity of the dim training configurations, which are weighted higher by the objective function: we can see in Figure 19 that the two BCSDF models match the red channel of the dim configurations better than brighter ones. As such, the behaviors are likely to be caused by the models’ additional expressiveness and our choice of objective function.

However, in Figures 21 and 24, the three models look very similar in retroreflection configurations, and the cyan tint of the highlight is very hard to notice. Thus, the suboptimal behaviors of the BCSDF

⁴The photographs and the corresponding renderings are not aligned horizontally. As a result, we shift the photograph in the x-axis so that the peak in the $\phi_i - \phi_o = 162^\circ$ configuration aligns with the peaks of the renderings.

models do not have a significant negative impact, but the inability of the volume/microflake model to become bright in forward scattering configurations takes away a large part of the overall appearance.

Velvet. Velvet was a challenging material to fit because its fitting residuals were higher than for other materials. As a result, the optimization with the settings used for the other materials did not consistently recover the highlights and speckled appearance that characterize it. These effects are more likely to arise from the R mode, given the dark base color of the material, so we constrained C_R to be at least 0.04 for Velvet, to force the high-residual fit towards a region more likely to produce acceptable results.

In Figure 18, the fiber/BCSDF model matches the training data better than the two volumetric models, which are too bright in many configurations. Nevertheless, Figure 20 suggests that all models fit poorly to validation examples. This indicates that more training data might lead to parameters that generalize better to the observed data.

In all draped configurations (Figures 15, 21, and 25), the largest difference between the models is the brightness of the highlights. Note though that all the highlights appear at the same locations, where the fabric turns away from the camera. Moreover, the 10° , 126° , and 162° columns of Figure 23 are evidence that all models can predict the locations of bright highlights under cylindrical configurations despite magnitudes being off. This suggests that all three models, together with a consistent direction field, can model Velvet's appearance, but the appearance matching process needs to be improved to obtain better parameter values.

Still, we note that Velvet is a tricky material because its pile fibers move when touched, and we did not control their directions both when the micro CT scanning was performed and when the photographs were taken. While using more training examples can improve results post hoc, controlling the microgeometry might be required to get truly good results.

Twill. According to Figure 20, our BCSDF models can quantitatively match the training data better than the volume/microflake. However, there are three problems with the result. The first is that all models generalize rather poorly to the validation configurations.

The second is that our models are worse at reproducing the fabric's texture than the volume/microflake. The photographs feature alternating bright yellow and dark brown stripes, but our models do not yield as much color contrast. On the other hand, the volume/microflake produces more contrast between stripes but does not capture the BRDF well in Figure 18.

Low texture contrast was caused by the fibers' being made transparent, as can be seen by low C_R values and high C_{TT} values. However, the fibers still reflect back some light in reality. Indeed, setting C_R higher makes the texture more prominent but worsens the overall color matches (see Figure 16). A cost function that considers contrast between different image parts might be able to trade BRDF correctness with texture correctness and is left for future work.

The third is that there is a feature in photographs of the cylindrical configurations that none of the models could capture. The photographs get brighter around the cylinder's edges in the 10° , 50° , and 90° configurations while the models get dimmer.

The fact that our models can fit the training data well but not the validation data might be explained by overfitting. However, the fact that we have to trade BRDF correctness with texture contrast and that none of the models can account for the bright cylinder's edges suggest rather that the models are not expressive enough to represent all aspects of Twill's appearance. The reason may be that

Fig. 16. An alternative set of scattering parameters of the fiber/BCSDF model for Twill. These parameters were derived manually from the fitted parameters with C_R changed to (0.472, 0.460, 0.218) and C_{TT} changed to (0.540, 0.435, 0.310). Under configurations used for parameter fitting, we compare (a) the photographs with (b) the alternative parameters and (c) the parameters fitted by our process. Additionally, we provide a photo and renderings of the fabric in the cylindrical configuration where $\phi_o = \phi_i = 10^\circ$. For completeness, we also provide the volume/microflake model (d) in the cylindrical configuration. While the alternative parameters and the microflake are inferior at the BRDF level, they produce more contrast in the texture.

the warp and weft yarns, while appearing to be dyed the same color, might actually have different optical properties.

Cotton. Photographs in Figure 21 show that Cotton is a rather diffuse material, and all three models can capture this appearance as well as provide realistic geometric details. Nevertheless, measurements in Figure 18 indicate that the fabric gets extremely bright at the 80° grazing configurations, and none of the models could imitate this behavior. Moreover, they could not match the color of the photographs in any of the remaining configurations. Compared to the other materials, the recovered Cotton fiber geometry and volume are unusually sparse within the threads of the material, suggesting that too many small fibers were thresholded away by the volume processing, limiting the ability of the models to reflect sufficient light back to the viewer.

Highlights. While we matched against a sparse sample of the fabric's BRDF under flat configurations, Figures 22 and 23 suggest that our model produced highlights at the right locations and could imitate the complex highlight shapes such as those of Silk and Velvet in the 10° configurations. Our models always match the magnitude of the highlights better in the grazing configurations than the volume/microflake model. Nonetheless, there are some cases where the latter performed better in nongrazing configurations; for example, Fleece and Silk in the blue and green channels. However, as can be seen in Figure 21, cases where our models perform worse than the volume/microflake do not lead to drastic differences from the reference photographs like the dim grazing highlights of the volume/microflake.

The lobe widths our system produces may not be accurate. In Table IV, many lobe width values are at their extremes, especially the β_R and β_{TT} values of the two BCSDF models. We observed that decreasing β_R and β_{TT} led to brighter images in grazing configurations, so these parameter values tend to get pushed to their lower bounds. We also surmise that the extra degrees of freedom introduced by the albedo α in the Volume/BCSDF model caused more of the model's β_R and β_{TT} values to be at the extremes. As such,

adding more training examples might alleviate this problem. Still, note that, while the individual lobe widths might not be accurate, the overall appearances are good in cases as Fleece, Gabardine, and Silk. Moreover, the problems of Twill and Cotton do not seem to be related to lobe widths at all.

We now draw conclusions from these discussions.

Comparison with the microflake model. For Fleece, Gabardine, Twill, and Silk, our new models—the fiber/BCSDF and the volume/BCSDF—are able to produce results that match the far-field photographed appearance across the images used for fitting. By comparison, the volume/microflake model produces renderings that vary less in intensity from image to image, and fail to recover the intense highlights as the camera and light approach grazing angles. Thus, we conclude that our models are superior to the microflake model. We believe the Fresnel term, which makes the R mode considerably brighter at grazing angles, accounts for much of its advantages.

Fiber versus volume. Except for Velvet, the fiber/BCSDF and the volume/BCSDF models produce very similar results. Nevertheless, in Figure 18, the fiber/BCSDF model generally produces slightly brighter images in grazing configurations. Still, the difference in brightness of grazing highlights is very hard to notice in draped configurations of Figures 21, 24, and 25. We therefore conclude that, in practice, the fiber geometry and volume geometry are similarly good when used with our light scattering model.

Areas for improvement. While we consider Fleece, Gabardine, and Silk to be success cases, the results on the other three fabrics indicate a few potential areas of improvement.

- Number of training examples.* Velvet indicates the small number of training examples may yield parameters that generalize poorly. Few training examples might lead to lobe width values being pushed to their extremes, especially when there are extra degrees of freedoms in the scattering model.
- Texture correctness.* Our objective function only considers correctness at the BRDF level, ignoring texture correctness. As seen in Twill, optimizing according to the function led to parameters that do not reproduce the material’s texture contrast.
- Microgeometry correctness.* We currently have no means to test whether the microgeometry is accurate, and the inaccuracy could have caused problems in Velvet and Cotton.
- Model correctness.* Our models seem to be not expressive enough to capture Twill’s appearance. Moreover, while performing much better than the Volume/microflake, our models still could not precisely reproduce the extremely bright grazing highlights of Silk, Twill, and Cotton.
- Nonphysical behaviors.* The cyan tint in Silk’s grazing highlight, while not a major issue on the overall fitted appearance, shows that our light scattering model can produce nonphysical appearance. This issue arises from our allowing the R mode to have color to achieve better matching results. A model where expressiveness is retained and physics-based behaviors are strictly observed is a possible future work.

7.4 Optical Thickness

As we focused on reflection from fabric rather than transmission, our results do not address the correctness of back-lit appearance. The density scale for the volumes was arbitrarily fixed because it

Fig. 17. Appearance of Gabardine and Twill with and without the underlying black meshes.

had little effect on the appearance from the illuminated side. (This was also done in [Zhao et al. 2011].) As a result, the renderings may appear optically thicker or thinner than the actual fabrics. Also, while our fiber reconstruction algorithm tries to make the generated fibers cover the micro CT volume well, some volume is always lost because the algorithm throws away fibers to remove noise.

The optical thickness has a significant impact on appearance when the fabric is not draped over an opaque object. Figure 7 illustrates the effect of density on the appearance of volume models. Figure 17 shows the change in appearance of the draped fiber models when the underlying black meshes are removed. Notice that the fabrics become much brighter after removal due to multiple scattering in and through their layers.

We emphasize that transmission is not a fundamental limitation but rather a part of the appearance space we have not yet measured. It is entirely within the scope of our optimization and appearance models, but it is currently unknown how well the models will fit under back-lit configurations.

While the fiber and volume microgeometry representations are similar in their abilities to capture fabric reflectance, the optical thickness of the volume geometry can be controlled easily just by setting the density scale. The density of fiber-based geometry can be manipulated by changing the fiber radius, but how to compute derivatives with respect to it is unclear. Therefore it is currently not possible to optimize radius in our system, so the volume geometry is advantageous in this regard. How to control the parameters to achieve the right optical thickness is left for future work.

8. CONCLUSIONS

Reproducing the appearance of fabrics is critical for many applications. While progress has been made on increasingly sophisticated appearance models for fabrics, matching the appearance of real fabrics remains very hard. In this work, we made two contributions to creating fabric renderings that match real fabrics. First, we introduced a novel appearance matching framework based on differentiation and optimization to match rendered images with photographs. Second, we introduced a simple fiber-based scattering

model (BCSDF) and coupled this with new approaches to reconstruct fiber-based geometry and better volumetric models of fabrics. Finally, we matched the appearance of these new models against real photographs and evaluated their strengths and weaknesses. We found that having a fiber-based BCSDF scattering model was critical to matching appearances in grazing configurations. Once we used such a scattering model, both the fiber-based geometry and volume models were approximately similar in quality. Both these approaches proved superior over the prior state-of-the-art volume-based models with microflake scattering.

Additionally, we believe the approach we propose maps a way forward in the field of appearance models based on microgeometry. With our new methods for differentiation and optimization, different models can be systematically tested and compared on an equal basis, providing a clear way to identify deficiencies in existing models and to evaluate a range of possible improvements in order to design the next generation of models for a given material. This general approach can be applied to other problems where it is desirable to test the ability of a model to match measurements, but there is a complicated global illumination process in between the parameters and the data.

The results of our application to cloth appearance also provide crucial knowledge about which models work best, which can be leveraged by future work in this area. In the future, our framework can be extended to handle more cases, for instance, to reason about parameters, such as fiber radius or other geometric parameters, that cause discontinuous changes to path contributions.

ELECTRONIC APPENDIX

The electronic appendix to this article can be accessed in the ACM Digital Library.

APPENDIX

A. NORMALIZATION FACTOR FOR LONGITUDINAL LOBE

In Section 4 the normalization factor G is a function designed to satisfy the following energy conservation constraint:

$$\int_{-\pi/2}^{\pi/2} \bar{g}(\theta; \mu, \sigma) \cos^2 \theta \, d\theta \leq 1.$$

We define $G(\mu, \sigma)$ by approximating $\cos^2 \theta$ from above by a polynomial $Q(\theta)$:

$$\begin{aligned} Q(\theta) &= 0.002439\theta^8 - 0.04301\theta^6 + 0.3322\theta^4 \\ &\quad - 0.999745\theta^2 + 1.0001 \\ &\geq \cos^2 \theta, \end{aligned}$$

and setting:

$$G(\mu, \sigma) = \int_{-\pi/2}^{\pi/2} g(\theta; \mu, \sigma) Q(\theta) \, d\theta.$$

The values of G can be computed by analytically integrating the product of a Gaussian function with a polynomial. The formula for the integral can be found in the supplementary material.

Table IV. Scattering Model Parameters Fitted by Appearance Matching Process for all Materials and Rendering Models

Material	Fiber scattering model, fibers					Fiber scattering model, volume					Microflake, volume		
	C_R	C_{TT}	β_R	β_{TT}	γ_{TT}	C_R	C_{TT}	β_R	β_{TT}	γ_{TT}	α	α	γ
Fleece	0.040, 0.087, 0.087	0.452, 0.725, 0.948	7.238	10.000	25.989	0.032, 0.049, 0.055	0.759, 0.622, 0.999	3.786	10.000	21.865	0.631, 0.840, 0.972	0.137, 0.416, 0.812	0.013
Gabardine	0.185, 0.047, 0.069	0.999 , 0.330, 0.354	2.141	10.000	23.548	0.110, 0.035, 0.048	0.868, 0.633, 0.592	5.034	10.000	23.902	0.993, 0.651, 0.698	0.967, 0.160, 0.230	0.013
Silk	0.745, 0.008, 0.070	0.620, 0.553, 0.562	1.000	10.000	19.823	0.992, 0.001 , 0.034	0.002, 0.690, 0.570	1.000	10.000	13.900	0.940, 0.746, 0.773	0.936, 0.051, 0.214	0.013
Velvet	0.044, 0.040 , 0.040	0.076, 0.058, 0.057	1.577	24.933	44.881	0.969, 0.985, 0.986	0.006, 0.003, 0.003	10.000	10.000	45.000	0.388, 0.310, 0.316	0.490, 0.381, 0.412	0.005
Twill	0.001 , 0.024, 0.987	0.987, 0.975, 0.825	1.367	23.509	26.419	0.001 , 0.001 , 0.016	0.999 , 0.999 , 0.693	1.000	19.759	21.156	0.974, 0.969, 0.878	0.843, 0.812, 0.463	0.013
Cotton	0.989, 0.959, 0.874	0.999 , 0.999 , 0.999	1.000	27.197	38.269	0.447, 0.486, 0.251	0.171, 0.125, 0.279	10.000	10.000	41.464	0.999 , 0.999 , 0.999	0.999 , 0.999 , 0.999	0.013

Note: Parameter values that are at an extreme of their permitted range are bolded. For volumes, the density multiplier d is fixed at 4000.

Fig. 18. Results for all materials and rendering methods on fitting configurations. For all of the viewing configurations used by the fitting process, we show (a) the photographs, (b) renderings produced by the fiber/BCSDF model, (c) renderings produced by the volume/BCSDF model, and (d) renderings produced by the volume/microflake model. The icons indicate the orientation of the material (bolded edges), the light source (brown), and the camera (blue); the normal angles to the light and camera are given beneath. Comparisons across all 492 photographed configurations are in the supplementary material.

Fig. 19. Scatter plots of the average values of the photographs versus those of renderings of Fleece, Gabardine, and Silk. The first row of each material shows the plots for the training configurations, which are the images in Figure 18. The second row shows those for the validation configurations, whose images are available in the supplementary material. Because there are 492 validation configurations, we plot the contours of the points instead.

Fig. 20. Scatter plots of the average values of the photographs versus those of renderings of Twill, Velvet, and Cotton. The setting of these plots are the same as in Figure 19.

Fig. 21. Comparison between photographs and renderings produced by the three fabric rendering methods for the six fabrics. The fabric is wrapped around a cylinder of radius 1.5cm whose axis is vertical and corresponds to the longitudinal angle $\theta = 0^\circ$. The camera was fixed in all images with the light source arranged so that its location spans from a retrorreflection configuration ($\phi_i - \phi_o = 10^\circ$) to being close to the opposite of the camera ($\phi_i - \phi_o = 162^\circ$). The longitudinal angles of both the camera and the light source were set to $\theta = 80^\circ$ in all images, except those with $\phi_i - \phi_o = 162^\circ$, where the light source was lifted to $\theta = 76^\circ$ to avoid being seen by the camera.

Fig. 22. Plots of per-column average pixel values of photographs and renderings in Figure 21 for Fleece, Gabardine, and Silk fabrics. Plots in the same columns are to scale with one another.

Fig. 23. Plots of per-column average pixel values of photographs and renderings in Figure 21 for Twill, Velvet, and Cotton. Plots in the same columns are to scale with one another.

Fig. 24. Renderings of Fleece, Gabardine, and Silk fabric in a simple draped configurations. In all the images, the fabrics are draped over a black mesh to reduce the effect of light passing through the fabrics themselves. In the first row of each fabric, the light is on the same side as the camera while, in the second row, the light is on the opposite side.

Fig. 25. Renderings of Twill, Velvet, and Cotton in the same configurations as in Figure 24.

ACKNOWLEDGMENTS

We thank Frederick von Stein and Mark Riccio for the scanning of cloth samples, Albert Liu for the gencamcal camera calibration system which we used as a part of the system to observe a number of cloth fibers' scattering behavior, and Wenzel Jakob for the Mitsuba renderer which we used to implement our appearance matching system.

REFERENCES

- N. Adabala, N. Magnenat-Thalmann, and G. Fei. 2003. Visualization of woven cloth. In *Proceedings of the 14th Eurographics Workshop on Rendering (EGRW'03)*. Eurographics Association, Aire-la-Ville, Switzerland, 178–185.
- P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. 2000. In vivo fiber tractography using DT-MRI data. *Magn. Reson. Med* 44, 625–632.
- L. Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of the 19th International Conference on Computational Statistics (COMPSTAT'10)*, Y. Lechevallier and G. Saporta, Eds. Springer, Paris, France, 177–187.
- M. Chai, L. Wang, Y. Weng, X. Jin, and K. Zhou. 2013. Dynamic hair manipulation in images and videos. *ACM Trans. Graph.* 32, 4 (July), 75:1–75:8.
- Y. Chen, Y. Xu, B. Guo, and H.-Y. Shum. 2002. Modeling and rendering of realistic feathers. *ACM Trans. Graph.* 21, 3 (July), 630–636.
- M. F. Cohen, J. Shade, S. Hiller, and O. Deussen. 2003. Wang tiles for image and texture generation. *ACM Trans. Graph.* 22, 3 (July), 287–294.
- E. d'Eon, G. Francois, M. Hill, J. Letteri, and J.-M. Aubry. 2011. An energy-conserving hair reflectance model. *Comput. Graph. Forum* 30, 4, 1181–1187.
- D. E. Drake and S. Hougardy. 2003. A simple approximation algorithm for the weighted matching problem. *Inf. Process. Lett.* 85, 211–213.
- I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin. 2013. Inverse volume rendering with material dictionaries. *ACM Trans. Graph.* 32, 6 (Nov.), 162:1–162:13.
- M. Hašan and R. Ramamoorthi. 2013. Interactive albedo editing in path-traced volumetric materials. *ACM Trans. Graph.* 32, 2.
- T. L. Herrera, A. Zinke, and A. Weber. 2012. Lighting hair from the inside: A thermal approach to hair reconstruction. *ACM Trans. Graph.* 31, 6 (Nov.), 146:1–146:9.
- P. Irawan and S. Marschner. 2012. Specular reflection from woven cloth. *ACM Trans. Graph.* 31, 1 (Feb.), 11:1–11:20.
- W. Jakob, A. Arbree, J. T. Moon, K. Bala, and S. Marschner. 2010. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Trans. Graph.* 29, 4 (July), 53:1–53:13.
- W. Jakob, J. T. Moon, and S. Marschner. 2009. Capturing hair assemblies fiber by fiber. *ACM Trans. Graph.* 28, 5 (Dec.), 164:1–164:9.
- J. T. Kajiya and T. L. Kay. 1989. Rendering fur with three dimensional textures. In *Proceedings of SIGGRAPH'89*. 271–280.
- N. Kang and J. Zhang. 2005. White matter fiber tractography via anisotropic diffusion simulation in the human brain. *IEEE Med. Imaging* 24, 9 (Sept.), 1127–37.
- T. Lindeberg. 1998. Feature detection with automatic scale selection. *Int. J. Comput. Vision* 30, 79–116.
- L. Luo, H. Li, and S. Rusinkiewicz. 2013. Structure-aware hair capture. *ACM Trans. Graph.* 32, 4 (July).
- S. R. Marschner, H. W. Jensen, M. Cammarano, S. Worley, and P. Hanrahan. 2003. Light scattering from human hair fibers. In *Proceedings of SIGGRAPH'03*. 780–791.
- S. Paris, H. Briceño, and F. Sillion. 2004. Capture of hair geometry from multiple images. *ACM Trans. Graph.* 23, 3 (Aug.), 712–719.
- G. J. Parker, C. A. Wheeler-Kingshott, and G. J. Barker. 2002. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Trans. Med. Imaging* 21, 5 (May), 505–512.
- I. Sadeghi, O. Bisker, J. De Deken, and H. W. Jensen. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.* 32, 2 (Apr.), 14:1–14:12.
- M. Sattler, R. Sarlette, and R. Klein. 2003. Efficient and realistic visualization of cloth. In *Proceedings of the Eurographics Symposium on Rendering*.
- K. Schröder, R. Klein, and A. Zinke. 2011. A volumetric approach to predictive rendering of fabrics. *Comput. Graph. Forum* 30, 4, 1277–1286.
- K. Schröder, A. Zinke, and R. Klein. 2014. Image-based reverse engineering and visual prototyping of woven cloth. *IEEE Trans. Visual. Comput. Graph. PP*, 99.
- T. Shinohara, J. Ya Takayama, S. Ohya, and A. Kobayashi. 2010. Extraction of yarn positional information from a three-dimensional CT image of textile fabric using yarn tracing with a filament model for structure analysis. *Textile Res. J.* 80, 7, 623–630.
- K. Srebrateep and E. L. J. Bohez. 2006. Computer aided modeling of fiber assemblies. *Comput. Aid. Des. Appl.* 3, 1–4, 367–376.
- K. Ward, F. Bertails, T. Yong Kim, S. R. Marschner, M. Paule Cani, and M. C. Lin. 2006. A survey on hair modeling: Styling, simulation, and rendering. *IEEE Trans. Visual. Comput. Graph.* 13, 2, 213–234.
- Y.-Q. Xu, Y. Chen, S. Lin, H. Zhong, E. Wu, B. Guo, and H.-Y. Shum. 2001. Photorealistic rendering of knitwear using the lumislice. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'01)*. ACM, New York, NY, 391–398.
- J. Zhang, G. Baci, D. Zheng, C. Liang, G. Li, and J. Hu. 2013. IDSS: A novel representation for woven fabrics. *IEEE Trans. Visual. Comput. Graph.* 19, 3, 420–432.
- S. Zhao, W. Jakob, S. Marschner, and K. Bala. 2011. Building volumetric appearance models of fabric using micro CT imaging. *ACM Trans. Graph.* 30, 44:1–44:10.
- A. Zinke and A. Weber. 2007. Light scattering from filaments. *IEEE Trans. Visual. Comput. Graph.* 13, 2, 342–356.

Received November 2014; revised July 2015; accepted August 2015