Specular Reflection from Woven Cloth

PITI IRAWAN and STEVE MARSCHNER
Cornell University

The appearance of a particular fabric is produced by variations in both large-scale reflectance and small-scale texture as the viewing and illumination angles change across the surface. This article presents a study of the reflectance and texture of woven cloth that aims to identify and model important optical features of cloth appearance. New measurements are reported for a range of fabrics including natural and synthetic fibers as well as staple and filament yarns. A new scattering model for woven cloth is introduced that describes the reflectance and the texture based on an analysis of specular reflection from the fibers. Unlike data-based models, our procedural model doesn't require image data. It can handle a wide range of fabrics using a small set of physically meaningful parameters that describe the characteristics of the fibers, the geometry of the yarns, and the pattern of the weave. The model is validated against the measurements and evaluated by comparisons to high-resolution video of the real fabrics and to BTF models of two of the fabrics.

Categories and Subject Descriptors: I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Color, shading, shadowing, and texture*; I.4.7 [**Image Processing and Computer Vision**]: Feature Measurement—*Texture*

General Terms: Design, Experimentation, Measurement, Verification

Additional Key Words and Phrases: Cloth, measurement, realism, reflection, texture

ACM Reference Format:

Irawan, P. and Marschner, S. 2012. Specular reflection from woven cloth. ACM Trans. Graph. 31, 1, Article 11 (January 2012), 20 pages.
DOI = 10.1145/2077341.2077352
http://doi.acm.org/10.1145/2077341.2077352

Funding was provided by the National Science Foundation under the awards CCF-0347303 and CCF-0541105, and by the Unilever Corporation.

Authors' addresses: P. Irawan and S. Marschner (corresponding author), Department of Computer Science, Cornell University; email: srm@cs.cornell.edu.

DOI 10.1145/2077341.2077352
http://doi.acm.org/10.1145/2077341.2077352

1. INTRODUCTION

Cloth is an important material to render convincingly because it is so often important in computer graphics scenes, especially those involving virtual humans in everyday environments. Accurate fabric appearance is particularly important in applications of computer graphics in the textile, garment, and fabric care industries. Our goal is to develop a simple, easy-to-use procedural model for the appearance of cloth that efficiently captures the important features of its appearance based on physically meaningful parameters.

In scenes rendered for computer graphics, two aspects of cloth appearance are important to capture in an appearance model. The directional reflectance, which describes the total light reflected from a large (at least several mm across) area of fabric, determines the overall shading. At the same time, the texture of the weave pattern is visible in more close-up views. Each weave has its own distinctive texture that is an important part of its appearance.

We assume that a general-purpose cloth model needs to be realistic at image resolutions up to a few pixels per yarn, when yarns are resolved but not individual fibers. Resolutions higher than this are in the realm of macrophotography and need to be rendered using a complete model of the cloth's three-dimensional structure.

This work makes two contributions: a set of measurements and a model to fit them. We present new, detailed measurements of the anisotropic Bidirectional Reflectance Distribution Function (BRDF) of six fabrics representing four textile fibers and the three most common weave patterns, as well as texture measurements for some of the fabrics. To study the appearance of the fabrics in context, we also took high-resolution video of the fabrics in a draped configuration under controlled conditions.

Our second contribution is a new reflection model for woven fabrics. The model is based on an analysis of specular scattering from fibers that are spun into yarns and then woven into fabric based on a given weaving pattern. Its parameters are all physically meaningful, describing the scattering properties of the fibers and the geometry of the yarns and weave. The model predicts both BRDF and, by a simple mapping of specular highlights onto the cloth surface, the texture of specular highlights. It defines a spatially varying BRDF that fits into standard realistic rendering systems and can be integrated over incident light using standard methods.

Our model, then, is a physics-based model. In contrast, it is popular to render cloth using data-based approaches, such as Bidirectional Texture Functions (BTFs). While physics-based models are derived from analysis of first principles, data-based models gather the reflection data by taking many pictures of the material to be modeled, store those data in a database, and query the database for the appropriate reflection data at render time. We shall compare the two distinct approaches in the next section, and compare their quality in the results.

We validate our model against our measurements and compare to renderings using BTFs. While the resulting appearance is not perfect, lacking some shadowing/masking effects and the irregularities of real cloth, our model predicts many key features both of the directional reflectance distribution and of the evolution of texture with viewing and illumination angle. Its accuracy compares favorably

with BTFs in many respects. Because all the directional variation in the model is due to specular reflection, an implication of this work is that specular reflection plays a more important role in the appearance of even quite matte fabrics than previously appreciated.

2. PRIOR WORK

While most of the work on modeling cloth for computer graphics has focused on motion rather than appearance, several researchers have addressed the problem of rendering cloth.

2.1 BRDF and BTF

The fundamental descriptions of appearance used for rendering cloth are the *Bidirectional Reflectance Distribution Function* (BDRF) and the *Bidirectional Texture Function* (BTF). The BRDF $f_r(\omega_i, \omega_r)$ is the ratio of radiance exiting a surface in the exitant direction ω_r to the irradiance arriving on the surface from an infinitesimal solid angle about the incident direction ω_i [Nicodemus et al. 1977; Dutré et al. 2003]. The BRDF is symmetric with respect to exchanging its arguments; that is, $f_r(\omega_i, \omega_r) = f_r(\omega_r, \omega_i)$. The BTF is a similar description but for texture: it gives the texture that appears in an image of a surface as a function of the incident and reflection directions [Dana et al. 1999]. Since BRDF is a function of four variables, BTF is a function of six variables; the additional two specify the texture coordinate.

2.2 Scattering from Fibers

Textiles are made out of fibers, and scattering from fibers has also been studied in the context of rendering hair, fur, and other materials. Models for scattering from cylinder-like dielectric structures have been developed [Kajiya and Kay 1989], some specialized for hair [Marschner et al. 2003; Zinke and Weber 2007]. Fiber scattering has also been used to create BRDF models for other fibrous materials, such as wood [Marschner et al. 2005]. All these models share the same basic specular-cone reflection geometry that we use in this article to predict specular reflection from textile fibers.

2.3 Measurements and Studies of Cloth BRDF

In the textile research community, luster was defined as a function of the ratio between specular reflection and diffuse reflection [Hunter and Herald 1987]. Buck and McCord provide some of the earliest quantitative measurements of luster of textiles [1949]. Among their findings are: fabrics made of filament fibers exhibit the greatest luster, yarn twist tends to reduce luster, and knitted fabrics exhibit less luster than woven fabrics.

Tao and Sirikasemlert measured specular reflection from single-jersey knitted fabrics made from monofilament yarns and developed a theoretical model of the reflection based on three parameters: fiber refractive index, yarn cross-sectional shape, and incident light angle [1999]. The model was later expanded to knitted fabrics made from twistless multifilament yarns [Sirikasemlert and Tao 1999]. Both models were developed to match goniophotometric measurements and no texture analysis or rendering was done.

In the computer vision community there is work involving reflection from fibers and woven materials. Lu et al. [1998] presented a measurement and study of velvet BRDF. They discovered that velvet cloth has a matte and diffuse reflectance with specular reflectance near grazing angles and retroreflection. The same team [Lu et al. 2000] later presented an analysis of the shape of specular highlights on fiber-covered surfaces based on geometric considerations similar to those we used to derive our model.

Ngan et al. [2005] measured velvet and two satin fabrics and fit analytical BRDF models to the measurements. They observed that the BRDFs of velvet and satin "far exceed the expressive power of simple analytical models" and approximated the cloth BRDF using a microfacet-based BRDF generator [Ashikhmin et al. 2000] with a tabulated microfacet distribution based on the measurements. This, however, requires high-resolution measurements of the cloth being modeled.

Pont and Koenderink [2003] presented a qualitative analysis of reflection from woven structures, emphasizing the double peak that is observed in some woven materials (including the polyester cloth measured in the present article). That work, unlike ours, did not aim to present a complete BRDF model or to predict texture.

2.4 Modeling Cloth BRDF

Cloth often appears as an example of an unusual BRDF. Westin et al. [1992] computed BRDFs for velvet and nylon by ray tracing models of the microstructure. In that work velvet was modeled as a collection of thin cylinders with randomly perturbed orientation. Yarns in the nylon cloth were modeled as flat cylinders and were interwoven according to the standard plain weave pattern.

Similarly, Volevich et al. [1997] ray traced a plane of interwoven yarns to study scattering from a piece of artificial silk. Unlike Westin et al. [1992], in that work the yarns were modeled as bundles of textile fibers, which in turn were modeled as very long and thin cylinders parallel to one another. These, therefore, were attempts to understand the appearance of woven (filament) cloth by explicitly modeling the structure of the cloth.

In their work on efficient rendering of Spatial Bidirectional Reflectance Distribution Functions (SBRDF), McAllister et al. [2002] measured anisotropic upholstery fabric and represented each texel using two Lafortune BRDF lobes.

Ashikhmin et al. [2000] dispensed with explicit models and used a combination of two cylindrical Gaussian slope distributions to model satin as an example of their microfacet-based BRDF generator. Velvet was another example, modeled using an "inverse Gaussian" heightfield.

Yasuda et al. [1992] presented a microfacet-based model that is compared to incidence-plane reflection measurements.

2.5 Modeling Cloth Texture

Other works have focused on the structure and texture of fabric. Adabala et al. [2003] presented a method based on a microfacet model and procedural textures that is capable of rendering cloth with a variety of weave patterns at different levels of detail. Without data to support the model, however, it is hard to judge its correctness.

Drago and Chiba [2004] modeled woven painting canvases with spline surfaces shaded by a procedural texture.

Xu et al. [2001] used a volume rendering approach called *lumislice* rendering to produce realistic close-ups of coarse knit fabrics. Their approach is related to our work because both consider a yarn as made up of helical fibers and take a volumetric approach to calculating the scattered light. The goals are different, however: our aim is an analytical model that works when yarns are barely resolved, whereas the lumislice was designed for closeups in which yarns are well resolved and fibers are prominent. Also, we focus on specular, rather than diffuse, reflection.

Wang et al. [2008] and Dong et al. [2010] both use a spatially varying BRDF based on tabulated normal distributions to represent a variety of measured materials including embroidered silk satin. The model does an excellent job capturing the spatially varying anisotropy of the embroidery, but it is not demonstrated on other

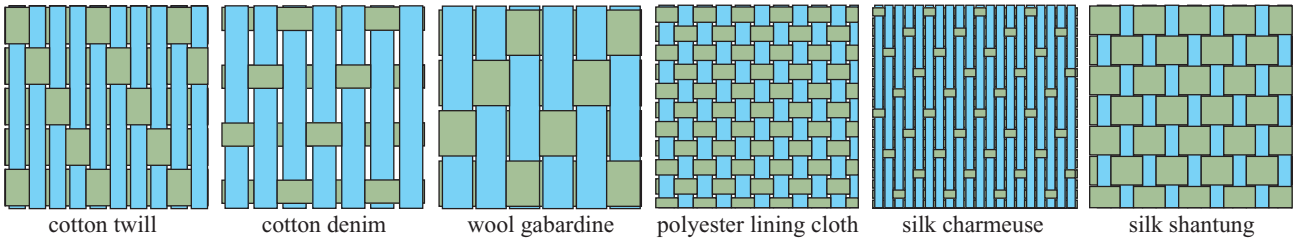

Fig. 1. Weave patterns of our sample fabrics.

fabrics or used at a resolution where the weave texture is visible, and it is not proposed or evaluated as a general model for fabric appearance.

A simpler method for generating cloth texture that is related in its structure to our model, but does not produce the same fidelity of BRDF or texture, is described in a technical report [Irawan and Marschner 2006].

2.6 Data-Based Approaches to Cloth Modeling

Because of its unusual BRDF and texture, simple analytical models often fail to represent cloth appearance well. One class of approach in cloth modeling abandons analytical models in favor of data-based ones. Data-based models, such as those of Daubert et al. [2001], Sattler et al. [2003], and Müller et al. [2005], start by taking many pictures of the cloth to be modeled and store the images as compressed Bidirectional Texture Function (BTF) [Dana et al. 1999] data in a database. At render time, the appropriate BTF data are then retrieved from the database.

By their nature, data-based models require large storage space and are able to model only the specific fabrics that have been captured and stored in the database. Compression methods have steadily improved, allowing higher resolutions to be used with lower runtime memory use [Havran et al. 2010], but BTFs still incur considerable acquisition time and storage cost compared to analytical models. Although some research has addressed the problems of editing BTFs [Kautz et al. 2007; Pellacini and Lawrence 2007; Wu et al. 2011], measurements of very similar materials will continue to be required, and the BTFs still cannot be controlled by parameters describing the structure of the fabrics. The comparisons we show in the results reveal that while BTFs produce a photographic appearance, undersampling and loss of grazing angles limits their quality for many fabrics.

Because our model was built from first principles and is analytical in nature, it doesn't require any data at render time. Measurements of cloth BRDF and texture discussed in this article were used only for study and verification of our model; the model itself does not require any data.

3. OVERVIEW

In this section we will talk about the structure of woven cloth and proceed to the overview of the model.

3.1 Structure of Woven Cloth

Woven cloth is constructed by interlacing two sets of parallel yarns, known as the *warp* and *weft*, at right angles to each other. In the process of weaving, warp yarns are raised or lowered and weft yarns (also known as *fillings*) are inserted in the resulting space. The pattern in which the warp and weft are interleaved varies greatly, but the majority of fabrics are made in one of the three simplest weave patterns: plain weave, twill, and satin [Parker 1993].

Different textile fibers are in common use in woven cloth: natural fibers (for example: cotton, wool, and silk); synthetic fibers (for example: polyester, nylon, and acrylic); and cellulosic fiber (for example: rayon and acetate). These fibers can be classified into two types: staple fibers and filament fibers.

Staple fibers—such as cotton and wool—are relatively short. To make staple yarns, staple fibers are twisted around one another so that they hold together by friction [Welford 1967]. Because of this twisted structure, the fibers on the surface of a fabric made from staple yarns appear in a diagonal arrangement, usually with alternating directions for exposed parts of the warp and weft. We use the term "staple" to refer to twisted staple yarn.

In contrast, filament fibers—such as silk and many synthetic fibers—are very long. As the result, filament yarns do not need to be twisted together in order to hold together. In this case the fibers lie parallel to the overall axis of the yarn. We use the term "filament" to refer to untwisted filament yarn.

Weaving creates a complex, regular geometry that can be considered, for purposes of appearance, to consist of a repeating pattern of visible segments of yarn (Figure 1). A warp yarn segment begins where the yarn emerges from behind one weft yarn, and continues until it next passes below another weft yarn (and similarly for weft segments). Interyarn forces make segments bend into curved shapes, convex toward the visible side. The degree of curvature is important to the appearance, and it depends on the stiffness of the yarn, the length of the segment, and the tension in the yarn and in the other yarns it interacts with. For instance, satin and twill weaves include long warp segments that will tend to lie flat and exhibit lower curvature than the shorter weft segments. A plain weave fabric may have similar yarn properties and tension in the warp and weft, leading to warp and weft segments with similar shape (e. g., the polyester fabric we measured); or it may be made with dissimilar yarns and/or tension, causing dissimilar segment shapes (e. g., the silk shantung fabric).

3.2 Overview of the Model

The idea behind our model is as follows: yarn segments are modeled as curved cylinders (Figure 2) made of spiralling fibers that reflect light specularly. As we will see later, specular reflection from the fibers forms a curved specular highlight on the surface of the segment. To get the total contribution to the BRDF from specular reflection, we can either integrate the reflection along the yarn segment (u direction) or around the yarn segment (v direction). Thus our BRDF model has two equations depending on how we choose to integrate the reflection.

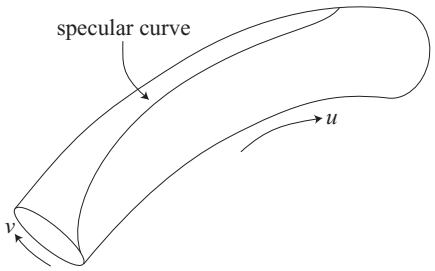

Fig. 2. A yarn segment with a specular reflection curve.

The amount of light that is reflected at one point on the specular reflection curve is $G_u f_c A$ or $G_v f_c A$ (depending on how we choose to integrate the reflection), which consists of the following terms.

(1) The geometry factor G_u or G_v. This is determined by the geometry of the yarn segment (including radius of curvature, size of the yarn segment, and change in specular reflection with change in illumination direction) and is discussed in Section 6.4.

(2) The phase function f_c. This function describes the local behavior of the fibers, and it should be chosen according to the actual behavior of the fibers being modeled. In this work we use a phase function that is the sum of a constant and a forward-directed lobe detailed in Section 6.5.

(3) The attenuation function A. This function describes the attenuation of light by other fibers on the way into and out of the yarn; it depends on the characteristics of the fibers as well as their microscopic arrangement. In this work we choose to use Seeliger's law as our attenuation function; this is described in Section 6.6.

Our model has two distinct incarnations: the reflectance model and the texture model. The reflectance model $f_r(\omega_i, \omega_r)$ is used when only the BRDF of the cloth is important (for example, in distant views of a large piece of cloth). The texture model $T(x, y, \omega_i, \omega_r)$, as the name implies, is used when the texture of the cloth is also important. Both models were built on top of the same set of assumptions and have the same average BRDF, which allows seamless switching between the two. (ω_i is incident direction, ω_r is exitant direction, and (x, y) is a point on the surface of the cloth.)

Our reflectance model consists of the following two functions.

$$f_{r,s}(\omega_i, \omega_r) = \int_{-u_{\max}}^{u_{\max}} G_v f_c A \, du$$

$$f_{r,s}(\omega_i, \omega_r) = \int_{0}^{2\pi} G_u f_c A \, dv$$

Similarly, our texture model consists of the following two functions.

$$T(x, y, \omega_i, \omega_r) \propto \chi G_u f_c A \frac{1}{\Delta x}$$

$$T(x, y, \omega_i, \omega_r) \propto \chi G_v f_c A \frac{1}{\Delta y}$$

The function χ equals 1 if the point (x, y) lands in the band of width Δx or Δy centered on the specular curve and 0 otherwise (see Figure 12). We shall elaborate on both models in later sections.

4. MEASUREMENTS

We made three types of measurements: reflectance (BRDF) measurements, closeup texture (BTF) measurements, and turntable videos. The BTF measurements were made to understand the behavior of the highlights; in this work we use them primarily for illustrative purposes. The BRDF measurements and turntable videos are used to validate our reflectance and texture model.

The fabrics we measured were:

(1) Black cotton fabric in a 3–1 twill weave.
(2) Denim, a cotton fabric with blue weft and white warp in a 2–1 twill weave.
(3) Red gabardine, a wool fabric in a 2–1 twill weave.
(4) Red polyester lining cloth with filament yarns in a very symmetric plain weave.
(5) Red charmeuse, a filament silk fabric in a satin weave.
(6) Red shantung, a filament silk fabric with red weft yarns and much finer dark gray warp yarns in a plain weave.

The weave patterns of the fabrics we measured can be seen in Figure 1. In this work, we follow the convention that the warp yarns run vertically in the figures.

4.1 Reflectance

To measure the BRDFs of our materials, we illuminated them with a light source of small solid angle (a DC regulated fiber-optic illuminator) and measured the reflected radiance by photographing them with a scientific CCD camera (QImaging Retiga 1300i, with frame-sequential RGB filter). The positions of the source and camera were controlled by a four-axis spherical goniometer. The linearity of the camera and stability of the source have been verified.

From the resulting images we computed the average of a small rectangle positioned at the center of rotation of the camera and source motion. The position in the image and with respect to the source were constant, eliminating the need for flat field calibration of the source or the camera, and the measured area was small enough to avoid significant variation in light source distance or incident angle over the measured area. The values were corrected for the cosine of the incident angle and normalized to a single measurement (per color channel) of a BRDF standard (Spectralon).

We measured datasets consisting of 225 incident directions for each of seven exitant directions. The incident directions are on a grid covering the hemisphere out to approximately 75 degrees, and the viewing directions coarsely cover the hemisphere (with the assumption of 180° rotational symmetry) out to 60 degrees.

The BRDF measurements can be seen in the left column of Figures 16, 17, 18, 19, 20, and 21. Each incident hemisphere is plotted in projection onto the tangent plane, with the warp direction vertical, and the hemispheres are arranged to indicate the exitant direction, which may also be seen by the shadow of the light source in the data. In the plots there is an obvious difference between filament yarns, which produce a pair of fairly classic anisotropic linear highlights (one from the warp yarns and one from the weft), and staple yarns, which produce still quite directional patterns but not distinct linear highlights. The BRDFs of staple fabrics are also asymmetric, even when the view direction is aligned with the warp or weft, because of the twist in the yarns. Also note that only the polyester is well balanced in the contribution of warp and weft; the others are all warp-dominated except shantung, which is heavily weft-dominated.

The plain weave filament fabrics both exhibit bright edges on the specular highlight, which are most noticeable on the polyester

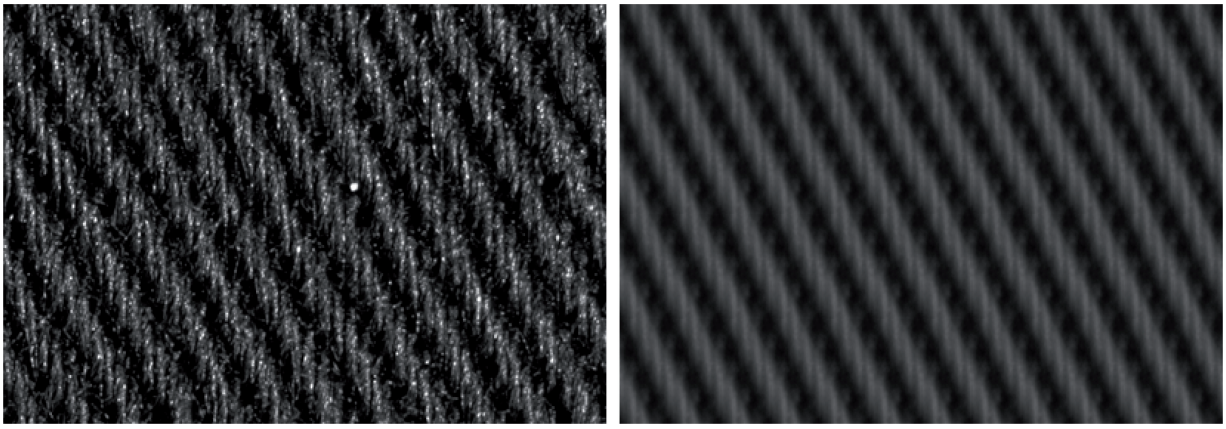

Fig. 3. Raw (left) and averaged (right) texture measurements of black cotton twill.

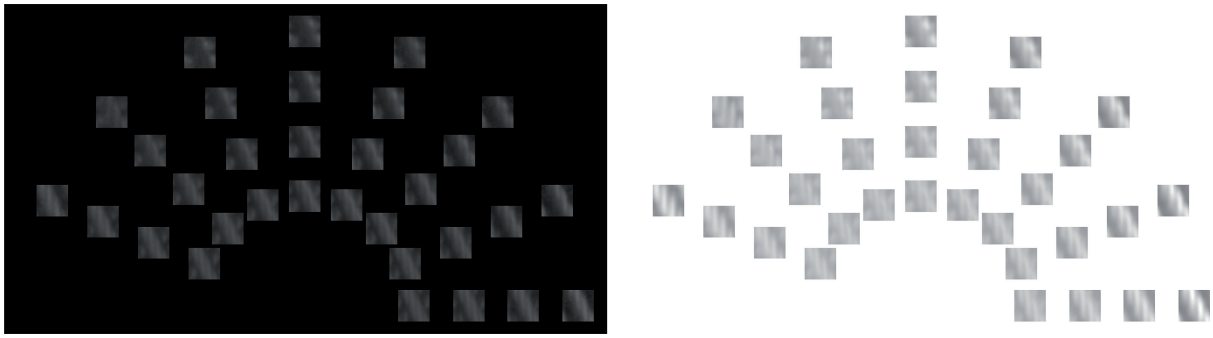

Fig. 4. Texture of black (left) and white (right) cotton twill under various illumination directions. The images are arranged in four half-circles, each represents the elevation angle of the illumination direction (30°, 45°, 60°, and 75°).

but also present on the warp component of shantung. This phenomenon has been explained by Pont and Koenderink [2003] as an effect of varying curvature of the yarns, with lower curvature towards the ends of the visible segment, and has also been observed by others [Ngan et al. 2005]. Most of the materials exhibit some retroreflection; and in particular the polyester shows a very sharp retroreflective peak that runs across the highlight (it is most noticeable in the 30° data). We believe that this is a result of interreflections between fibers of circular cross-section, but the phenomenon requires further study.

4.2 Texture

The second set of measurements was made using the same setup but with a macro lens at a magnification that enabled the yarns to be clearly discerned.

A representative frame from the measurements of a piece of black cotton twill cloth is shown in Figure 3. In the photographs, the overall pattern is difficult to discern because of the natural irregularities of the yarns. To remove this random variation and make the systematic pattern more visible, we computed a regularly tiled pattern by averaging all the unit tiles in the measured image. The averaged image of the same piece of black cotton twill cloth under the same condition is shown in the same figure.

Figure 4 shows the texture of black and white cotton twill under various illumination directions. A particularly interesting feature of this measurement is the similarity of texture between low and high reflectance fabrics. One might expect to see a similar specular component with a much larger diffuse component for white; in fact, the specular peaks in white are between 9 and 25 times brighter than those in black. This suggests that the light contributing to the specular highlights is not simply due to surface reflection (which should be unaffected by dyeing the fibers) but also includes substantial multiple scattering from well-aligned fibers, which, as has been observed in other materials [Marschner et al. 2005], continues to obey specular reflection geometry.

The similarity of these two textures suggests that specular reflection (including specular multiple scattering) is the main contributor to the texture of cloth. This is contrary to the commonly accepted notion that textures on matte-looking fabrics result primarily from diffuse reflection and shadowing–masking.

4.3 Turntable Sequences

To test our model in a more realistic context, we recorded high-resolution video of the same fabrics that were measured for BRDF under controlled conditions that allowed for comparison to renderings. To isolate the optical behavior from the confounding differences in appearance due to draping characteristics, we built a rigid form by coating draped black canvas with epoxy resin. The fabric samples were draped over the form in turn, ensuring that all the samples were photographed with the same geometry. The form also served to absorb the majority of transmitted light, thereby isolating reflection from transmission. We scanned the form with

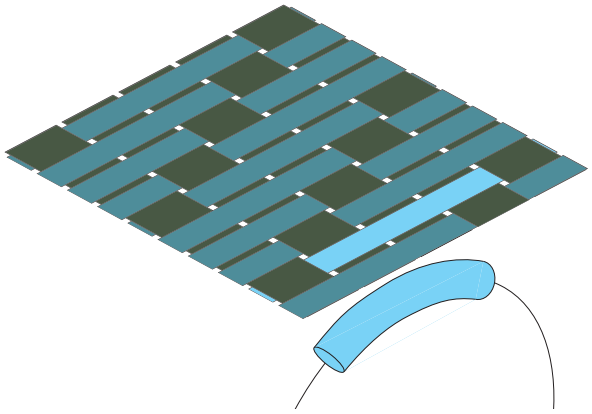

Fig. 5. A segment rectangle with its curved cylinder.

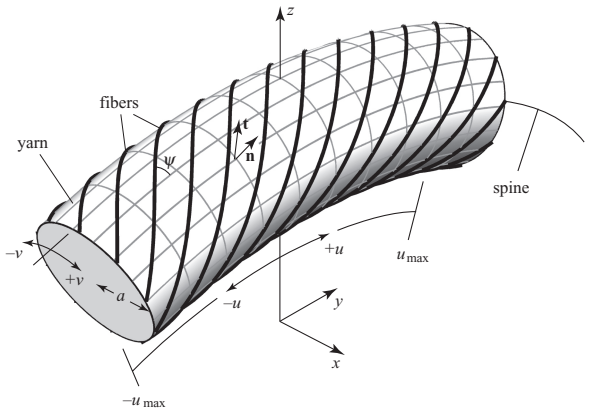

Fig. 6. A yarn segment modeled as a curved cylinder parameterized by $-u_{\max} \le u \le u_{\max}$, $-\pi \le v \le \pi$, and $0 < r < a$. Textile fibers form helices around the cylinder with a constant twist angle $-\pi/2 < \psi < \pi/2$.

a laser range scanner and fit a surface that was used for rendering the comparisons. The video was captured using stop motion with a high-resolution still camera (Canon EOS 20D). The motion sequence includes a segment where the object rotates with the light and camera stationary, and a sequence where the light moves with the object and camera stationary. The turntable sequences are described in Section 8.

5. GEOMETRY

A piece of fabric can be thought of as a collection of *segment rectangles*—short visible segments of yarn on the surface of the fabric—arranged in a particular position and orientation relative to one another according to the weave pattern. Each of these segment rectangles represents a yarn segment, which, in turn, is modeled as a curved cylinder made up of fibers spiraling around its axis. Figure 5 shows a segment rectangle with its curved cylinder. This section describes the geometry of the curved cylinder in detail.

5.1 Assumptions

A yarn is made up of relatively long fibers that may be twisted together. When a staple yarn is straight, we assume that the fibers are aligned with helices spiraling around the yarn axis and that the vectors tangent to fibers near the surface of the yarn all make the same angle with the yarn axis. When the staple yarn is bent into a curved configuration, we assume that it takes on the shape of a tube with curved spine and circular cross-section. We assume that the fibers' directions rotate with the cross-section, remaining at the same angle to the spine. Since filament yarns are not twisted, the fibers are simply parallel to the yarn axis.

5.2 Geometry of a Yarn Segment

The geometry of a yarn segment (Figure 6) is defined in a coordinate system that has z parallel to the overall normal to the fabric surface, y parallel to the relevant weaving direction (the warp or weft direction), and x completing the right-handed orthonormal basis.

We model a yarn segment as a curved cylinder: a circular cross-section with radius a swept perpendicularly along a spine curve $\mathbf{x}_0(u)$ in the y-z plane from $u = -u_{\max}$ to $u = u_{\max}$; here, $u_{\max}$ is the maximum inclination angle. We will discuss the parameter u a few paragraphs later.

Normally, the spine is a circular arc; resulting in a yarn segment in the shape of a torus segment; though some materials may require a different spine curve. The shape of the spine, however, only enters into the analysis through its curvature, denoted $R(u)$. When the spine is a circular arc, the radius of curvature of the spine is a constant R.

The yarn is parameterized by three variables: u, v, and r. The variable $-u_{\max} \le u \le u_{\max}$ is the angle between the spine's tangent and the y-axis (or, alternatively, between the spine's outward-directed normal and the z-axis). Parameterizing the spine by the angle u of course requires that each tangent angle occur only once. The variables $-\pi \le v \le \pi$ and $0 < r < a$ parameterize the circular cross-section for each u in polar coordinates.

The normal to the yarn surface $\mathbf{n}$ is a function of u and v

$$\begin{aligned}\mathbf{n}(u, v) &= R_x(-u)R_y(v)\begin{bmatrix}0\\0\\1\end{bmatrix}\\ &= \begin{bmatrix}\sin v\\ \sin u \cos v\\ \cos u \cos v\end{bmatrix}\end{aligned}$$

and the parameterization of the segment can be written as follows.

$$\mathbf{x}(u, v, r) = \mathbf{x}_0(u) + r\mathbf{n}(u, v)$$

As explained earlier, we assume that the tangents of the fibers are carried along with the cross-section. Like $\mathbf{n}$, they also rotate with v

$$\begin{aligned}\mathbf{t}(u, v) &= R_x(-u)R_y(v)\begin{bmatrix}-\sin\psi\\ \cos\psi\\ 0\end{bmatrix}\\ &= \begin{bmatrix}-\cos v \sin\psi\\ \cos u\cos\psi + \sin u \sin v \sin\psi\\ -\sin u\cos\psi + \cos u \sin v \sin\psi\end{bmatrix},\end{aligned}$$

where $-\pi/2 < \psi < \pi/2$ is the twist angle of the fibers. At $\psi = 0$ we have a filament yarn, in which $\mathbf{t}(u, v) = (0, \cos u, -\sin u)^T$ is independent of v.

5.3 Relation between a Segment Rectangle and Its Curved Cylinder

Recall that we break a piece of fabric into segment rectangles, each of which represents a yarn segment. Also recall that we use

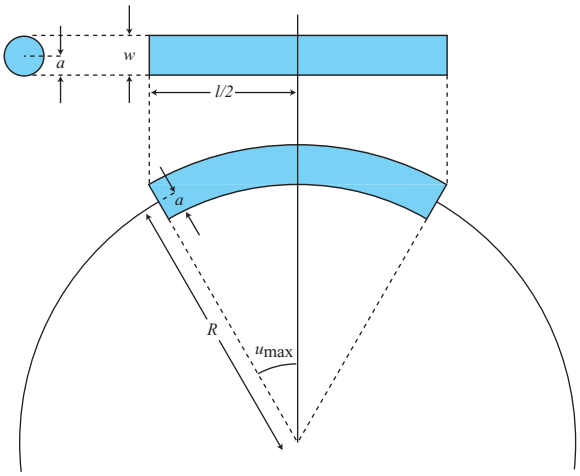

Fig. 7. Relation between a toroidal yarn segment and the segment rectangle.

a curved cylinder to model the yarn segment. We have examined the geometry of the curved cylinders, but we haven't discussed its relationship with the segment rectangle. This subsection illustrates the relationship for the simpler case involving circular spines (resulting in toroidal yarn segments). Details of this relationship for the general case involving noncircular spines are discussed in the next subsection.

Let w and l be the width and length of a segment rectangle. Given the maximum inclination angle $u_{\max}$, our goal is to find R (the radius from the center of the torus hole to the center of the torus tube) and a (the radius of the torus tube). We do this by choosing the largest torus segment whose projection fits in the segment rectangle. Figure 7 illustrates this concept.

From the figure, we can see that $a = \frac{w}{2}$ and thus

$$R = \frac{0.5l - a \sin u_{\max}}{\sin u_{\max}}.$$

Note that this imposes the following constraint: $\frac{w}{2} \sin u_{\max} < \frac{l}{2}$.

5.4 Spine Curves and Radius of Curvature

When the spine of the curved cylinder is a circular arc, the yarn segment is a segment of a torus. In reality, however, the shape of the spine curve depends on the weave pattern and the tension between the yarns of the fabric. Yarn segments in a satin cloth are usually flatter overall and more curved at the ends, while yarn segments in a plain weave cloth are usually more curved at the center. Figure 8 shows the cross-section of yarns arranged in plain weave and satin patterns. This section describes a way to adjust the curvature of the spine of the curved cylinder to control the shape of the yarn segment.

As in the previous subsection, the projection of the curved cylinder must fit in the segment rectangle (this implies that $a = \frac{w}{2}$ as shown in Figure 7). The spine curve is further constrained such that its tangent direction at the ends is $u_{\max}$. What we want is the ability to control the curvature of the yarn segment between its two ends.

We use conic sections to define the spine of the segment: ellipses for segments that are more curved at the ends, and ellipses, a parabola, or hyperbolas for segments that are less curved at the ends (that is, more curved at the center). The position and orientation of

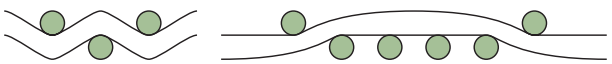

Fig. 8. Cross-section of yarns arranged in plain weave (left) and satin patterns (right).

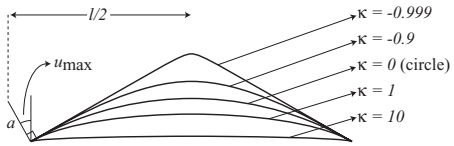

Fig. 9. Effect of κ on the shape of the spine curve.

the conic sections are not important since we care only about the radius of curvature of the segment.

Curvature of the segment is controlled by the spine curvature parameter $-1 < \kappa < \infty$. The spine curve is a segment of a circle (and the yarn segment becomes a segment of a torus) for $\kappa = 0$. The more negative κ is, the more curved the segment is at the center. The more positive κ is, the more curved the segment is at the ends. The effect of κ on the shape of the spine curve is shown in Figure 9.

Given κ, we compute the axis ratio $\hat{r}$ as follows.

$$\hat{r} = 1 + \kappa(1 + \cot u_{\max})$$

This variable determines whether the spine curve is a hyperbola, a parabola, or an ellipse; $\hat{r} < 0$ specifies a hyperbola, $\hat{r} = 0$ a parabola, and $\hat{r} > 0$ an ellipse ($\hat{r} = 1$ specifies a circle). Additionally, $\hat{r}$ relates the ellipse or hyperbola's semimajor axis $\hat{a}$ and semiminor axis $\hat{b}$ in the following way.

$$\hat{r} = \frac{\hat{b}}{\hat{a}}$$

Given $\hat{a}$ and $\hat{b}$, we can compute the radius of curvature $R(u)$ required in evaluating the scattering model. We have

$$R(u) = \begin{cases} \frac{(\hat{b}^2 \cos^2 t(u) + \hat{a}^2 \sin^2 t(u))^{1.5}}{\hat{a}\hat{b}} & \text{if } \hat{r} > 0 \\ 2\hat{b}(1 + t(u)^2)^{1.5} & \text{if } \hat{r} = 0 \\ -\frac{(\hat{b}^2 \cosh^2 t(u) + \hat{a}^2 \sinh^2 t(u))^{1.5}}{\hat{a}\hat{b}} & \text{if } \hat{r} < 0 \end{cases}$$

where

$$\hat{b} = \begin{cases} \frac{0.5l - a \sin u_{\max}}{\sin t_{\max}} & \text{if } \hat{r} > 0 \\ \frac{0.5l - a \sin u_{\max}}{2t_{\max}} & \text{if } \hat{r} = 0 \\ \frac{0.5l - a \sin u_{\max}}{\sinh t_{\max}} & \text{if } \hat{r} < 0 \end{cases}$$

and $\hat{a}$ is computed from $\hat{b}$ and $\hat{r}$.

The derivation of these formulae may be found in Section 6.4 of Irawan [2007].

6. REFLECTION

Recall from the previous section that we think of a piece of fabric as a collection of yarn segments, each modeled as a curved cylinder made up of fibers spiraling around its axis. Light that strikes these fibers reflects specularly into a cone centered on the local fiber axis (Figure 10). Different fibers reflect light that comes from the same direction into different cones, and by summing over all the fibers we can describe the scattering due to an entire yarn segment. The light scattered from the whole fabric is then simply a weighted sum of

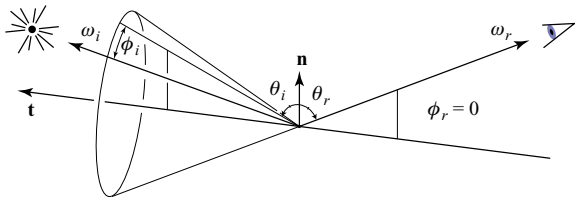

Fig. 10. Geometry of specular reflection from a fiber. Reflected light depends only on incident light within the specular cone.

the light scattered by the different segments, together with a diffuse component.

By summing their contributions in this way, interactions between segments, including masking, shadowing, and inter-reflection, are disregarded. The model nonetheless succeeds in capturing important visual features of the fabrics we have studied, which suggests that the local reflection geometry accounts for much of the appearance of woven cloth.

In this section, we describe light scattering from a yarn segment, derive the scattering function, and explain the various components of the function. While this material is important to the development of our model, readers interested only in the model itself may prefer to skip this section and go to Section 7 instead.

6.1 Assumptions

Our reflection model for fabric is based on some simple assumptions about the scattering behavior of the yarns that it is made from.

Since the fibers in a yarn are not tightly packed, the yarn must be treated as a volumetric medium, rather than as a reflecting surface. We do assume, however, that all important scattering happens close enough to the surface that the fiber direction is the same as for fibers on the surface.

Because most textile fibers are fairly specular and locally well aligned, we assume that local reflection from the fibers is ideally specular: all light from a single incident direction is reflected into the cone that has the same inclination to the fiber tangent. Of course imperfections in the fibers and random variations in fiber orientation mean the highlight will not be perfectly sharp, but other aspects of the geometry serve to blur the highlight into a smooth distribution, and as long as that blur is large compared to the width of the actual distribution it is safe to use the ideal specular model. This assumption is important because it restricts significant contributions to the scattering integral to happen only under certain geometric conditions, significantly simplifying the model.

A second simplification about local reflection is that all scattering that happens outside of a local area of well-aligned fibers is diffuse. This means that all directional effects are treated as single scattering.

6.2 Scattering from a Yarn Segment

The goal of this subsection is to compute the scattering function of a yarn segment, in isolation from the rest of the cloth. The scattering function, $f_s(\omega_i, \omega_r)$, describes the contribution of incident irradiance arriving from the direction ω_i to scattered intensity exiting in the direction ω_r. The total scattered intensity is the integral of the scattering function over incident light from the entire sphere (denoted by "4π" in the following).

$$I_r(\omega_r) = \int_{4\pi} f_s(\omega_i, \omega_r) L_i(\omega_i)\, d\omega_i$$

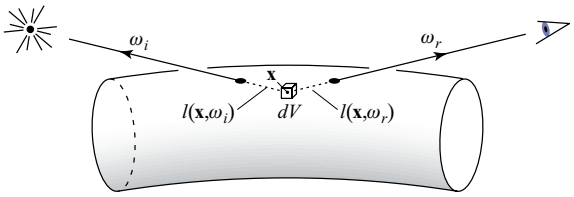

Fig. 11. Scattering from a volume in a yarn segment.

Under the assumption that a yarn acts like a single-scattering medium, we can compute the contribution of a volume element $dV(\mathbf{x})$ to the intensity scattered in direction ω_r by integrating over the incident radiance distribution $L_i(\omega_i)$. We have

$$\frac{dI_r(\omega_r)}{dV} = \int_{4\pi} \sigma_s f_p(\omega_i, \mathbf{t}(\mathbf{x}), \omega_r) e^{-\sigma_t l(\mathbf{x},\omega_i)} e^{-\sigma_t l(\mathbf{x},\omega_r)} L_i(\omega_i)\, d\omega_i,$$

where σ_s and σ_t are the volume scattering and attenuation coefficients, f_p is the phase function, and $l(\mathbf{x}, \mathbf{v})$ is the distance from the point $\mathbf{x}$ to the outside of the volume in the direction $\mathbf{v}$ (Figure 11).

We have stated the assumption that local scattering from the fibers is ideally specular, as illustrated in Figure 10. This makes this integral simpler than over the whole sphere because only light on the specular cone can contribute to the overall scattering. To write this integral we introduce a spherical coordinate system aligned with $\mathbf{t}$, where $\omega_i = (\theta_i, \phi_i)$ and $\omega_r = (\theta_r, \phi_r)$. As seen in the figure, $\sin\theta_i = \omega_i \cdot \mathbf{t}$ and $\phi_i = 0$ when ω_i is coplanar with $\mathbf{t}$ and $\mathbf{n}$. Similarly, $\sin\theta_r = \omega_r \cdot \mathbf{t}$ and $\phi_r = 0$ when ω_r is coplanar with $\mathbf{t}$ and $\mathbf{n}$. We denote the difference $\phi_r - \phi_i$ as ϕ. Ideal specular reflection occurs exactly when $\mathbf{h} \cdot \mathbf{t} = 0$, where $\mathbf{h}$ is the *half vector*, the bisector of the directions ω_i and ω_r. In this coordinate system, it means that light is only reflected from ω_i to ω_r when $\theta_i = -\theta_r$, as can be seen from Figure 10. This assumption about the phase function f_p can be expressed mathematically as a statement about the local scattering integral.

$$\int_{4\pi} f_p(\omega_i, \mathbf{t}, \omega_r) L_i(\omega_i)\, d\omega_i = \int_0^{2\pi} f_c(\theta_r, \phi) L_i(-\theta_r, \phi_i)\, d\phi$$

That is, the radiance scattered locally from the fibers is an integral of the incident radiance only over the specular cone; the rest of the incident sphere does not contribute. The function f_c is the "circular phase function," which describes how scattered light is distributed over the specular cone. (If we were to write an expression for f_p it would involve the product of f_c with a delta function in terms of θ.) We are assuming for simplicity that f_c depends on $\phi = \phi_r - \phi_i$ rather than on ϕ_i and ϕ_r separately. Thus, we obtain

$$\frac{dI_r(\omega_r)}{dV} = \int_0^{2\pi} \sigma_s f_c(\theta_r, \phi) e^{-\sigma_t (l(\mathbf{x},\omega_i) + l(\mathbf{x},\omega_r))} L_i(-\theta_r, \phi_i)\, d\phi.$$

To obtain the total scattered intensity for a segment of yarn viewed at a distance from direction ω_r, we simply integrate this expression over the segment's volume.

$$I_r = \iint_0^{2\pi} \sigma_s f_c(\theta_r, \phi) e^{-\sigma_t (l(\mathbf{x},\omega_i) + l(\mathbf{x},\omega_r))} L_i(-\theta_r, \phi_i)\, d\phi\, dV(\mathbf{x})$$

Note that θ_r, ϕ_i, and ϕ depend on $\mathbf{x}$.

For a segment parameterized as described in Section 5.2, we have

$$dV = dr \cdot r\, dv \cdot (R(u) + r\cos v) du.$$

Therefore

$$I_r = \int_{-u_{\max}}^{u_{\max}} \int_0^{2\pi} \int_0^{a} \int_0^{2\pi} \sigma_s f_c(\theta_r, \phi) e^{-\sigma_t l(\mathbf{x},\omega_i)} e^{-\sigma_t l(\mathbf{x},\omega_r)} \times L_i(-\theta_r, \phi_i)\, d\phi\, r(R(u) + r\cos v)\, dr\, dv\, du$$

where $R(u)$ is the radius of curvature of the spine. With the assumption that scattering happens near the surface, $\mathbf{t}$ doesn't depend on r, and we can replace the volume element $r(R(u) + r\cos v)$ with its value at $r = a$, leaving the attenuation $e^{-\sigma_t(l(\mathbf{x},\omega_i)+l(\mathbf{x},\omega_r))}$ as the only quantity depending on r. Let A, called the attenuation function, be defined as follows.

$$A(\omega_i, u, v, \omega_r) = \int_0^a \sigma_s e^{-\sigma_t(l(u,v,r,\omega_i)+l(u,v,r,\omega_r))} dr$$

Using this notation we arrive at

$$I_r = \int_{-u_{\max}}^{u_{\max}} \int_0^{2\pi} \int_0^{2\pi} f_c(\theta_r, \phi) L_i(-\theta_r, \phi_i) A(\omega_i, u, v, \omega_r) \times d\phi\, a(R(u) + a\cos v)\, dv\, du.$$

This integral is in the coordinates (u, v, ϕ), but these variables can be computed from ω_i and u or from ω_i and v. If we reparameterize this integral by (ω_i, u) or by (ω_i, v), we can move the integral over ω_i to the outside, then extract a scattering function from the equation. To reparameterize the integral with u on the outside, we need to express (ϕ, v) as a function of (ω_i, u) and find the Jacobian $|\partial(\phi, v)/\partial\omega_i|$. The integral then becomes

$$I_r = \int_{-u_{\max}}^{u_{\max}} \int_{4\pi} f_c L_i A \left| \frac{\partial(\phi, v)}{\partial \omega_i} \right| a(R(u) + a\cos v)\, d\omega_i\, du.$$

There will be zero, one, or two (ϕ, v) that satisfy $\mathbf{h} \cdot \mathbf{t} = 0$ for a given u and ω_i and, in general, we need to sum over the different solutions. However, the particular attenuation function A we use has the implication that at most one has a nonzero contribution.

To simplify, we introduce the *geometry factor*

$$G_v(\omega_i, u, \omega_r) = \left| \frac{\partial(\phi, v)}{\partial \omega_i} \right| a(R(u) + a\cos v)$$

and rearrange the equation into the form of a scattering integral

$$I_r(\omega_r) = \int_{4\pi} \left[\int_{-u_{\max}}^{u_{\max}} G_v f_c(\theta_r, \phi) A(\omega_i, u, v, \omega_r)\, du \right] L_i(\omega_i)\, d\omega_i$$

from which we can read off the scattering function as follows.

$$f_s(\omega_i, \omega_r) = \int_{-u_{\max}}^{u_{\max}} G_v f_c(\theta_r, \phi) A(\omega_i, u, v, \omega_r)\, du \qquad (1)$$

Similarly, if we reparameterize with v on the outside we have

$$f_s(\omega_i, \omega_r) = \int_0^{2\pi} G_u f_c(\theta_r, \phi) A(\omega_i, u, v, \omega_r)\, dv, \qquad (2)$$

where G_u is defined analogously to G_v.

These two integrals are equivalent except where the reparameterization fails. In particular, we cannot use u as the parameter for filament yarns (with $\psi = 0$) because $\mathbf{t}$ does not depend on v and therefore v cannot be written as a function of ω_i and u. We integrate over u for staple yarns and over v for the filament case.

Eqs. (1) and (2) define our model for the scattering function of a yarn segment. In the following subsections we provide detailed expressions for the terms therein, and then in Section 7 we explain how these pieces are used in evaluating our BRDF and texture models.

6.3 Finding Ideal Specular Reflection

In order to compute the integral in the previous subsection, we need to be able to express v as a function of u and vice versa. Geometrically, we want to find the value of v at which the ideal specular reflection takes place given a value of u and vice versa.

Recall that ideal specular reflection occurs exactly when $\mathbf{h} \cdot \mathbf{t} = 0$, where $\mathbf{h}$ is the half vector. This means that, for given incoming and exitant directions, only a one-dimensional subset of the surface of the yarn segment contributes to the specular reflection. Since the surface of the yarn segment is parameterized using u and v, we can write v as a function of u, incoming direction ω_i, and exitant direction ω_r. Similarly, we can express u as a function of v, ω_i, and ω_r.

Solving the equality $\mathbf{h} \cdot \mathbf{t} = 0$ for v given u, ω_i, and ω_r results in the following equation.

$$v(\omega_i, u, \omega_r) = \arctan(-\mathbf{h}_y \sin u - \mathbf{h}_z \cos u, \mathbf{h}_x) \pm \arccos(D) \qquad (3)$$

where

$$D = \frac{\mathbf{h}_y \cos u - \mathbf{h}_z \sin u}{\sqrt{\mathbf{h}_x^2 + (\mathbf{h}_y \sin u + \mathbf{h}_z \cos u)^2}} \cot\psi$$

If $|D| > 1$, no fiber tangent reflects light from ω_i to ω_r. It can be shown that only one of the two reflections satisfies $\mathbf{h} \cdot \mathbf{n} > 0$. The derivation of these formulae may be found in Section 7.3 of Irawan [2007].

For the filament case ($\psi = 0$), solving the equality $\mathbf{h} \cdot \mathbf{t} = 0$ for u given v results in the following equation.

$$u(\omega_i, v, \omega_r) = \arctan(-\mathbf{h}_z, \mathbf{h}_y) \pm \frac{\pi}{2} \qquad (4)$$

Similarly, it can be shown that only one of the two reflections occurs on the front side of the yarn (that is, $-\frac{\pi}{2} < u < \frac{\pi}{2}$). The derivation of these formulae may be found in Section 7.3 of Irawan [2007].

6.4 Geometry Factor

Computing the geometry factors in Eq. (1) or Eq. (2) requires evaluating the Jacobian of (ϕ, v) or (ϕ, u) with respect to ω_i and the curvature R.

For the Jacobian, we begin by observing that the allowed variation in ω_i is only in directions tangent to the unit sphere (since ω_i is a direction vector that cannot change length). Furthermore, ϕ is unchanged by a small change in ω_i perpendicular to the reflection cone, and u or v is unchanged by a small change in ω_i along the cone. So the determinant of the derivative is the product of the two directional derivatives

$$\left| \frac{\partial(\phi, v)}{\partial \omega_i} \right| = \left| \frac{\partial \phi}{\partial \mathbf{e}_1} \right| \left| \frac{\partial v}{\partial \mathbf{e}_2} \right|$$

$$\left| \frac{\partial(\phi, u)}{\partial \omega_i} \right| = \left| \frac{\partial \phi}{\partial \mathbf{e}_1} \right| \left| \frac{\partial u}{\partial \mathbf{e}_2} \right|$$

where $\mathbf{e}_1$ is the unit vector perpendicular to the cone at ω_i and $\mathbf{e}_2$ is the unit vector tangent to the unit sphere and to the cone at ω_i.

Because the derivation is fairly involved, here we present only the final results and refer interested readers to Section 7.4 of Irawan [2007].

$$G_v(\omega_i, u, \omega_r) = \frac{a(R(u) + a\cos v)}{|\omega_i + \omega_r|(\mathbf{n} \cdot \mathbf{h})|\sin\psi|} \qquad (5)$$

$$G_u(\omega_i, v, \omega_r) = \frac{a(R(u) + a\cos v)}{|\omega_i + \omega_r||(\mathbf{t} \times \mathbf{h})_x|} \qquad (6)$$

where $(\mathbf{t} \times \mathbf{h})_x$ means the x component of $\mathbf{t} \times \mathbf{h}$.

6.5 Phase Function

The phase function is a physical property of a particular type of fiber. Note that the desired phase function is not the phase function of an individual fiber but a phase function describing the effects of multiple scattering events occurring nearby in the yarn, all encountering the same fiber tangent. Since the fibers share the same tangent, the multiply scattered light will still stay in the specular cone, but will be more spread out around the cone.

Investigating the scattering properties of individual yarns and fibers in isolation to discover and model their behavior is an important research topic that is beyond the scope of the current work. Instead we use a generic phase function with the appropriate general properties that can be tuned to model different fibers. Preliminary measurements of single-fiber scattering, together with experience fitting the model to data, suggest that the phase function should be predominantly forward scattering, with a smaller uniform component. To this end we use a phase function that is the sum of a constant and a forward-directed lobe; we use the von Mises distribution [Evans et al. 2000], evaluated for the angle between the incident and exitant directions, for the lobe

$$f_c(\theta_r, \phi) = \alpha + g(-\omega_i \cdot \omega_r, \beta), \tag{7}$$

where

$$g(\cos x, b) = \frac{\exp(b \cos x)}{2\pi I_0(b)},$$

where α is the uniform scattering parameter, β is the forward scattering parameter, and $I_0(x)$ is a modified Bessel function of the first kind of order 0 [Arfken et al. 1995]. We chose the von Mises function because it is continuous around the circle and has proven to work well in practice.

6.6 Attenuation Function

The attenuation function A describes the attenuation of light by other fibers on the way into and out of the yarn. Our framework allows A to let light scatter through the fiber, even when the scattering point is not facing both the light source and the camera. After some experiments with sophisticated models for A, we found that a very simple model, which is the limit of the more general case for shallow penetration depths, worked well. In this limit the curvature of the yarn surface may be neglected and Seeliger's law, which describes scattering from a medium below a flat surface [Hanrahan and Krueger 1993], applies

$$A(\omega_i, u, v, \omega_r) = \frac{\sigma_s}{\sigma_t} \frac{(\omega_i \cdot \mathbf{n})(\omega_r \cdot \mathbf{n})}{\omega_i \cdot \mathbf{n} + \omega_r \cdot \mathbf{n}}, \tag{8}$$

where $\mathbf{n} = \mathbf{n}(u, v)$ and the dot products are all clamped to nonnegative values. This is the attenuation function we used for the results. The albedo σ_s/σ_t is unimportant because it can be absorbed into the specular coefficient.

For filament fibers, because the highlight will maintain full intensity right up to the moment it falls off the end of the segment (when $u(v)$ becomes greater then $u_{\max}$), it's necessary to include some form of smoothing at the ends of the integration domain, to simulate the gradual disappearance of the imperfect highlight in a real material (as contrasted with the sudden disappearance of the ideally sharp highlight in the model). We simply use a smoothstep cubic to fade out the contribution to the integral smoothly over an interval leading up to $u = u_{\max}$

$$A_s(u) = A(u)\left(1 - s\left(\frac{|u| - (1-\delta)u_{\max}}{\delta u_{\max}}\right)\right), \tag{9}$$

where $s(x)$ is a smooth step function that is 0 for $x \le 0$ and 1 for $x \ge 1$ and smooth in value and derivative in between, and $0 \le \delta < 1$ is filament smoothing parameter ($\delta u_{\max}$ is the size of the range over which the contribution ramps down).

7. REFLECTANCE MODEL AND TEXTURE MODEL

This section describes our two physically based appearance models for woven cloth: the reflectance model and the texture model. The reflectance model is used when only the reflectance of the fabric matters (for example, when the fabric is far enough from the camera that the texture is not visible, or when validating the model against BRDF measurements). In contrast, as the name implies, the texture model is able to model the texture of the fabric. Both models are based on the results developed in the previous two sections. Because of this, both models have the same BRDF and, therefore, switching between the models doesn't require any additional adjustments.

7.1 Reflectance Model

In the previous section, we derived the scattering functions $f_s(\omega_i, \omega_r)$ and explained the various components of the function. However, the function we actually need for rendering is the BRDF, $f_r(\omega_i, \omega_r)$, which describes the contribution of incident irradiance falling on the cloth surface from the direction ω_i to reflected radiance leaving the surface in the direction ω_r.

With no consideration for correlated shadowing–masking or inter-reflection, we can derive f_r from f_s by assuming that light scatters from a segment according to f_s regardless of where it hits the segment, and also that the scattered light has the same probability of escaping the surface regardless of where it leaves the segment. We apportion the incident irradiance uniformly to all the segments, so that each segment receives an average radiance of $L_i(\omega_i)(\omega_i)_z$ where $(\omega_i)_z$ is the z component of ω_i. The fraction of light escaping is also assumed to be proportional to $(\omega_i)_z$, and since the projected area over which it escapes is proportional to $(\omega_i)_z$, the radiance is simply proportional to the intensity of the segments.

This makes the relationship between f_r and f_s very simple: $f_r(\omega_i, \omega_r)$ is directly proportional to $f_s(\omega_i, \omega_r)$, and the constant of proportionality can be absorbed into the specular coefficient. Therefore f_s will serve directly as the specular component of our reflectance model.

There are, therefore, two ways to compute the BRDF, depending on how we parameterize the integral.

$$\begin{aligned} f_{r,s}(\omega_i, \omega_r) &= \int_{-u_{\max}}^{u_{\max}} G_v f_c A \, du \\ f_{r,s}(\omega_i, \omega_r) &= \int_0^{2\pi} G_u f_c A \, dv \end{aligned} \tag{10}$$

As noted in Section 6, we integrate over u for staple yarns and over v for filament yarns.

7.2 Texture Model

In order for a fabric to look realistic, the distinctive texture of reflections from individual yarns must be reproduced when the cloth is rendered at high enough magnification. All that is required for good results is to very roughly predict the position and shape of the highlight; if the magnification is high enough to resolve details within a yarn, a more detailed model such as lumislice rendering [Xu et al. 2001] must be used.

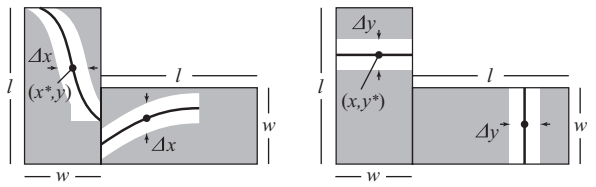

Fig. 12. The specular highlight in the texture is a fixed-width band around the ideal highlight curve. In the staple case (left) the width is constant in x, the direction across the yarn; in the filament case (right) the width is constant in y, the direction along the yarn.

Since our reflectance model already computes the highlight location in order to evaluate the various geometry-dependent terms, we can make use of this information to "unroll" the integrand into a texture in a way that satisfies the constraint that the average brightness in texture space equals the value of the BRDF. We do this by mapping u and v linearly to the segment rectangle on the cloth surface. In the texture space, the segment rectangle is parameterized by $-w/2 \le x \le w/2$ and $-l/2 \le y \le l/2$ (recall that w and l are the width and length of the segment rectangle). To unroll the yarn surface we map (x, y) to (u, v) using

$$\begin{aligned} u &= \frac{2u_{\max}}{l} y \\ v &= \frac{\pi}{w} x. \end{aligned} \tag{11}$$

This approach ignores visibility and foreshortening effects, but it nonetheless produces a realistic highlight texture.

The scattering model predicts an infinitely thin highlight, whose shape is defined by the function $v(u, \omega_i)$ or $u(v, \omega_i)$. We widen this curve into a band of constant width in the dependent coordinate: a constant width Δx for staple yarns and constant width Δy for filament yarns. Therefore, the texture model returns a nonzero value only if the point (x, y) lands inside this band of constant width. This process is illustrated in Figure 12.

We can find whether (x, y) lands inside the band as follows: first use Eq. (11) to map y (or x) to get u (or v). Then compute the location of ideal specular reflection $v(\omega_i, u, \omega_r)$ (or $u(\omega_i, v, \omega_r)$) using Eq. (3) (or Eq. (4)). Next use Eq. (11) to remap $v(\omega_i, u, \omega_r)$ (or $u(\omega_i, v, \omega_r)$) to get $x(v)$ (or $y(u)$). Finally, clamp $x(v)$ to the range $\pm(w - \Delta x)/2$ (or clamp $y(u)$ to the range $\pm(l - \Delta y)/2$). The point (x, y) lands inside the band if and only if $|x - x(v)| < \Delta x/2$ (or $|y - y(u)| < \Delta y/2$). We can encode this in a function $\chi(x, y, \omega_i, \omega_r)$ defined as follows.

$$\begin{aligned} \chi &= \begin{cases} 1 & \text{if } |x - x(v)| < \frac{\Delta x}{2} \\ 0 & otherwise \end{cases} \\ \chi &= \begin{cases} 1 & \text{if } |y - y(u)| < \frac{\Delta y}{2} \\ 0 & otherwise \end{cases} \end{aligned} \tag{12}$$

Given the function χ that defines the location of the highlight, it remains to determine the brightness. Recall that we have the constraint that the average brightness of the texture $T(x, y, \omega_i, \omega_r)$ in texture space equals the value of the BRDF $f_{r,s}(\omega_i, \omega_r)$. Mathematically, this is expressed as follows.

$$f_{r,s}(\omega_i, \omega_r) = \frac{1}{lw} \int_A T(x, y, \omega_i, \omega_r)\, dA$$

Fig. 13. The texture of black cotton twill at three different magnifications, each a factor of two from its neighbor. The simple, blocky shape of individual highlights is sufficient to represent the appearance of the ribs found in twill cloth.

Depending on how we parameterize the preceding equation, we have

$$\frac{1}{lw} \int_{-l/2}^{l/2} T(x, y, \omega_i, \omega_r) \Delta x\, dy$$

or

$$\frac{1}{lw} \int_{-w/2}^{w/2} T(x, y, \omega_i, \omega_r) \Delta y\, dx$$

for the staple or filament case. The brightness of the specular reflection, which varies along the highlight but not across it, is calculated to match the average value of the texture to the BRDF. To make these averages match Eq. (10) we need

$$\begin{aligned} T(x, y, \omega_i, \omega_r) &= \chi l w G_u f_c A \left|\frac{du}{dy}\right| \frac{1}{\Delta x} \\ &= \chi l w G_u f_c A \frac{2u_{\max}}{l} \frac{1}{\Delta x} \\ &= \chi 2 w u_{\max} G_u f_c A \frac{1}{\Delta x} \\ T(x, y, \omega_i, \omega_r) &= \chi l w G_v f_c A \left|\frac{dv}{dx}\right| \frac{1}{\Delta y} \\ &= \chi l w G_v f_c A \frac{\pi}{w} \frac{1}{\Delta y} \\ &= \chi \pi l G_v f_c A \frac{1}{\Delta y}. \end{aligned} \tag{13}$$

With the BTF defined in this way, the average value of the texture over a region of the image with constant shading geometry will match the value of the BRDF: in essence, the antialiasing filter of the rendering system is performing the integral that is done by quadrature in the reflectance model.

The result of the texture model for the black twill cloth can be seen in three different magnifications in Figure 13. Despite the simple, blocky shape of individual highlights, together they form an accurate representation of the ribs found in twill cloth and only at a very large magnification do they look artificial.

7.2.1 *Noise.* Since most textiles are not perfectly regular, we introduced two simple noise sources to improve the appearance of the renderings. Although noise is ad hoc and essentially separate from the model, the randomness is very important for visual quality.

To model irregularities in fiber structure, we scale the brightness of the specular component by a fixed noise texture with values drawn from the exponential distribution (between 0 and ∞ with mean 1). The noise is constant over each of a grid of k by k rectangles in a yarn segment. The parameter k controls the coarseness of the noise. This signal-independent multiplicative noise will not affect the average BRDF.

The shape of a yarn segment depends on the stiffness and the tension of the yarn, as well as the stiffness and the tension of

Table I. Parameters of the Reflectance Model

Parameter	Purpose	Typical Values
Fiber properties		
α	uniform scattering	0 to 0.1
β	forward scattering	2 to 5
δ	filament smoothing	0 to 1
Yarn geometry		
ψ	fiber twist angle	$-\pi/2$ to $\pi/2$
$u_{\max}$	maximum inclination angle	0 to $\pi/2$
κ	spine curvature	-1 to ∞
Weave pattern		
w	width of segment rectangle	0.1 mm to 1 mm
l	length of segment rectangle	0.1 mm to 1 cm
Coefficients		
k_s	specular coefficient	0 to ∞
k_d	diffuse coefficient	0 to ∞

the yarn crossing under the segment. In some materials these yarn properties vary significantly but slowly along the yarns, leading to a distinctive cross-hatch texture traditionally seen in linen and silk materials. We define a 1D Perlin noise function along each yarn and modulate $u_{\max}$ for a segment based on the noise values for its yarn and also the yarns crossing it.

7.3 Computing the Models

Our model defines a function of ω_i and ω_r based on the parameters in Table I, which describe the fibers, the yarns, and the weave pattern. All these parameters (other than the specular and diffuse coefficients) are directly meaningful in terms of the physical model of the fabric. A complete description of a fabric starts with a single set of fiber parameters and diffuse coefficients. Then for each distinct type of yarn segment in the weave pattern, we need a set of yarn and weave parameters and a specular coefficient. All the examples in this article have two distinct segment types, one warp and one weft.

The models are defined as the sum of a diffuse component and a specular component for each segment

$$k_d + \sum_j k_{s,j} f_{r,j}(\omega_i, \omega_r),$$

where k_d and the $k_{s,j}$ are the diffuse coefficient and the specular coefficients.

If one is interested only in the reflectance of the fabric but not in the texture, the reflectance model is sufficient. In the reflectance model, the BRDF $f_r(\omega_i, \omega_r)$ is evaluated by computing the integral in Eq. (10) using the parameters for the j^{th} segment type (normally one for warp and one for weft). In most cases this must be done numerically. The plots for this work were computed using the default quadrature routine in MATLAB. In practice, however, simple numerical integration methods (such as Trapezoidal Rule with 11 samples) are sufficient since the integrand is well behaved and no special precautions are required in integrating it.

If, instead, the texture of the fabric is desired, one has to use the texture model. In the texture model, as stated earlier, the integration in the BRDF $f_r(\omega_i, \omega_r)$ is performed automatically by the antialiasing filter of the rendering system. That is, only the integrand needs to be evaluated and no quadrature is required.

7.3.1 *Staple Yarn.* The algorithm to compute the integrand (a function of y) is as follows: Compute u from y using Eq. (11) and then find the point where the ideal specular reflection occurs $v(\omega_i, u, \omega_r)$ using Eq. (3). Next, we evaluate the geometry factor G_v using Eq. (5), the phase function f_c using Eq. (7), and the attenuation function A using Eq. (8). Multiply the three together and we have the integrand. For the reflectance model, depending on the numerical integration method used, this process is repeated several times on different points on the fabric. The pseudocode for the reflectance model on staple yarn can be seen in Algorithm 1.

ALGORITHM 1: Reflectance model on staple yarn

```
function StapleYarnReflectance ( ω_i, ω_r )
    f_r(ω_i, ω_r) = integrate(StapleIntegrand, [−u_max, u_max]);
    return k_d + k_s f_r(ω_i, ω_r);

function StapleIntegrand ( y )
    Compute u from y using (11);
    Compute v(ω_i, u, ω_r) using (3);
    if |v(ω_i, u, ω_r)| < π/2 then
        Compute G_v using (5);
        Compute f_c using (7);
        Compute A using (8);
        return G_v f_c A;
    else
        // ideal specular reflection is not visible
        return 0;
```

ALGORITHM 2: Reflectance model on filament yarn

```
function FilamentYarnReflectance ( ω_i, ω_r )
    f_r(ω_i, ω_r) = integrate(FilamentIntegrand, [−π/2, π/2]);
    return k_d + k_s f_r(ω_i, ω_r);

function FilamentIntegrand ( x )
    Compute v from x using (11);
    Compute u(ω_i, v, ω_r) using (4);
    if |u(ω_i, v, ω_r)| < u_max then
        Compute G_u using (6);
        Compute f_c using (7);
        Compute A_s using (8) and (9);
        return G_u f_c A_s;
    else
        // ideal specular reflection is not visible
        return 0;
```

For the texture model, given texture coordinates on the cloth, first find the type of yarn segment, j, that the shading point falls into and the (x, y) coordinates relative to the center of that segment's rectangle. Follow the same computation as for the integrand, then compute $x(v)$ from $v(\omega_i, u, \omega_r)$ using Eq. (11) and clamp it to the range $\pm(w - \Delta x)/2$. Compute χ using Eq. (12). Finally, compute the BTF $T(x, y, \omega_i, \omega_r)$ using Eq. (13) and return the sum of $k_s T(x, y, \omega_i, \omega_r)$ and k_d. The pseudocode for the texture model on staple yarn can be seen in Algorithm 3.

7.3.2 *Filament Yarn.* The algorithm to compute the integrand (a function of x) is as follows: Compute v from x using Eq. (11) and then find the point where the ideal specular reflection occurs $u(\omega_i, v, \omega_r)$ using Eq. (4). Next we evaluate the geometry factor G_u using Eq. (6), the phase function f_c using Eq. (7), and the

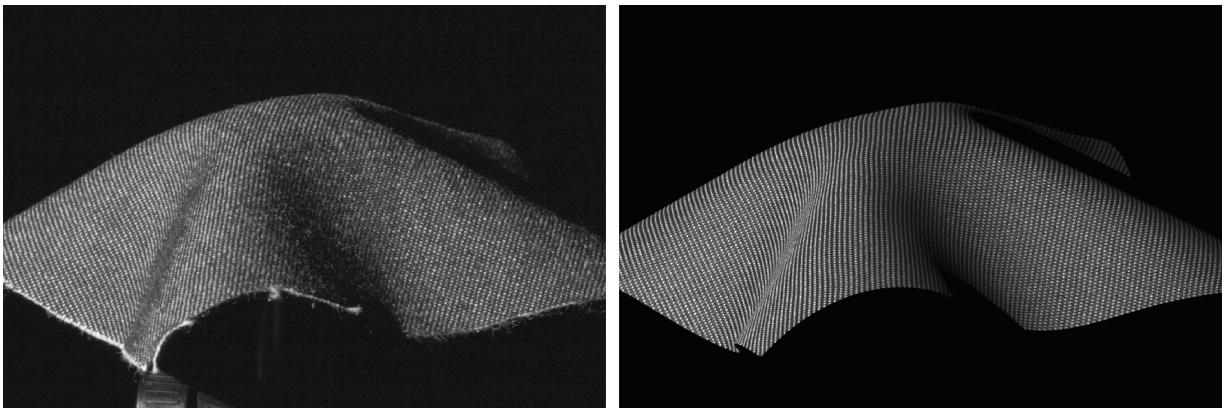

Fig. 14. The darkening effect of shadowing and masking can be seen at grazing viewing angle on denim (left image). Because our model lacks a shadowing and masking term, our rendered image (right image) doesn't have this darkening effect. This demonstrates a case where shadowing and masking clearly affect the BRDF of cloth.

ALGORITHM 3: Texture model on staple yarn

function `StapleTexture` (x, y, ω_i, ω_r)
- Compute u from y using (11);
- Compute $v(\omega_i, u, \omega_r)$ using (3);
- **if** $|v(\omega_i, u, \omega_r)| < \pi/2$ **then**
  - Compute G_v using (5);
  - Compute f_c using (7);
  - Compute A using (8);
  - Compute $x(v)$ from $v(\omega_i, u, \omega_r)$ using (11);
  - Clamp $x(v)$ to the range $\pm(w - \Delta x)/2$;
  - Compute χ using (12);
  - Compute BTF $T(x, y, \omega_i, \omega_r)$ using (13);
  - **return** $k_d + k_s T(x, y, \omega_i, \omega_r)$;
- **else**
  - `// ideal specular reflection is not visible`
  - **return** 0;

ALGORITHM 4: Texture model on filament yarn

function `FilamentTexture` (x, y, ω_i, ω_r)
- Compute v from x using (11);
- Compute $u(\omega_i, v, \omega_r)$ using (4);
- **if** $|u(\omega_i, v, \omega_r)| < u_{max}$ **then**
  - Compute G_u using (6);
  - Compute f_c using (7);
  - Compute A_s using (8) and (9);
  - Compute $y(u)$ from $u(\omega_i, v, \omega_r)$ using (11);
  - Clamp $y(u)$ to the range $\pm(l - \Delta y)/2$;
  - Compute χ using (12);
  - Compute BTF $T(x, y, \omega_i, \omega_r)$ using (13);
  - **return** $k_d + k_s T(x, y, \omega_i, \omega_r)$;
- **else**
  - `// ideal specular reflection is not visible`
  - **return** 0;

attenuation function A_s using Eq. (8) and Eq. (9). Multiply the three together and we have the integrand. For the reflectance model, depending on the numerical integration method used, this process is repeated several times on different points on the fabric. The pseudocode for the reflectance model on filament yarn can be seen in Algorithm 2.

For the texture model, given texture coordinates on the cloth, find the type of yarn segment, j, that the shading point falls into and the (x, y) coordinates relative to the center of that segment's rectangle. Follow the same computation as for the integrand, then compute $y(u)$ from $u(\omega_i, v, \omega_r)$ using Eq. (11) and clamp it to the range $\pm(l - \Delta y)/2$. Compute χ using Eq. (12). Finally, compute the BTF $T(x, y, \omega_i, \omega_r)$ using Eq. (13) and return the sum of $k_s T(x, y, \omega_i, \omega_r)$ and k_d. The pseudocode for the texture model on filament yarn can be seen in Algorithm 4.

8. RESULTS

We implemented our reflectance model in MATLAB for data fitting and the texture model in a Monte Carlo ray tracing system written in Java for rendering. In the renderer, the cloth model acts as a spatially varying BRDF. It receives texture coordinates, a shading frame, and incident and exitant directions, and it uses the texture model to compute a BRDF value that is returned to the system.

The general behavior of the model can be understood starting from the degenerate case of $\psi = 0$ and $u_{\max} = \epsilon$ (for a small nonzero ϵ), which describes a surface covered with parallel, perfectly specular fibers and would produce a very bright and thin anisotropic highlight (like a machined metal surface). As $u_{\max}$ is increased, the range of tangents present expands, causing the highlight to spread out. The distribution of intensity across the highlight is controlled by the shape of the yarn segment. A circular torus creates a fairly uniform highlight; a shape that is straighter near the ends leads to bright edges (as seen in the polyester); a shape that is flatter near the middle would lead to a highlight that falls off smoothly with long tails. In this filament mode, the model behaves somewhat like (though not identically to) a microfacet BRDF with a long, narrow facet normal distribution. Increasing ψ from zero also causes the highlight to broaden, but in a different and asymmetric way. The fiber parameters control the intensity distribution *along* the highlight. The weave parameters principally serve to establish the texture and to balance the brightness of warp and weft,

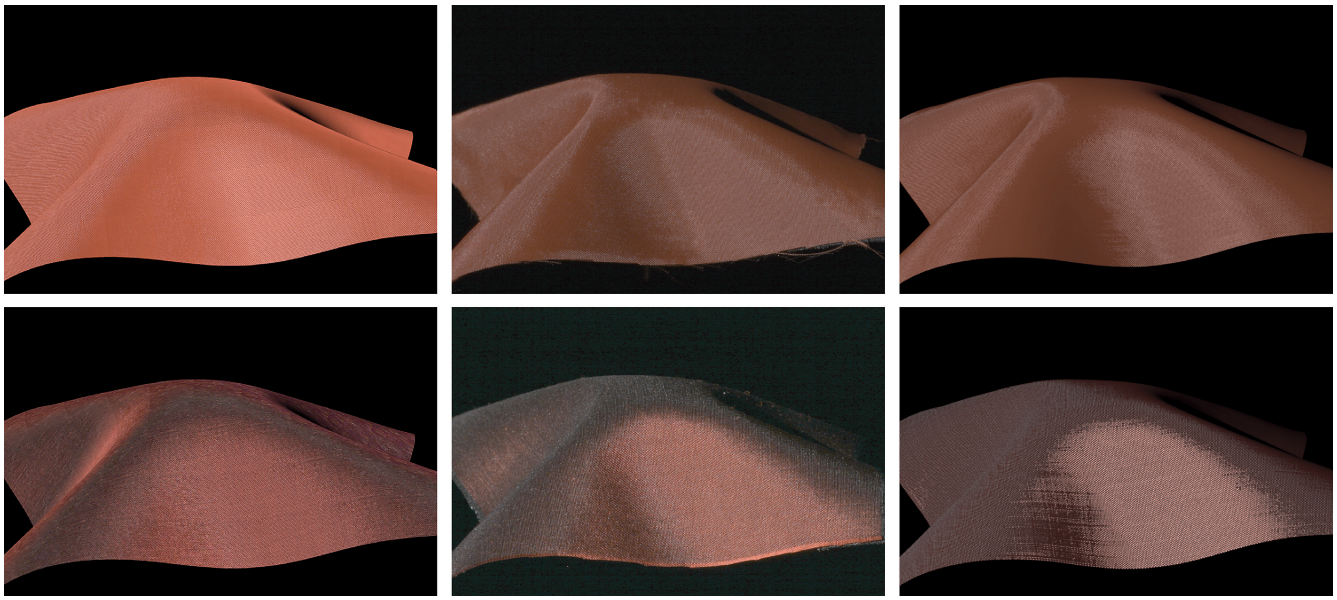

Fig. 15. Comparisons between renderings from minimally compressed 6561-image captured BTFs (left), the reference photos (center), and our analytic model (right). The captured BTF data produces a more photographic appearance, particularly for the silk shantung fabric, but it produces smooth highlights that miss some details and quite wrong appearance at grazing viewing and illumination angles. (Colors are not directly comparable because the BTFs were captured using a different camera and color processing pipeline.)

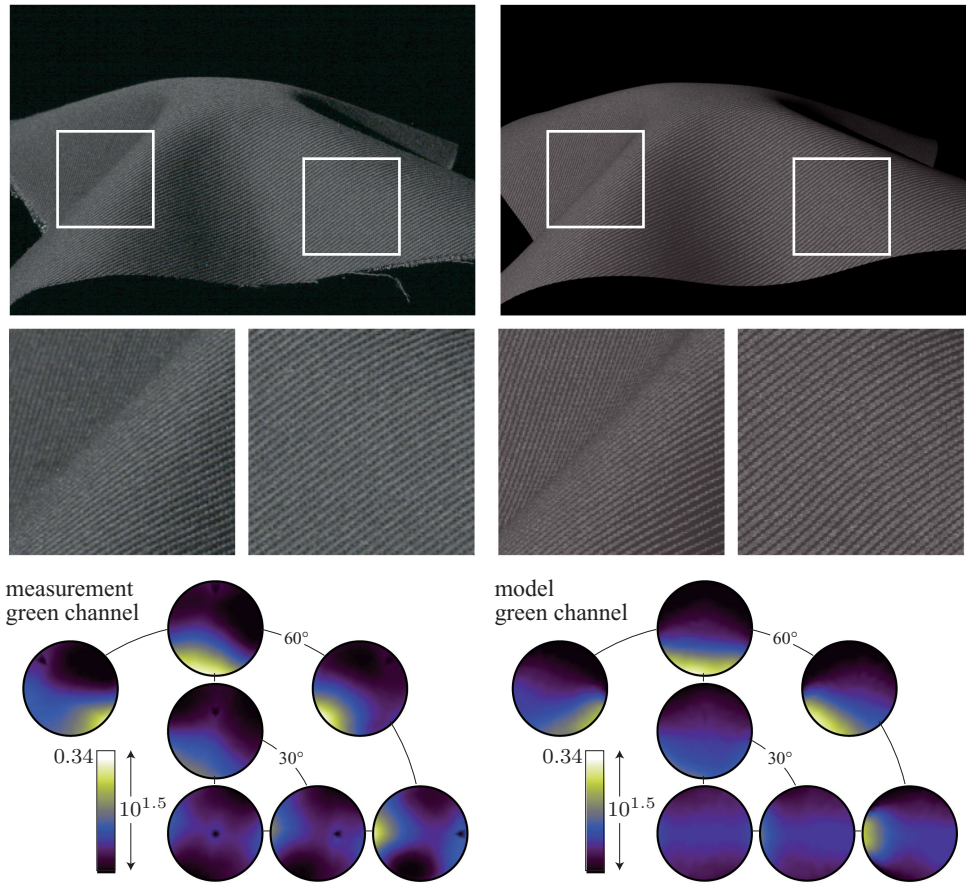

Fig. 16. Cotton twill comparison. This solid black fabric shows a moderately directional highlight and a diagonal texture.

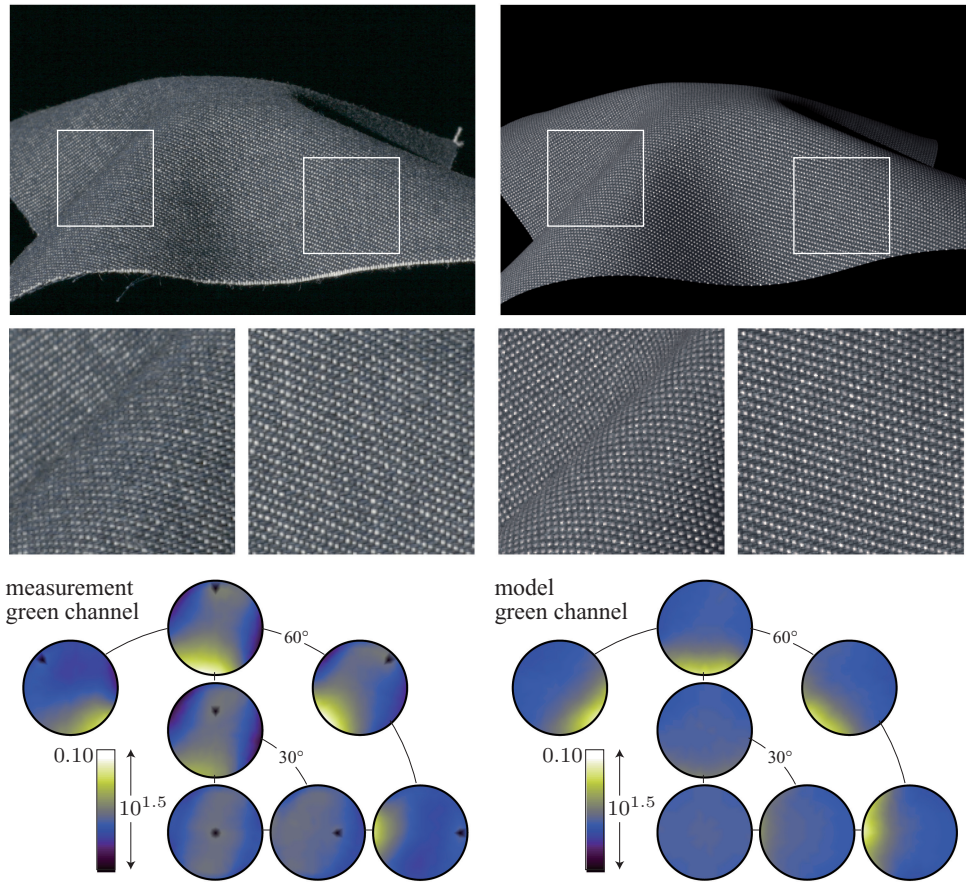

Fig. 17. Cotton denim comparison. The texture of white weft dots and blue warp highlights is well simulated.

though they do subtly affect the reflection pattern by affecting the relationship between R and a.

To compare our model to the BRDF data, we selected parameters by a combination of direct measurement and manual and automatic fitting. The model produces complex multi-component BRDFs, and fully automatic fitting proved unreliable because of the difficulty of balancing fit in highlights against more diffuse regions, fitting a weak weft component underneath the much brighter warp residual, and balancing texture appearance against BRDF fit.

The weave dimensions and approximate values of ψ were measured by observing the samples under a microscope, then $u_{\max}$, ψ, α, and β were chosen by plotting the model BRDF for a coarse grid of parameters and picking parameters to yield a good match to the measurement BRDF. For each setting of these parameters, automatic linear fitting was used to obtain specular and diffuse coefficients to match the data. Comparisons between the reality and the model can be seen in Figures 16, 17, 18, 19, 20, and 21.

The median relative error was 21%, 18%, 12%, 19%, 40%, and 28% for cotton twill, cotton denim, wool gabardine, polyester lining cloth, silk charmeuse, and silk shantung, respectively. The 80$^{\text{th}}$ percentile relative error was 36%, 34%, 27%, 39%, 69%, and 54%. Note that our fitting procedure did not attempt to minimize these error metrics.

Using the same parameters as in the BRDF comparisons, we rendered animation sequences to match the known viewing, illumination, and surface geometry from the turntable videos (described in Section 4). To compare them to the videos (captured with a different camera), we computed a color space transformation from photographs of a standard color chart and applied it to put the rendered images in the color space of the photographs. Please refer to the accompanying video for the complete turntable sequences. Weave patterns of the fabrics we analyzed are shown in Figure 1.

8.1 Staple Fabrics

The BRDFs of staple fabrics are asymmetrical with a forward-scattering lobe. This general structure is well matched by our model (though our forward-scattering lobe is not as strong as the reality). One interesting feature in the BRDFs is the darkening at grazing angles due to shadowing and masking. This effect can be observed in the accompanying turntable video, in particular on denim. However, it is relatively minor (see Figure 14) and suggests that shadowing and masking have less influence on the BRDF of staple cloth than is commonly accepted.

Because of the fairly subtle BRDFs, texture is often the prominent feature of staple fabrics. Our model is able to replicate the twill ribs in cotton twill, including the small dots of reflections between the ribs. For cotton denim, we are able to model the white dots seen in the photograph and also the thin slivers of blue reflections between the white dots. The wool gabardine is coarser than the cotton twill with a higher ratio of weft to warp area and the red dots between

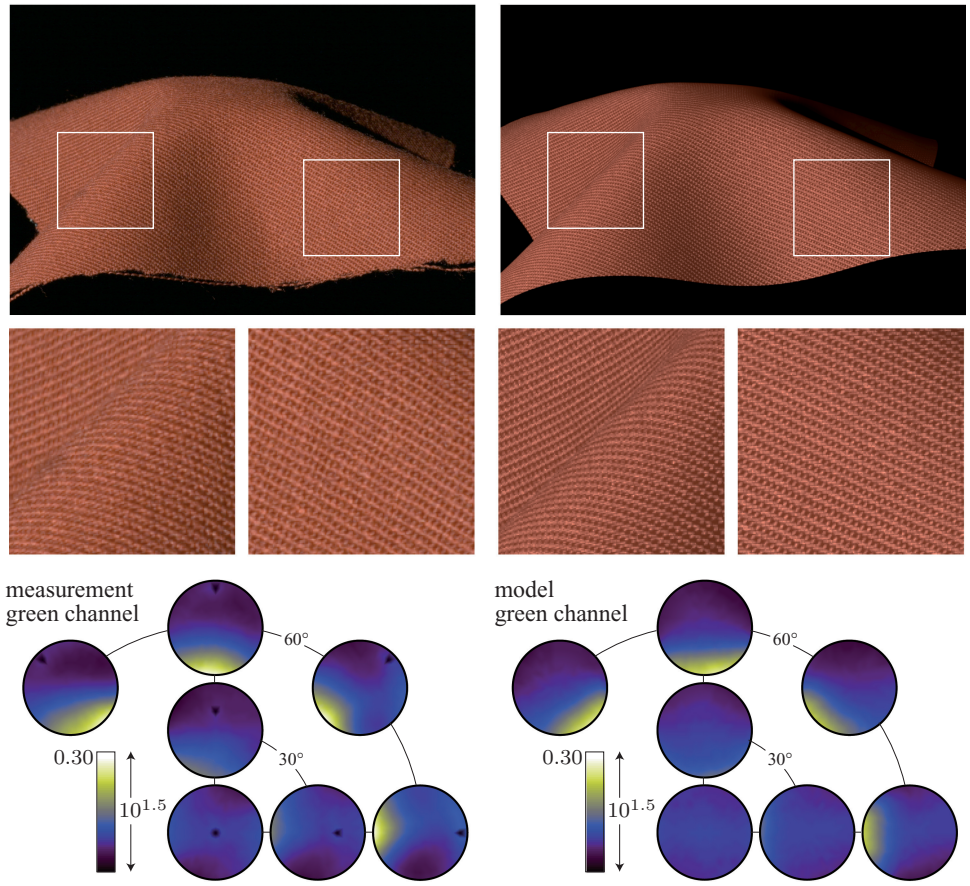

Fig. 18. Wool gabardine comparison. The texture of this quite matte material still is well predicted by specular highlights.

the twill ribs are easier to notice here. Even though our texture model operates purely in a 2D manner (simple, blocky highlights on rectangles), the three fabrics rendered using our model give an illusion of depth and three-dimensional structure.

8.2 Filament Fabrics

The contribution of the warp and weft yarns to the BRDFs of filament fabrics can be easily discerned from looking at the BRDF plots. Filament fabrics with warp and weft yarns of the same color and brightness, such as our polyester lining cloth, have a cross-shaped BRDF. In warp-dominated filament fabrics, such as our silk charmeuse, one of the bars of the cross is less prominent; conversely, in weft-dominated fabrics, such as our silk shantung, the other bar is less prominent.

Polyester lining cloth demonstrates a superposition of two highlights, which is also predicted by our model. It also shows some irregularities along the edge of the highlights, which we modeled using the correlated noise described in the previous section. The edge-brightening effect seen in the photograph and the BRDF is modeled effectively by using a hyperbolic shape for the segment spine (as is appropriate for a tightly woven plain weave fabric). Silk charmeuse has much thinner and brighter highlights with no apparent edge-irregularity. Because this is a warp-dominated satin weave, we expect the warp segments to be prominent and the weft segments less so, which can be seen both in the BRDF and in the video. Silk shantung is woven with bright red weft yarns and darker warp yarns. Up close, the texture looks like a grid of red dots. Shantung shows strong irregularities along the edge of its highlights in form of cross-hatch formations. Similar to the polyester lining cloth, shantung is a plain weave fabric and shows a mild edge-brightening effect. The lack of good match for highlight profile in silk can be attributed to the properties of silk fibers (such as their unusual cross-section) to be studied in future work.

8.3 Comparison to BTF renderings

A common approach to rendering complex materials like cloth is to use Bidirectional Texture Functions (BTFs) [Dana et al. 1999; Sattler et al. 2003], in which photographs of the cloth surface under the full range of illumination and viewing conditions are interpolated to produce view- and illumination-dependent textures. BTFs inherently produce photographic-looking textures, but they have several inherent drawbacks. Because they interpolate between discrete view and illumination directions, they are prone to sampling artifacts in directional space, which take the form of smoothed highlight structure and highlight strobing. Also, because capture setups are limited in their ability to approach grazing viewing and illumination conditions, grazing angles are not captured and must be extrapolated from the data. And finally, like other image-based approaches, BTFs have large storage costs and limited editability.

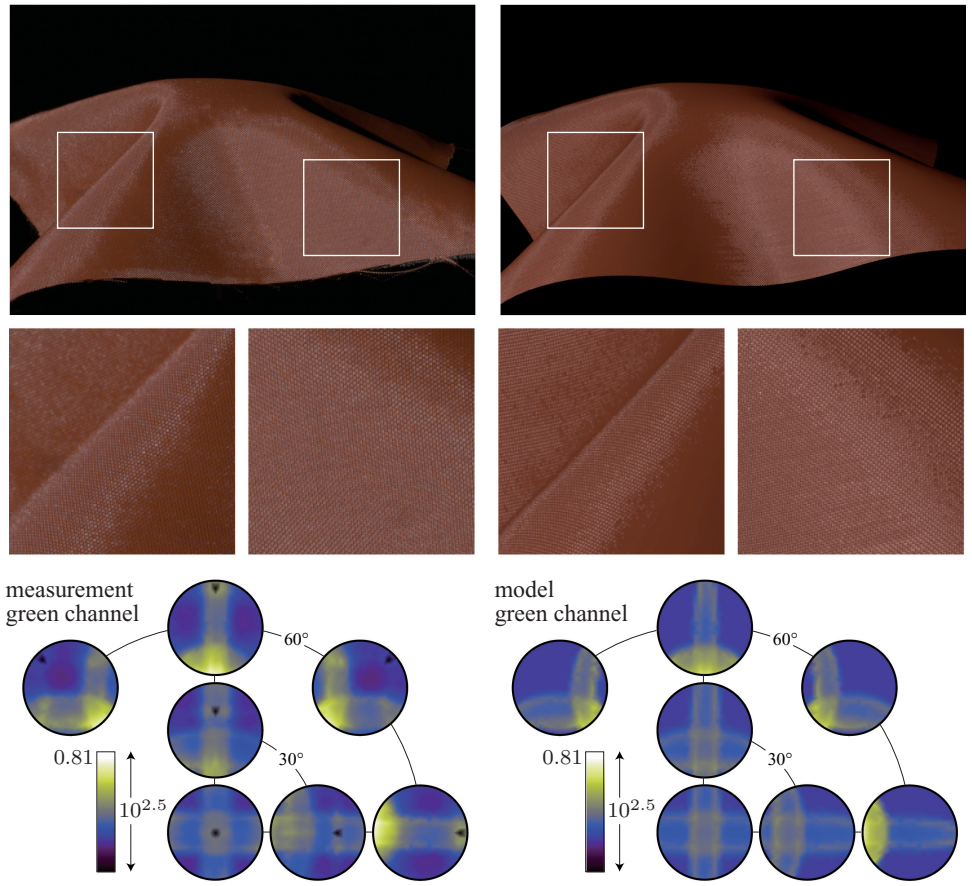

Fig. 19. Polyester lining cloth comparison. A spine with lower curvature at the ends of the segment produces the bright-edged highlight.

To evaluate the relative performance of BTFs compared to our analytic model, we captured BTFs for two of our materials, using the same goniometer to position the camera and source. (The camera, lens, and color processing software for this setup are different, which causes some color difference between the datasets.) Following previous work [Sattler et al. 2003; Müller et al. 2004], we used a grid of 81 camera and 81 source positions covering the incident and exitant hemispheres out to 75°, totaling 6561 images per material.[1] The images were registered using projective warps, and the aligned images were cropped to 256 × 256 textures with a resolution of roughly 25 pixels per millimeter. Our GPU BTF renderer stores the images in texture memory and uses a GLSL shader to interpolate the textures for rendering. Only minimal per-image texture compression was used to fit the BTF into texture memory; the results are equivalent in quality to renderings using uncompressed textures. Note that tiling artifacts in the BTF images could be alleviated by the common procedure of processing the textures to hide the boundaries, a step which we omitted in order to avoid any possible impact to the quality of the results. Also, our BTF implementation does not compute cast shadows.

[1] Although some special setups have achieved higher angular sampling rates for BTFs (e.g., 151 by 151 cameras/lights [Müller et al. 2005]), we feel our measurements are representative of common practice.

The results of these comparisons can be seen for a single view in Figure 15 and for varying view and lighting in the accompanying video. The BTF model produces a more photographic look than the analytic model, but its limitations are quite apparent for these materials. In particular, the complex structure of the highlight in the polyester fabric is significantly blurred; the warp and weft highlights are less distinct, and the bright-edged structure is nearly blurred away. Also, because the BTF data ends at 75°, there is no information for grazing angles, and the appearance is quite wrong along the silhouette and for grazing illumination. This material has a highly saturated red color that reproduced differently in the different camera systems used for the two types of model; the difference in base color should not be taken as a difference between the two techniques. Comparing the results for silk shantung, the BTF model reproduces the inhomogeneous texture of the fabric better, and produces smoother results. The subtle gray warp highlight (visible best in the video), which is predicted well—though rendered too homogeneously—by the analytic model, is almost completely missing in the BTF renderings. The appearance of the BTF is quite wrong around grazing angles. The characteristic glitter of this material is missed by both models.

From these comparisons we conclude that the BTF and analytic models both have their strengths and weaknesses; neither is uniformly better in comparison to the true appearance. Besides the quality differences it must be remembered that the BTF model comes with much larger memory requirements, requires an

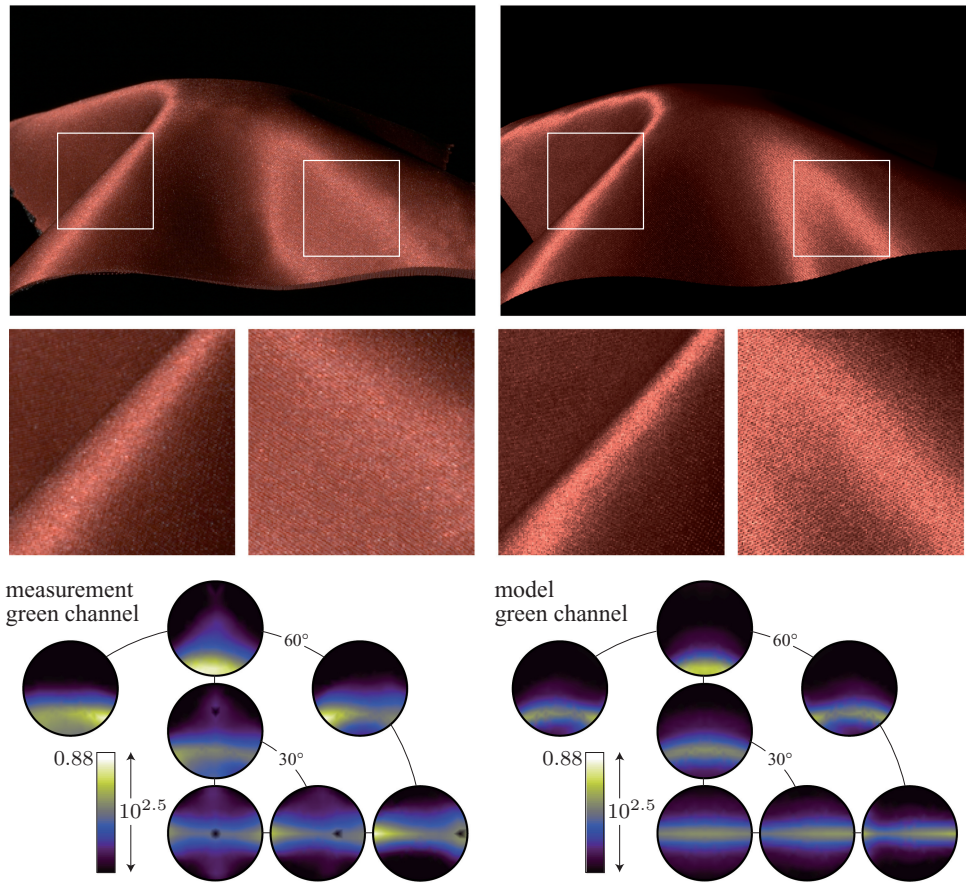

Fig. 20. Silk charmeuse comparison. For this fine, smooth fabric the texture is barely visible.

elaborate capture process, and can only be used for materials that have been captured; on the other hand, the analytic model is much easier to tune but requires a more elaborate shader and is harder to antialias efficiently.

9. CONCLUSION

This work has presented an extensive study of light reflection from woven fabrics. Our reflectance measurements show a variety of phenomena, ranging from sharp anisotropic highlights to asymmetric non-Lambertian diffuse patterns, and our model demonstrates that most of these features can be explained as resulting from specular reflection, once the structured geometry of the material is taken into account. The textures we produce, again using only specular highlights, capture key appearance characteristics over a remarkable range of conditions. These results are in contrast with the prevailing assumption that the most important features of the reflectance, and especially the texture, of fabrics with generally matte appearance are due to diffuse reflection and shadowing–masking.

We expect our model will be useful in practice wherever realistic cloth appearance is needed. Although the derivation is fairly involved, the texture model itself is not difficult to evaluate. Although the BRDF model requires a numerical integration, the integrand is well behaved and the integration can be done with a simple numerical integration routine.

One major advantage of using our model is that it doesn't require any data and thus can be used to model an arbitrary piece of fabric, even one that is not available or has not been manufactured. In contrast, data-based models do require BTF data of the fabric to be modeled. Not only does this require a large storage space, but it also is able to model only the specific fabric that has been captured and stored in the database. Our model also comes with physically meaningful parameters, which is very important for the users of this model to be able to tweak the appearance of the fabric to suit their particular need. This connection to the fabric properties is also important for textile applications.

The model presented in this article has certain limitations, but the framework we have established can be used to build other, more sophisticated fabric scattering models. The empirical forward-scattering phase function we use is very generic; research should be done into appropriate phase function models for various types of textile fibers. More sophisticated models for attenuation can be used. We have ignored interyarn interactions in order to concentrate on specular reflection, but some effects of these interactions are visible in the data. For instance, the white weft dots in denim disappear in the turntable photographs at grazing angles, and the warp component of charmeuse shows a sharpening for grazing angles that we conjecture is a shadowing effect (vertical features in the center column of the data). Also, shadowing lends additional contrast to the texture of the twill ribs in gabardine and denim. In order to

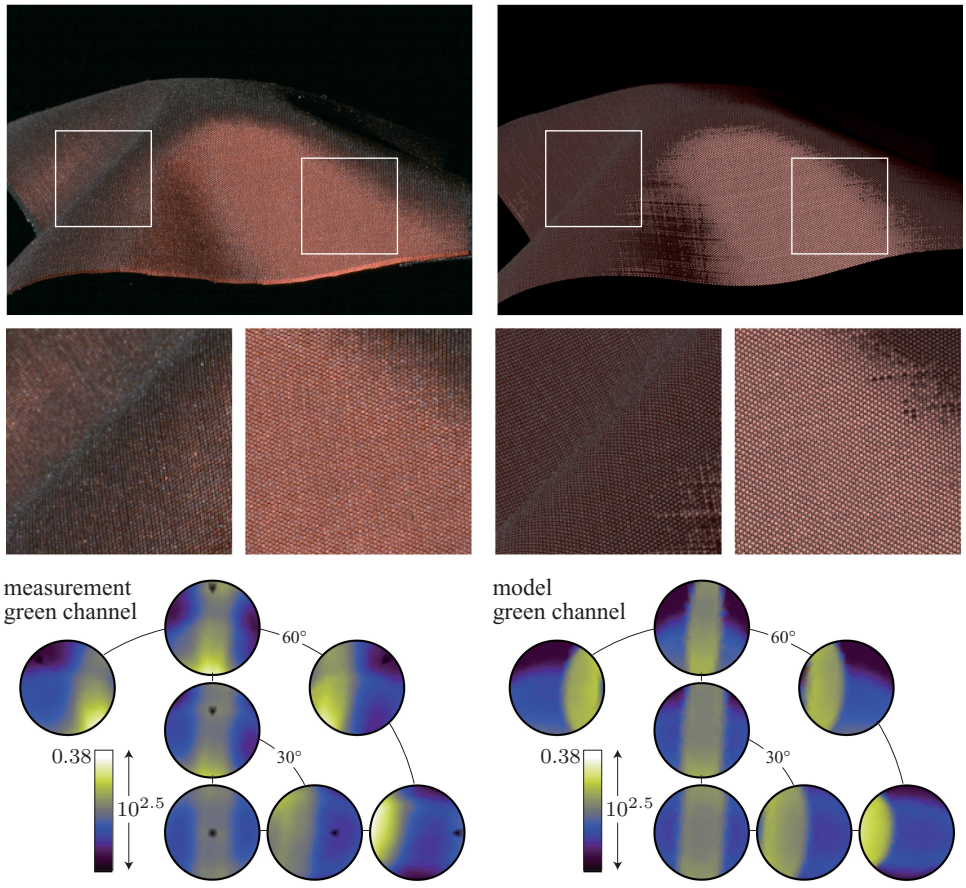

Fig. 21. Silk shantung comparison. This fabric is weft-dominated and displays a yarn-correlated texture.

correctly predict the appearance of materials, such as shot silk, with dissimilar warp and weft colors, a model for shadowing/masking is required. The noise we used is ad hoc, and a better model for cloth irregularities along with a model for visible fibers will improve the appearance of the renderings. Since most fabrics are thin and transmit a significant amount of light, a model for translucency is needed to better predict their appearance from the side opposite the light, and to properly compute global illumination due to light transmitted through the cloth. Our model is also not suitable for pile fabrics (e.g., velvet and corduroy), knitted fabrics, and nonwoven fabrics.

In this article we used the model to fit measurements of our example fabrics, which served our purpose of evaluating the model by comparing to photographs. However, measurements are not required to use the model; since the parameters are in terms of the construction of the cloth, it is quite feasible to simply type in a description of some new cloth, existing or not, in order to simulate its appearance. An implementation of the model, including a domain-specific language for describing weave patterns and rendering parameters, is available as part of the open-source *Mitsuba* rendering system [Jakob 2011].

Finally, the BTF comparison shows that some aspects of appearance are well captured by BTFs while others are better captured by the analytic model; this suggests that the best possible cloth model would combine the two approaches, using an analytic model to provide accurate highlight structure and photographic data to provide rich texture and irregularity.

ACKNOWLEDGMENTS

The authors warmly thank Manolis Savva, whose BTF capture and rendering system was used to make the comparison in Section 8.3; he also captured the BTF data for that comparison. His work was instrumental in bringing this article to press. We also thank Wenzel Jakob for implementing our model in his *Mitsuba* renderer, which increases its availability to the graphics community.

REFERENCES

ADABALA, N., MAGNENAT-THALMANN, N., AND FEI, G. 2003. Visualization of woven cloth. In *Proceedings of the Eurographics Workshop on Rendering*. 178–185.

ARFKEN, G. B., WEBER, H. J., AND WEBER, H.-J. 1995. In *Mathematical Methods for Physicists*, 4th Ed. Academic Press, Chapter 11.5.

ASHIKHMIN, M., PREMOŽE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Proceedings of the SIGGRAPH Conference*. 65–74.

BUCK, G. S. AND MCCORD, F. A. 1949. Luster and cotton. *Textile Res. J. 19,* 11, 715–754.

DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph. 18,* 1, 1–34.

DAUBERT, K., LENSCH, H. P. A., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Efficient cloth modeling and rendering. In *Proceedings of the Eurographics Workshop on Rendering*. 63–70.

DONG, Y., WANG, J., TONG, X., SNYDER, J., LAN, Y., BEN-EZRA, M., AND GUO, B. 2010. Manifold bootstrapping for SVBRDF capture. *ACM Trans. Graph. 29,* 4.

DRAGO, F. AND CHIBA, N. 2004. Painting canvas synthesis. *Vis. Comput. 20,* 5, 314–328.

DUTRÉ, P., BEKAERT, P., AND BALA, K. 2003. *Advanced Global Illumination*. A K Peters, Natick, MA.

EVANS, M., HASTINGS, N., AND PEACOCK, B. 2000. *Statistical Distributions*, 3rd Ed. Wiley-Interscience, New York, Chapter 41.

HANRAHAN, P. AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the SIGGRAPH Conference*. 165–174.

HAVRAN, V., FILIP, J., AND MYSZKOWSKI, K. 2010. Bidirectional texture function compression based on multi-level vector quantization. *Comp. Graph. Forum 29,* 1, 175–190.

HUNTER, R. S. AND HERALD, R. W. 1987. *The Measurement of Appearance*, 2nd Ed. Wiley-Interscience, New York.

IRAWAN, P. 2007. Appearance of woven cloth. Ph.D. thesis, Cornell University.

IRAWAN, P. AND MARSCHNER, S. R. 2006. A simple, accurate texture model for woven cotton cloth. Tech. rep. PCG-06-01, Cornell University.

JAKOB, W. 2011. Mitsuba physically based renderer. mitsuba-renderer.org.

KAJIYA, J. T. AND KAY, T. L. 1989. Rendering fur with three dimensional textures. In *Proceedings of the SIGGRAPH Conference*. 271–280.

KAUTZ, J., BOULOS, S., AND DURAND, F. 2007. Interactive editing and modeling of bidirectional texture functions. *ACM Trans. Graph. 26,* 3.

LU, R., KOENDERINK, J. J., AND KAPPERS, A. M. L. 1998. Optical properties (bidirectional reflection distribution functions) of velvet. *Appl. Opt. 37,* 25, 5974–5984.

LU, R., KOENDERINK, J. J., AND KAPPERS, A. M. L. 2000. Specularities on surfaces with tangential hairs or grooves. *Comput. Vis. Image Underst. 78,* 3, 320–335.

MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. *ACM Trans. Graph. 22,* 3, 780–791.

MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Trans. Graph. 24,* 3, 727–734.

MCALLISTER, D. K., LASTRA, A., AND HEIDRICH, W. 2002. Efficient rendering of spatial bi-directional reflectance distribution functions. In *Proceedings of the Conference on Graphics Hardware*. 79–88.

MÜLLER, G., BENDELS, G. H., AND KLEIN, R. 2005. Rapid synchronous acquisition of geometry and BTF for cultural heritage artefacts. In *Proceedings of the Symposium on Virtual Reality, Archaeology and Cultural Heritage*. 13–20.

MÜLLER, G., MESETH, J., SATTLER, M., SARLETTE, R., AND KLEIN, R. 2004. Acquisition, synthesis and rendering of bidirectional texture functions. In *Eurographics 2004, State of the Art Reports*. 69–94.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*. 117–126.

NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. Geometric considerations and nomenclature for reflectance. Monograph 161, National Bureau of Standards (US).

PARKER, J. 1993. *All About Cotton: A Fabric Dictionary & Swatchbook*. Rain City Pub.

PELLACINI, F. AND LAWRENCE, J. 2007. AppWand: editing measured materials using appearance-driven optimization. *ACM Trans. Graph. 26,* 3, 54.

PONT, S. C. AND KOENDERINK, J. J. 2003. Split off-specular reflection and surface scattering from woven materials. *Appl. Opt. 42*, 1526–1533.

SATTLER, M., SARLETTE, R., AND KLEIN, R. 2003. Efficient and realistic visualization of cloth. In *Proceedings of the Eurographics Workshop on Rendering*. 167–177.

SIRIKASEMLERT, A. AND TAO, X. 1999. Effects of fabric parameters on specular reflection of single-jersey knitted fabrics. *Textile Res. J. 69,* 9, 663–675.

TAO, X. AND SIRIKASEMLERT, A. 1999. A three-dimensional analysis of specular reflection from single-jersey knitted fabrics. *Textile Res. J. 69,* 1, 43–51.

VOLEVICH, V. L., KOPYLOV, E. A., KHODULEV, A. B., AND KARPENKO, O. A. 1997. An approach to cloth synthesis and visualization. In *Proceedings of the GRAPHICON Conference*.

WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph. 27,* 3, 41:1–41:9.

WELFORD, T. 1967. *The Textiles Student's Manual*. Pitman, London.

WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Proceedings of the SIGGRAPH Conference*. 255–264.

WU, H., DORSEY, J., AND RUSHMEIER, H. 2011. A sparse parametric mixture model for BTF compression, editing and rendering. *Comput. Graph. Forum 30,* 2, 465–473.

XU, Y.-Q., CHEN, Y., LIN, S., ZHONG, H., WU, E., GUO, B., AND SHUM, H.-Y. 2001. Photorealistic rendering of knitwear using the Lumislice. In *Proceedings of the SIGGRAPH Conference*. 391–398.

YASUDA, T., YOKOI, S., ICHIRO TORIWAKI, J., AND INAGAKI, K. 1992. A shading model for cloth objects. *IEEE Comput. Graph. Appl. 12,* 6, 15–24.

ZINKE, A. AND WEBER, A. 2007. Light scattering from filaments. *IEEE Trans. Vis. Comp. Graph. 13,* 2, 342–356.

Received February 2011; revised June 2011; accepted August 2011