

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

A measurement-based appearance model for fabrics

Permalink

<https://escholarship.org/uc/item/1n4501jb>

Author

Bisker, Oleg

Publication Date

2012

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

A Measurement-Based Appearance Model for Fabrics

A thesis submitted in partial satisfaction of the
requirements for the degree
Master of Science

in

Computer Science

by

Oleg Bisker

Committee in charge:

Henrik Wann Jensen, Chair
David Kriegman
Jurgen Schulze

2011

Copyright
Oleg Bisker, 2011
All rights reserved.

The thesis of Oleg Bisker is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2011

DEDICATION

To my parents, for seeing greater things for me than I could see for myself.

TABLE OF CONTENTS

Signature Page		iii
Dedication		iv
Table of Contents		v
List of Figures		vii
List of Tables		x
Acknowledgements		xi
Abstract of the Thesis		xii
Chapter 1	Introduction	1
	1.1 History	2
	1.2 Problem statement	3
	1.3 Main Contributions	3
	1.4 Thesis Organization	4
Chapter 2	Properties of Threads and Fabric	5
	2.1 Textile Thread Characteristics	5
	2.2 Weave Patterns	6
	2.3 Microscope Observations	6
	2.3.1 Linen Plain	7
	2.3.2 Silk Crepe de Chine	8
	2.3.3 Polyester Satin Charmeuse	8
	2.3.4 Wool Flannel	9
Chapter 3	Scattering Measurement	11
	3.1 BRDF	11
	3.2 Inverse Reflectometry	12
	3.2.1 Setup	13
	3.2.2 CCD Sensor	16
	3.2.3 Color	17
	3.2.4 Noise	18
	3.2.5 HDR	19
	3.2.6 Light Calibration	21
	3.2.7 BSDF Processing	22
	3.3 Results	24
	3.3.1 Incidence Plane BRDF	24
	3.3.2 Hemisphere BSDF	27

	3.3.3 Scattering Features	29
Chapter 4	Light Scattering From Threads	33
	4.1 Surface Scattering	33
	4.2 Smooth Cylinder Model	34
	4.3 A Light Scattering Model for Thread	36
	4.3.1 Model Validation	39
Chapter 5	Modeling the Appearance of Cloth	41
	5.1 Previous Approaches	41
	5.2 Weave Pattern Representation	42
	5.2.1 Thread Curve Definition	43
	5.2.2 Weave Definition	43
	5.2.3 Local Shadowing	44
	5.3 Image Synthesis	46
	5.4 Results Validation	48
Chapter 6	Conclusion	50
	6.1 Future Work	52
Appendix A	Additional Renders	53
Bibliography	56

LIST OF FIGURES

Figure 2.1:	Common weave patterns (left) plain, (middle) satin, (right) twill.	7
Figure 2.2:	Microscope pictures of the <i>linen plain</i> fabric. The fabric (left) is made out one type of thread with varying thickness.	7
Figure 2.3:	Microscope pictures of the <i>silk crepe de Chine</i> fabric. The fabric (left) is made out of two different types of thread: a set of dense spun threads (middle) and a series of filament shiny threads (right) which go above and below of the the other threads. . . .	8
Figure 2.4:	Microscope pictures of the <i>polyester satin charmeuse</i> fabric. The fabric (left) is made out of two different types of thread: A set of dense spun threads (middle) and a series of filament shiny threads (right) which go above and below of the the other threads.	9
Figure 2.5:	Microscope images of the <i>wool flannel</i> fabric. The fabric (left) is made out of one type of thread (spun). The right image shows a closeup of the thread structure where individual wool fibers are visible and their translucency is apparent.	10
Figure 3.1:	An image showing the thread BRDF measurement setup using the UCSD Hemispherical Gantry. It shows a vertically suspended thread sample being illuminated by a light source. . . .	13
Figure 3.2:	Schematic of BRDF measurement parameter space. The cylinder in the middle represents the thread sample, ω_i is the light source direction, and ω_o is the camera direction.	15
Figure 3.3:	A schematic diagram relating radiometric quantities to a physical CCD sensor.	17
Figure 3.4:	A simplified diagram of how real world lighting is encoded by a CCD imaging device.	17
Figure 3.5:	Diagram showing a color filter mounted on top of CCD sensor elements (left). Side by side comparison of a raw Bayer CFA image and its reconstruction (right).	18
Figure 3.6:	Low reflectance image dominated by noise(left), noise averaging result from two captures of the same image (middle), and joint denoising and demosaicing (right).	19
Figure 3.7:	The response curve plot for our SCOR-20SD CCD camera (left). The blue highlight outlines what a linear response should look like. (right) HDR capture samples showing a thread imaged at 4 exposure levels as well as the combined HDR result.	20
Figure 3.8:	Thread BSDF measurement notation. Longitudinal angles θ are computed with respect to the normal plane and the azimuthal angles ϕ are computed based on the local surface normal direction n	23

Figure 3.9:	HDR images of an in-plane acquisition sweep (top). Recovered reflectance values h for each of the HDR images (bottom). . . .	23
Figure 3.10:	BRDF measurements of rayon thread, plotted for a 15° incidence angle as a function of exitant angles (top). Also shown are the images that were used to produce the plot as well as their reflectance contribution(bottom).	25
Figure 3.11:	Polar plots of measured incidence plane BRDF for two different threads. The left plot contains the RGB average BRDF for a filament silk thread, and the right plot a filament polyester thread.	26
Figure 3.12:	Polar plots of measured incidence plane BRDF for two different threads. The left plot contains the RGB average BRDF for a staple linen thread, and the right plot a staple polyester thread.	27
Figure 3.13:	Hemisphere plots of the BSDF as a function of the scattering direction when the position of the light is held fixed at $\theta_i = 30^\circ$ incidence.	28
Figure 3.14:	BRDF visualization sets of backward scattering hemisphere (top row) and forward scattering hemisphere (bottom row). The labels on the left correspond to the yarn type and the labels on the right correspond to fiber type.	30
Figure 3.15:	Schematic relating a bright band in a hemisphere plot to a cone shaped reflectance from a cylinder.	32
Figure 4.1:	Angular notation for cylinder based models.	35
Figure 4.2:	Incidence plane thread BRDF measurements in the top row matched by thread model in the bottom row. The thread sample is twisted and was extracted from the <i>polyester satin charmeuse</i> cloth. The plots show scattering as a function of view angle. . .	39
Figure 4.3:	Incidence plane thread BRDF measurements in the top row matched by thread model in the bottom row. The thread sample is filament type (untwisted) and was extracted from the <i>polyester satin charmeuse</i> cloth. The plots show scattering as a function of view angle.	40
Figure 5.1:	The weaving pattern and a sample tangent curve for the <i>polyester satin charmeuse</i> fabric: (left top) the weaving pattern, (left bottom) a smallest patch, (right) the tangent curve for the two types of threads.	43
Figure 5.2:	Weave pattern encoding pseudocode (left), and the resultant patterns of satin (middle) and plain (right).	44

Figure 5.3:	Approximate thread-to-thread shadowing is shown (middle), compared to no shadowing (left), and the shadow contribution mask (right). The light is hitting the fabric at a glancing angle and causing higher lying threads to block their lower lying neighbors.	45
Figure 5.4:	Path traced images of physically based cloth shading. Each image synthesizes a real world cloth sample. Both the thread light-scattering properties and the exact weave structure are reproduced in our models, yielding a high fidelity simulation of the appearance of the real cloth samples. The fabrics pictured are: (a) <i>polyester satin charmeuse</i> , (b) <i>silk crepe de Chine</i> , (c) <i>shot fabric</i> , (d) <i>velvet</i> , (e) <i>linen</i> , and (f) <i>wool</i>	47
Figure 5.5:	Result validation setup, show the direction in which the cloth samples are wrapped.	48
Figure 5.6:	Comparison of reference photographs and the rendered results for (top) <i>silk crepe de Chine</i> , (middle) <i>polyester satin charmeuse</i> , and (bottom) <i>plain linen</i> , . The three cylinders for each fabric show three orientations as illustrated in Figure 5.5.	49
Figure A.1:	Closeup rendered image demonstrating the procedural texture found in fabrics. The texture arises from the weave pattern as well as the local thread shadowing.	53
Figure A.2:	Extreme closeup render (top) showing limitations of our approach where each thread is uniformly colored and lacking detail from the fibers that make it up. The second image (bottom) shows the moire pattern that arises due to the high frequency texture detail found in fabrics.	54
Figure A.3:	Clothing rendered using our appearance model. The precise control of thread parameters allows us to define multi-color fabrics (right).	55

LIST OF TABLES

Table 3.1:	A listing of the threads measured, organized by type and fabric that they were extracted from.	24
Table 3.2:	Scattering characteristics of thread samples based on data in Figure 3.14. Next to each thread type, (s) denotes spun, and (f) denotes filament. The first word in each entry identifies the magnitude of the scattering phenomenon (high-med-low), and the second word describes its scattering lobe shape (shiny-glossy-blurry). Plots in each column of the figure correspond to a light incidence angle θ_i as labeled at the very top.	29
Table 4.1:	Description of important symbols.	36

ACKNOWLEDGEMENTS

I would like to thank my advisor Dr. Henrik Wann Jensen for his encouragement to pursue computer graphics, his infectious curiosity for understanding visual phenomena, and his continuous support of my research goals. I would also like to thank my master's committee members Dr. David Kriegman and Dr. Jurgen Schulze who agreed to review this thesis.

I would especially like to thank my numerous colleagues in the computer graphics lab who help make it such a fruitful collaborative environment. In particular, I would like to thank Iman Sadeghi for all the hard work and collaboration that led to coauthoring a conference paper which is currently in submission. I am very grateful to Toshiya Hachisuka, for giving me feedback on my many half-baked ideas, and allowing me to tap into his encyclopedic knowledge of computer graphics. I would also like to thank Krystle De Mesa, Carlos Dominguez, and Daniel Tenedorio who at various stages contributed their expertise to my projects.

Outside of my work, I would like to thank my whole family and girlfriend, who encouraged and supported me throughout my academic and professional career. Last but not least, I am very thankful to my friends for making my time in San Diego so amazing and unforgettable.

Chapters 2, 3, and 4, contain material that has been submitted for publication. Sadeghi, Iman; Bisker, Oleg; De Dekken, Joachim; Wann Jensen, Henrik. The thesis author was the coauthor of this material.

ABSTRACT OF THE THESIS

A Measurement-Based Appearance Model for Fabrics

by

Oleg Bisker

Master of Science in Computer Science

University of California, San Diego, 2011

Henrik Wann Jensen, Chair

Accurate representations of real-world materials are a crucial prerequisite for realistic image synthesis. This thesis presents a method for rendering fabric based on a low-level simulation of light scattering from threads. Making use of reflectance acquisition techniques, we study the reflectance characteristics of a variety of threads extracted from real-world fabrics.

We introduce a model for light scattering from threads which simulates both the surface and subsurface transport of light. The model's accuracy is assessed against physical measurements. In order to generate an image of fabric, we propose a weave pattern and thread curve encoding that enables us to mimic the structure of a wide variety of fabrics. We apply this appearance model in a path tracing framework to realistically reproduce the appearance of fabrics.

Chapter 1

Introduction




Photo-realistic rendering of complex objects can be achieved at near interactive rates on modern commodity hardware. This has led to a large increase in the production and demand for computer generated (CG) imagery in consumer markets. This trend can be observed in the yearly CG production output of at least ten animated feature films, dozens of video games, and thousands of smaller visual effects achieved with computer graphics. As a result of this increased production output, there is a rising demand for more complex and more realistic visualization of objects and materials.

This thesis focuses on one such complex material, woven cloth, which exhibits a broad range of intricate visual phenomena and is at the same time pervasive in every day life. Despite its ubiquitous nature, cloth remains a challenging

problem for computer visualization. In large part, this is due to the sheer variety of available fabrics. They are manufactured from many raw materials and structured in multiple weaving patterns. As a result, there exists no comprehensive appearance model which can capture the look of a range of fabrics.

1.1 History

When humans first began wearing clothes, their garments were made from natural elements such as animal skin, furs, grasses and leaves. Clothing was an essential component of warmth, protection, and ultimately survival. From these early beginnings we fast forward many centuries of human progress to arrive at roughly 6000 B.C. where a weaving process first appears. By adapting basket weaving techniques, woven cloth was born and over the subsequent centuries emerged as one of humankind's most fundamental technologies.

Research indicates that cultures on every continent devised crude looms and methods for creating cloth. Despite their isolation from each other, there was great similarity in the looms constructed and in the weavings produced. We can trace the evolution and differentiation of clothing by looking at two iconic historical periods. In the classical age, the Greeks and Romans predominantly wore woven wool and linen clothes with an earthen color palette ranging from off-white to dark red and indigo. A simple plain-weave pattern, an over-under stitch, was used for the construction of their cloth.

By the end of the middle ages and into the height of the Renaissance, garments saw a huge explosion in visual intricacy thanks to embroidery, advances in bleaching and dyeing methods, and the widespread availability of silk and cotton as raw materials. Additionally, satin and twill weaving patterns became more common, yielding more fanciful and lustrous fabrics. In modern times, we have extended the variety of fabrics by introducing synthetic fibers and highly precise industrial fabrication methods. Clothing now occupies a position close in importance to food. Its value exceeds the nominal notions of utility and comfort. The differentiation of clothing has come to uniquely define cultures, religions, societal

stature, and the ever changing tides of fashion.

1.2 Problem statement

Synthesizing realistic images of cloth requires a geometric representation as well as an accurate description of how light interacts with the cloth surface. It is often sufficient to define the geometric model as a smooth polygonal shape. Alternatively, drape simulation can procedurally produce a realistically shaped surface.

Once we have chosen a geometric surface, our next challenge is to define a method for how it interacts with light. To achieve this goal we need to measure and characterize how cloth reflects light in the real world. A common technique for measuring reflectance uses a gonireflectometer under controlled conditions of position and orientation. The measured reflectance data then needs to be processed and encoded into a specific format usable in rendering. Finally, in order to be truly useful in a practical rendering setting, the measured results need to be explained in terms of light transport and encoded into either a physically based or a phenomenological approximation model.

1.3 Main Contributions

This thesis covers a range of challenges in studying the appearance and visualization of cloth. The primary contributions are summarized as follows:

- A measurement setup designed to capture directionally dependent light scattering, and a description of the hardware used.
- A BSDF processing pipeline that takes gigabytes of raw images as input, and computes accurate radiometric quantities needed to characterize light scattering by a material.
- An extension of the smooth dielectric cylinder model to approximate the scattering behavior of threads.

- A rendering algorithm which relies on independent descriptions of threads and weaving pattern to synthesize the appearance of cloth.

We first present a complete reflectance measurement system, which we use to acquire directionally dependent light scattering from single thread samples. We then leverage our measured data to identify important light scattering characteristics. Next, we present an empirical model for light scattering from different types of threads. We combine this thread-scale reflectance with a fabric weave pattern description to produce a model and explanation of how cloth scatters light.

1.4 Thesis Organization

The subsequent chapters are organized as follows:

- Chapter 2 gives a brief overview of several fabrics that we studied and their salient appearance characteristics.
- Chapter 3 starts by summarizing radiometric concepts that are relevant to measurement and ray tracing. Next, we describe the physical apparatus and sensors involved in our reflectance acquisition. We then outline a framework for processing and extracting the BRDF/BSDF. We conclude the chapter by presenting our measurements for a variety of threads.
- Chapter 4 discusses concepts relevant to light scattering by surfaces. It then motivates a novel scattering model, by giving an overview of previous work on smooth dielectric cylinders. Finally, we present an extension of the smooth cylinder model to approximate the scattering behavior of threads.
- Chapter 5 describes a rendering algorithm for synthesizing the appearance of cloth based on a procedurally generated weave pattern and the thread model presented in chapter 4.

Chapter 2

Properties of Threads and Fabric

Certain fabrics exhibit vibrant color, sheen and luster. These characteristics arise from the threads that make up the fabric as well as the pattern in which they are woven. The textile fabrication process consists of the following steps [Sho24]:

1. Formation of thread from an irregular mass of fibers.
2. Weaving or knitting of thread into a fabric.
3. Finishing, i.e. converting crude textile into saleable article.

The weave pattern and thread characteristics are largely orthogonal subjects which we address independently.

2.1 Textile Thread Characteristics

Threads can be made from a variety of natural and synthetic fibers. Linen and cotton fibers are coarse and stiff, resulting in strong threads that can be used for everything from sewing seams to making sails. Wool fibers are similar to human hair in that they have a scaly shiny surface. Silk is a super-fine fiber which can be made into threads of long continuous filament or spun together. In this thesis, we focus on this distinction of spun vs. filament threads as an important characteristic that determines how they scatter light.

Synthetic threads, such as those made of polyester, are very versatile, appearing in filament or spun form. They have a variety of cross-sections as well as smooth or bumpy surfaces. Additionally, a thread can contain several fiber types, resulting in an overall mixture of desirable properties that are superior to any one fiber type. One such popular mixture is a polyester core with cotton fibers spun on the outside.

An important property of threads that is relevant to light scattering is the shape. Discussion of shape can be simplified by separating the cross-sectional shape from the length-wise shape. The former may be circular, ellipsoidal, or race-track shaped, varying widely and depending on the weave pattern, the tightness of construction, and the finishing treatment [LA98]. Regardless of the cross-section, the length-wise shape can be identified by a center line through the thread that is continuously differentiable and forms a spatial curve. With this representation we can reconstruct a 3D thread shape by placing an orthogonal plane at a point on the spatial curve and projecting the thread's cross-sectional shape onto it.

2.2 Weave Patterns

The weave pattern of a fabric refers to how warp and weft yarns are interlaced. Warp and weft refer to orthogonal directions of thread, where the warp is usually the thread direction held in tension by a loom, and the weft is the thread inserted over-and-under the warp threads. Common weave patterns are plain, satin, and twill (see Figure 2.1). Plain weave repeats an over under thread path for both the warp and the weft. Satin is asymmetrical in that one warp thread passes over several weft threads, then under one weft thread. Twill weaves have multiple warp threads passing over multiple weft threads in a repeating pattern.

2.3 Microscope Observations

We obtained four fabric samples to study their light scattering properties in more detail. To understand the behavior of each fabric, we first investigated




Figure 2.1: Common weave patterns (left) plain, (middle) satin, (right) twill.

their structure under a microscope.

2.3.1 Linen Plain

The *linen* fabric with plain weaving pattern is made out of dense spun threads which exhibit a considerable amount of absorption (see Figure 2.2). In the weaving pattern the warp and weft threads have the same curve structure and a randomly varying thickness. The identical nature of the thread curves is due to the symmetry of the plain weaving pattern.




Figure 2.2: Microscope pictures of the *linen plain* fabric. The fabric (left) is made out one type of thread with varying thickness.

2.3.2 Silk Crepe de Chine

The *silk* fabric with the *crepe de Chine* weaving pattern is made out of two different threads (see Figure 2.3). One is a dense spun thread that is highly light absorbing. It is mainly responsible for the color of the fabric. In the weave pattern these threads do not change direction, but stay straight and uniformly spaced from each other. The second thread type is sparse and flat. This thread exhibits sharp surface reflection and very little absorption resulting in its translucent appearance. While capturing microscope images of the thread, we observed that two separate sections of the thread would get very bright as we moved the light source. This double specular reflection is due to the local slope of thread samples, (see Figure 2.3 right).




Figure 2.3: Microscope pictures of the *silk crepe de Chine* fabric. The fabric (left) is made out of two different types of thread: a set of dense spun threads (middle) and a series of filament shiny threads (right) which go above and below of the the other threads.

2.3.3 Polyester Satin Charmeuse

The *satin charmeuse* is made of polyester threads. Similar to the *silk crepe de Chine*, this fabric is also made out of two distinct threads (see Figure 2.4). The first type is a spun dense thread which determines the color of the fabric. These threads stay straight in one direction. The second type is a thin, transparent, filament thread with a sharp surface reflection.




Figure 2.4: Microscope pictures of the *polyester satin charmeuse* fabric. The fabric (left) is made out of two different types of thread: A set of dense spun threads (middle) and a series of filament shiny threads (right) which go above and below of the the other threads.

These threads go above and below the spun threads, but remain longer above than below. This asymmetry in the weaving pattern causes the fabric to have two different sides (see left-insert in Figure 2.4). By moving the light during our microscope observations, we noticed strong reflections in three different direction of light as can be seen in Figure 2.4 right.

2.3.4 Wool Flannel

This wool fabric is a plain weave though shares some characteristics with twill by having weft threads in pairs. The flannel fabric has been brushed, which is a mechanical process of forming loose fibers on the surface to increase "softness". In right frame of Figure 2.5, we can clearly identify the wool fibers that make up the spun wool thread. They appear as translucent cylinders with a shiny exterior.



Figure 2.5: Microscope images of the *wool flannel* fabric. The fabric (left) is made out of one type of thread (spun). The right image shows a closeup of the thread structure where individual wool fibers are visible and their translucency is apparent.

Chapter 3

Scattering Measurement

In this section we describe our BRDF measurement process. Threads pose unique challenges to the BRDF acquisition system due to their thin nature and high dynamic range light scattering properties. To address these issues we adapt previous approaches in BRDF measurement to suit our acquisition goals.

3.1 BRDF

To synthesize realistic images, a complete reflectance description is needed for all surfaces in a scene. A simplified form of this description was formalized as the bidirectional reflectance distribution function (BRDF) in [NRH⁺92]. This function contains wavelength-dependent information for scattering in any direction given any incoming direction over a hemisphere.

By assuming that all light arriving at a surface is either reflected or absorbed at the point of incidence x , the BRDF can be defined as:

$$f_r(x, \omega_i, \omega_r) = \frac{dL(x, \omega_r)}{dE(x, \omega_i)} \quad (3.1)$$

The BRDF is defined as the ratio of reflected radiance $L(x, \omega_r)$ from point x in the direction ω_r , to the irradiance $E(x, \omega_i)$ arriving at the same point x from direction ω_i . As mentioned earlier, this function is defined uniquely per wavelength, but for purposes of clarity we have omitted the wavelength parameter.

We can express the differential irradiance in terms of radiance as $dE(x, \omega_i) = L_i(x, \omega_i) \cos(\theta_i) d\omega_i$. Here, $d\omega$ is an area patch around direction ω defined as $d\omega = \sin(\theta) d\theta d\phi$. This allows us to evaluate the BRDF purely in terms of radiance.

$$f_r(x, \omega_i, \omega_r) = \frac{dL(x, \omega_r)}{dE(x, \omega_i)} = \frac{dL(x, \omega_r)}{L_i(x, \omega_i) \cos(\theta_i) d\omega_i} \quad (3.2)$$

For a BRDF to be physically plausible it must also fulfill the following two constraints: Helmholtz reciprocity, and energy conservation. These constraints imply that the view and light directions are reversible as well as that real materials can not propagate more light than they receive.

By assuming a spatially homogeneous BRDF, f_r is a four dimensional function with units of $[sr^{-1}]$. We can also extend the definition of a BRDF to describe scattering more generally. In fact, since a thread strand does not represent a surface patch per se, a bidirection scattering distribution function (BSDF) is what we really want to measure. This function observes the same exact constraints as a BRDF, except that it is defined over a sphere as opposed to a hemisphere.

3.2 Inverse Reflectometry

The goal of inverse reflectometry is to measure the BRDF/BSDF f_r described in the previous section. We will refer to BRDFs and BSDFs interchangeably throughout our discussion, with the only difference being a back-hemisphere function or a full-sphere function with respect to and incident light direction. To acquire f_r , we need to measure the reflectance as well as know the surface normal and information about the lighting. For purposes of this thesis, we focus our discussion on measuring yarn threads, however the techniques we use generalize to measuring other materials. In this section we describe the measurement process using a hemispherical gantry and a unique thread suspension apparatus. The BRDF results serve to validate an analytical model presented in the next chapter, as well as provide a qualitative basis for reasoning about threads and cloth in general.

3.2.1 Setup

Traditionally, the BRDF is measured with a gonioreflectometer, positioning an incident light source and a detector with respect to a flat sample. Since threads are more cylindrical than flat, we utilize a purpose-built setup that differs from conventional setups (see Figure 3.1).



Figure 3.1: An image showing the thread BRDF measurement setup using the UCSD Hemispherical Gantry. It shows a vertically suspended thread sample being illuminated by a light source.

Our system handles the four angular degrees of the BRDF as well as the spectral dimension through the use of a color encoding CCD chip. The idea is to vary the position of the light and the camera while taking photos at each position combination. The pixel values in the photos are proportional to directional flux. The main drawback of this kind of system is that it takes several hours to measure even a sparse sampling of a BRDF. A faster BRDF acquisition process was demonstrated by [War92] by capturing multiple BRDF samples simultaneously through the use of a fisheye lens and a hemispherical mirror. Since then, BRDF capture

systems by [DvGNK99, MWL⁺99] demonstrated approaches that exploit BRDF reciprocity as well as imaging non-planar patches to obtain multiple BRDF samples. Despite these improvements, we opted to only take advantage of reciprocity, retaining simplicity in both the physical apparatus and the BRDF acquisition.

We employed the UCSD Hemispherical Gantry to conduct our measurements. The system is composed of two robotic arms, one that holds a light source and the other a detector. The light source is connected by a fiber optic cable to a Dolan-Jenner DC950H Machine Vision Illuminator that uses a 150-watt quartz halogen bulb. The detector is a Point Grey Scorpion video camera SCOR-20SD, with a 12-bit, 1/1.8" Sony CCD sensor with a maximum resolution of 1628x1236 at 15 FPS. The mechanics and design of the gantry maintain a constant distance between the light/detector and the center where the measurement sample is placed. The distance from the center to the light is 66cm and from the center to the detector is 103cm. The gantry arms are controlled through software allowing them to be placed with an angular resolution of 0.01° . In our measurements we illuminate an 8 cm section of thread with a collimated light beam and collect radiance scattering measurements with the CCD camera. Our measurement setup can be seen in Figure 3.1.

Sample Preparation

To procure a thread sample, we first remove a single strand from a finished fabric. When a thread is removed from fabric it is no longer straight, but retains the shape that it had in the fabric. In order to obtain accurate scattering measurements, the thread must be extended to its maximal length. This type of procedure is common in fabric quality testing and requires standard tension, which has the general goal of non-destructively pulling on one end of the thread. In our experimental thread mount, we clamp one end of the thread to a poseable arm, and let the rest of it hang, weighed down by a magnetic set of spheres at the unclamped end. Hanging the thread in mid-air allows the gantry to measure a full 4D BSDF with minimal occlusions and minimal background to contaminate the measurements. Additionally, gravity provides a vertically straight thread orientation

which eliminates pose calibration.

BSDF Sampling

To facilitate discussion of measuring the BSDF, we establish a reference frame around the thread direction (see Figure 3.2). Here, t is the thread tangent direction, θ is a longitudinal angle, and ϕ is an azimuthal angle. Measuring one point sample of the BSDF involves moving the light to position (θ_i, ϕ_i) and moving the camera to (θ_o, ϕ_o) .




Figure 3.2: Schematic of BRDF measurement parameter space. The cylinder in the middle represents the thread sample, ω_i is the light source direction, and ω_o is the camera direction.

To acquire a sampling of a 4D BSDF at a coarse angular resolution of 10

degrees requires $36^4 = 1.68$ million photographs. We reduce this acquisition burden by making the assumption that the thread BSDF is largely invariant to rotation around its primary axis t . This results in collapsing distinct ϕ_i and ϕ_o angles into a single angle $\phi_d = \phi_i - \phi_o$. Additionally, due to Helmholtz reciprocity we can avoid performing redundant measurements, thus reducing our acquisition to a manageable $\frac{36^3}{2} = 0.02$ million photographs. In practice, we take multiple images per sample point as well as perform a non-symmetric sampling with respect to the angular dimensions. The reasons for this will become clear from discussions in the subsequent sections.

3.2.2 CCD Sensor

The most important component of reflectance measurement systems is the detector. Prior to solving the inverse reflectance problem, we need to have a thorough characterization of the detector in order to obtain real world light quantities from digital data. As mentioned in the previous section, we use a CCD based camera to detect reflectance. The CCD chip inside the camera contains a two dimensional array of two million sensor elements that convert arriving irradiance into digital values. While the shutter is open, light strikes these sensor elements causing them to release and store charge, where the amount of charge stored is proportional to the irradiance at the sensor element. In ray tracing terms, the sensor element integrates radiance over the solid angle subtended by the lens aperture $d\omega_i$. This radiance integral is equal to the irradiance multiplied by the sensor element's area or more precisely the total flux incident at the sensor element (see Figure 3.3).

At the end of the exposure, this stored charge is passed to an analog to digital converter (ADC) and digitized. The final step maps the linear quantized intensities to pixel values using the sRGB standard. Thus the pixel values stored in an image are not linearly related to the light intensity in a scene. A block diagram showing the path from real world into pixel world is show in Figure 3.4.

An important physical limitation of a CCD chip is its *dynamic range*, i.e. the ratio between the maximum and minimum amount of charge that it can encode.




Figure 3.3: A schematic diagram relating radiometric quantities to a physical CCD sensor.




Figure 3.4: A simplified diagram of how real world lighting is encoded by a CCD imaging device.

For example our 12-bit sensor can capture on the order of 4,096 electrons of charge before it saturates. When a sensor element saturates, blooming artifacts may occur, which result in incorrect (brightened) results for neighboring elements.

3.2.3 Color

Color acquisition in cameras is something that we take almost for granted, but it too has to be dealt with carefully. The digital images that the CCD sensor produces are effectively single channel, since each sensor element can only return a single value. The most common solution to this is to cover the sensor elements with a color filter array (CFA) arranged in a Bayer [Bay76] pattern (see Figure 3.5). This approach trades spatial resolution for color resolution, by effectively producing a sparser sampling of color than of image detail. A consequence of using a CFA is that a color reconstruction step (demosaicing) is required to produce the final color image (see Figure 3.5).

The lens system also plays an important role in capturing the reflectance from thin scene details such as threads. Lenses are used to focus light from the scene outside the camera onto the sensor elements of the CCD chip. As a result




Figure 3.5: Diagram showing a color filter mounted on top of CCD sensor elements (left). Side by side comparison of a raw Bayer CFA image and its reconstruction (right).

of the physical size of the sensor elements and the Bayer CFA, a scene detail can be correctly imaged only if it spreads over at least four sensor elements. If the scene detail is visible to less than four elements, then its color will be incorrectly reconstructed; worse yet, if the detail partially covers a single sensor element, then it may be imperceptible altogether. Our acquisition process faces these exact problems due to the poor sensor element coverage by the thin threads. We address this issue by using a higher magnification lens, a Cosmicalar C5028-M 50mm. Despite the higher magnification compared to a 35mm lens, raw images of threads contain only five pixels of coverage. By defocusing the lens, we are able to spread the radiance over more sensor elements and thus solve the color reconstruction error problem.

3.2.4 Noise

CCD sensors suffer from leakage current caused by chip defects and thermal variation. This results in a variety of noise patterns in the exposed image, with longer exposures producing more noise. Our acquisition process requires single photo exposure times ranging from 20 ms to 800 ms, which means that our images can be very noisy.

We first tackle this problem by taking multiple photos at a single exposure and averaging the results. This partially removes random noise which varies from

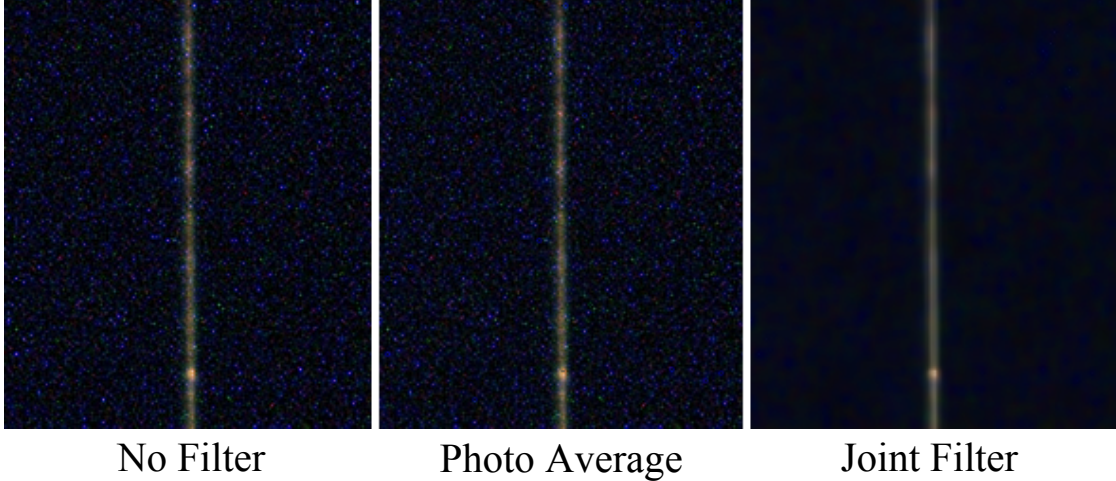


Figure 3.6: Low reflectance image dominated by noise(left), noise averaging result from two captures of the same image (middle), and joint denoising and demosaicing (right).

frame to frame. Next, we remove the remaining noise through a filtering process. We use the method of joint denoising and demosaicing introduced by [HP05] and generalized by [Con10] (see results in Figure 3.6). The reason for combining CFA reconstruction (demosaicing) with denoising is because demosaicing distorts the characteristics of the noise making subsequent denoising very complex. The approach we use generates a smooth chrominance image and a sharp luminance image from the raw image data. The noise is therefore isolated in the luminance image which is subsequently filtered to remove noise while maintaining accurate color reconstruction in the unaffected chrominance image.

3.2.5 HDR

As described in the previous section, our camera has a 12-bit dynamic range, i.e. it can encode 4,096 unique values per-pixel into the raw photo. The image formats available to us are 8 and 16 bits. For the practical purposes of using half the storage, we only use 8 bits in the raw photos. This limited dynamic range may result in dim areas all mapping to 0, or alternatively very bright areas all mapping to 255.

The idea behind high dynamic range (HDR) imaging [DM97, RBS99] is

to extend the limited dynamic range of a single photo by sequentially capturing multiple photos with different exposure times. Once the photos are captured, the goal is to combine them into a single high contrast image. Before that can be done the photos need to be converted to linear intensity so that numeric operations can be accurately performed on them.




Figure 3.7: The response curve plot for our SCOR-20SD CCD camera (left). The blue highlight outlines what a linear response should look like. (right) HDR capture samples showing a thread imaged at 4 exposure levels as well as the combined HDR result.

The camera’s response curve is a function which maps digital pixel values to linear intensity. We use the method of [RBS99], where we take multiple variable length exposures of a Macbeth chart to reconstruct a response curve I (see Figure 3.7 left). The recovered response curve for our vision-specific camera is mostly linear, becoming non-linear at pixel-values above 240. This type of response curve is rarely found in consumer grade cameras which have nonlinear, non-identical, response curves for the different color bands.

Having computed the response curve I , we compute the HDR linear intensity of a pixel p_j at image position j , by taking a weighted average of all the photos at position j according to:

$$p_j = \frac{\sum_e w_{ej} t_e I(y_{ej})}{\sum_e w_{ej} t_e^2} \quad (3.3)$$

Here, y_{ej} refers to a pixel in input photo at exposure e and position j . We pass this value into our response curve I to retrieve a linear intensity value equivalent. We sum these linear intensities weighting them by a Gaussian w_{ej} and

the exposure time t_e . The input photos for different exposure times and their combined result is shown in Figure 3.7 (right).

3.2.6 Light Calibration

To determine the intensity of our light source and therefore the irradiance at our thread sample, we use a calibration device called a spectralon. It is a translucent material made of compressed and sintered polytetrafluoroethylene powders [Lab98]. It is useful for BRDF calibration because of its high reflectivity, low absorption, and near-Lambertian behavior at near-normal illuminations. Our goal is to use a Labsphere Spectralon diffuse reflectance standard to measure and derive the irradiance from our light source.

The specifications of our reflectance calibration standard provide a bidirectional reflectance factor (BRF) that we can use to calibrate our measurements. A BRF is defined as the ratio of the flux scattered in a given direction by the sample, to that which would be scattered in that direction by the perfect reflecting diffuser, under identical conditions of illumination. At 8° incidence angle and normal viewing angle, the given BRF is $R = (0.989, 0.989, 0.988)$. For a single collimated light source, we can express the BRF as:

$$R = \frac{\pi L_{co}}{L_i \cos(\theta_i) d\omega_i} \quad (3.4)$$

and therefore incident radiance as:

$$L_i d\omega_i = \frac{\pi L_{co}}{R \cos(8^\circ)} \quad (3.5)$$

Here, L_{co} is the reflected radiance from the spectralon as captured in an image from normal incidence with the light positioned at 8° . We assume that the small patch of the spectralon that we are imaging is lit by uniform irradiance. We discard minor spatial variation and noise by averaging the pixel values in the spectralon image. The result, L_{co} , is a three component vector RGB representation of the light intensity relative to the spectralon. We can plug this into our definition of a BRDF to obtain:

$$f_r(x, \omega_i, \omega_o) = \frac{dL(x, \omega_o) R \cos(8^\circ)}{\pi L_{co} \cos(\theta_i)} \quad (3.6)$$

This final form of the BRDF can be intuitively explained as taking the linearized sample image, per-color channel multiplying it by (0.989,0.989,0.988), and per-color channel dividing by the linearized spectralon image (note that the solid angle $d\omega_i$ has conveniently dropped out). In the next section we give more detail on what constitutes an image.

3.2.7 BSDF Processing

We measure a complete 3D BSDF by varying the longitudinal angles θ_i , θ_r , and the azimuthal difference angle $\phi_d = \phi_i - \phi_r$. The diagram in Figure 3.8 shows the same notations as Figure 3.2 except from the perspective of the thread and in more detail.

As mentioned earlier, we did not measure a 4D BSDF because we assumed invariance of the BSDF with respect to individual ϕ angles and instead used a difference angle $\phi_d = \phi_i - \phi_r$. Since threads are not perfect cylinders, this assumption is somewhat violated, however, it allows us to capture less data while still observing the salient light scattering features.

Each HDR photo $H(\theta_i, \theta_r, \phi_d)$ contains an image of a thread on a black background cropped to a projected area (see Figure 3.9). The HDR images are computed using the methods in previous sections with special rejection heuristics for under and overexposed pixels. By thresholding the background in the images to zero, we can represent an HDR image as a single radiance value h :

$$h = \frac{1}{d\omega_r dA} \sum_j H_j \quad (3.7)$$

Here, H_j is an HDR pixel at position j and $d\omega_r$ is the solid angle subtended by the camera relative to the thread sample and dA is the sensor area. We are computing the differential quantity of radiance as an average over a finite interval of space and solid angle.




Figure 3.8: Thread BSDF measurement notation. Longitudinal angles θ are computed with respect to the normal plane and the azimuthal angles ϕ are computed based on the local surface normal direction n .



Figure 3.9: HDR images of an in-plane acquisition sweep (top). Recovered reflectance values h for each of the HDR images (bottom).

We acquired two types of reflectance measurements for seven different threads resulting in the following data:

- Detailed Slice: In-plane BRDF measurements for $(\phi_d = 0, \theta_r = [6 - 160])$ produces 5 exposures, each with 2 noise suppression photos at 2° steps, resulting in 780 captures (48.7 MB).
- Half BSDF: Hemisphere BSDF measurements for $(\phi_d = [0 - 180], \theta_r = [6 - 160], \theta_i = [6 - 90])$ produces 5 exposures, each with 2 noise suppression photos at 8° steps, resulting in 54,000 captures (3.29 GB).

The data processing can run concurrently with the acquisition, since processing is done one image at a time. The main processing steps to compute a single BRDF sample are:

1. Joint demosaic and denoise filter raw photos.
2. Combine filtered color photos into HDR images.
3. Sum the pixel values and normalize by solid angle and reflectance standard.

3.3 Results

We used the BRDF measurement approach described in the previous sections to acquire the scattering functions of seven different threads (see Table 3.1).

Table 3.1: A listing of the threads measured, organized by type and fabric that they were extracted from.

Fabric	Thread Type
light weight blue linen	spun
pink polyester satin charmeuse	spun filament
green silk crepe de Chine	spun filament
taupe rayon satin shantung	spun
tan wool flanel	spun

There are two types of measurements discussed: a detailed BRDF measurement to identify the fine-grained scattering shape, and a coarse BSDF measurement meant to characterize the overall scattering behavior. Important properties of both types of scattering measurements are analyzed for the purpose of building an analytical BSDF model in the next chapter.

3.3.1 Incidence Plane BRDF

The data collected by our acquisition process is plotted with three (RGB) color bands for a fixed incidence angle and a varying view angle in Figure 3.10. The plot shows the directions into which light scatters when a rayon thread sample, oriented along t , is illuminated at 15° incidence. The distance from the graph to the plot center denotes the magnitude of the BRDF. Note the log plot scale. Since the thread was not treated with any additional dyes and no polarizing filters were

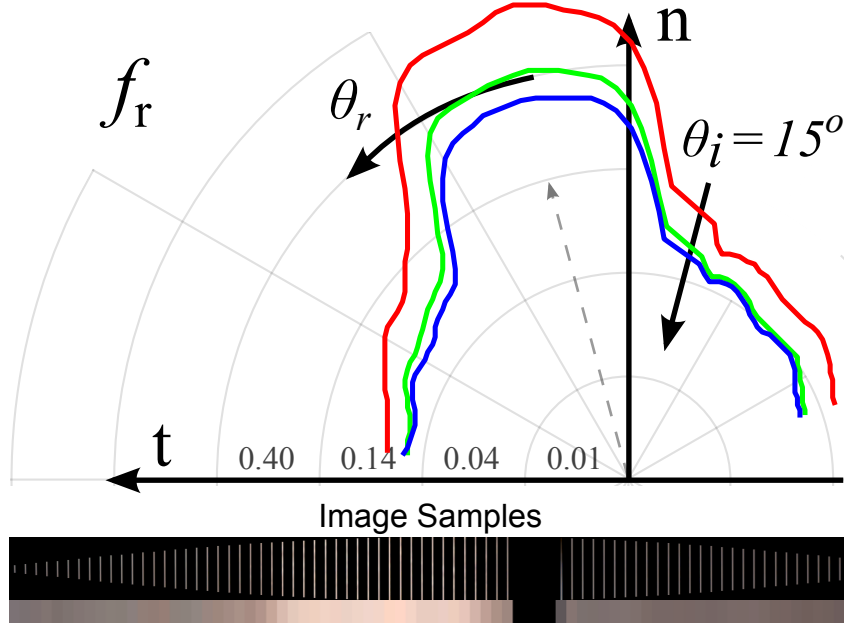


Figure 3.10: BRDF measurements of rayon thread, plotted for a 15° incidence angle as a function of exitant angles (top). Also shown are the images that were used to produce the plot as well as their reflectance contribution (bottom).

used on the camera, the BRDF plot represents the natural visible combination of surface reflection and internal scattering.

This particular thread shows a stronger red scattering component, which results in an overall pink hue to the thread and by extension the fabric. The dashed arrow in the plot shows the expected reflection direction for an ideal reflector. Ineed, we observe that there is a greater fraction of light scattering in this preferred direction than the rest of the hemisphere. Below the plot, the image samples that were used to generate the graph are shown, along with their extracted reflectance contribution.

To facilitate intuition about the plots in this section, we state the following: surface reflection results in a lobe in the specular reflection direction (ideal reflection) ($\theta_r = -\theta_i$), and internal scattering results in a wider lobe that is more decoupled from appearing in the specular direction. Applying these notions to Figure 3.10, we can observe that internal scattering, results in a diffuse BRDF of magnitude 0.04, and the surface reflection is added on top of it resulting in a combined magnitude of 0.4 around the specular direction. This semi-quantitative

analysis allows us to conclude that surface scattering is around 10x greater than internal scattering in the specular direction. Additionally, based on the shape of the plot, we can state that the width of the specular lobe is roughly 30° .

We present in-plane BRDF measurements for two different threads in Figure 3.11. As before, the BRDF is a function of two angles θ_i and θ_r , except now instead of three distinct color bands, each line graph represents the average intensity of the RGB bands. We present several θ_i angles and plot a continuous range of BRDF measurements for θ_r .




Figure 3.11: Polar plots of measured incidence plane BRDF for two different threads. The left plot contains the RGB average BRDF for a filament silk thread, and the right plot a filament polyester thread.

The plots demonstrate the similarity among filament threads and their disparity from spun threads in Figure 4.3. Both the polyester and the silk filament threads possess narrow specular lobes oriented at the exact specular reflection direction. This result can be attributed to their low surface roughness as well as minimal internal scattering. The fact that the lobe is oriented at the exact specular reflection direction means that, unlike hair, threads have no consistent cuticle that displaces their specular reflection. The polyester thread is the more specular of the two filament threads, as evidenced by its narrower and brighter reflection lobe. This can be attributed to the synthetic vs. organic fibers that they are composed of, where polyester has fewer natural imperfections and irregularities due to its industrial fabrication process.

At first glance, the spun threads in Figure 3.12 appear nearly identical. They both exhibit a characteristic wide scattering lobe that slowly increases as the




Figure 3.12: Polar plots of measured incidence plane BRDF for two different threads. The left plot contains the RGB average BRDF for a staple linen thread, and the right plot a staple polyester thread.

light goes to glancing angle as well as no distinctly identifiable specular lobe. Focusing on glancing incidence angles, we observe that the spun linen thread scatters more light in the non-specular directions. This type of BRDF can be attributed to either a very rough surface or isotropic internal scattering. We address these behaviors in the next chapter when we present our thread BRDF model.

3.3.2 Hemisphere BSDF

We extend our previous discussion of BRDFs to include the full sphere of scattering directions, namely the BSDF. By assuming rotational homogeneity with respect to the azimuthal angle, we can characterize the BSDF of a thread by measuring half of the scattering sphere around it, i.e $\phi_{id} = [0 - 180]$. Figure 3.13 (left) shows a schematic of the acquisition angles, where the light only moves along the θ keeping its ϕ angle at zero, and the camera moves along both ϕ and θ dimensions. Intuitively, when ϕ_{id} is zero, the light and the camera are in the sample plane, whereas when ϕ_{id} approaches 180, the light and the camera are on opposite sides of the thread. On the right portion of the figure, we show our raw sampled, and interpolated results for polyester thread illuminated from $\theta_i = 30^\circ$.

Based on our experiments, we observed that the BSDF varied more along the longitudinal dimension (θ) than it did along the azimuthal direction (ϕ). To reduce our BSDF measurement duration, while still producing the best possible

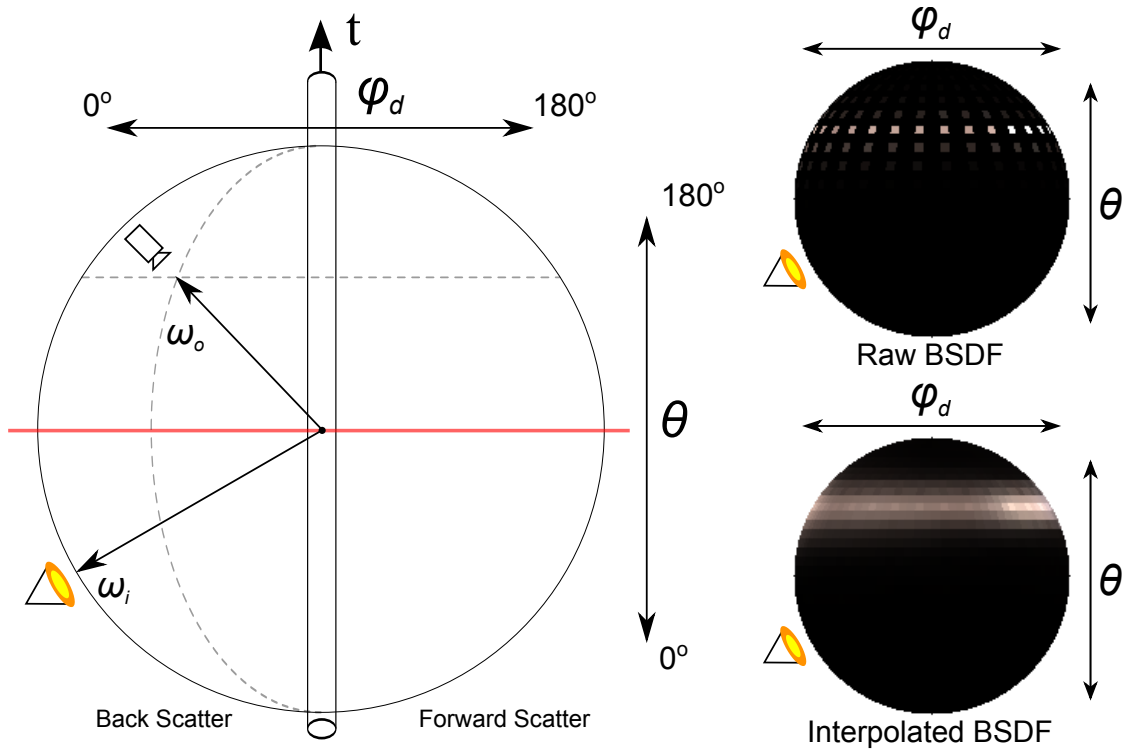


Figure 3.13: Hemisphere plots of the BSRDF as a function of the scattering direction when the position of the light is held fixed at $\theta_i = 30^\circ$ incidence.

fidelity, we chose to sample the longitudinal dimension at 4° intervals and the azimuthal dimension at 12° intervals. To derive a continuous representation of scattering, we interpolate our sampled measurements in three-dimensional BSRDF space (see top right of Figure 3.13). The interpolated BSRDF value p_i is expressed as:

$$p_i = \frac{\sum_{m=1}^n p_m / (1 + d_m)}{\sum_{m=1}^n 1 / (1 + d_m)} \quad (3.8)$$

where p_m is a closest measured BSRDF point, and d_m is the distance to that point from p_i . The idea is to compute a weighted sum of n of the nearest points, and normalize by the total distance. Using our smoothly interpolated sphere BSRDF, we can focus our observations more precisely by splitting forward scattering from back scattering. As shown in Figure 3.13, back scattering is light that gets scattered within 90° azimuthally, relative to the incident light. It is effectively the thread BRDF since it describes light returning to the same hemisphere from which it was

incident. Forward scattering on the other hand describes light that has transmitted through, or scattered through the thread and continues forward, away from the incident light hemisphere. Surface scattering is also a component in forward scattering since the thread cross-section is round and can reflect light at a glancing angle. We present all the measurements performed in this thesis in terms of forward (bottom row), and back (top row), scattering hemispheres in Figure 3.14. Here, each incident hemisphere is plotted in projection onto the tangent plane of the thread, with its direction vertical as in Figure 3.13. The hemispheres are arranged to indicate the incident direction of light on the thread, starting from being lit from below (left), to being lit from above (right). In the back-scattering data, the position of the light can be seen as a dark spot due to occlusion between the light and the camera (see wool). In the forward scattering hemispheres the dark spots are an artifact of overexposed/invalid measurements, which occur when the light is directly visible to the camera. The salient features of the data are qualitatively summarized in Table 3.2.

Table 3.2: Scattering characteristics of thread samples based on data in Figure 3.14. Next to each thread type, (s) denotes spun, and (f) denotes filament. The first word in each entry identifies the magnitude of the scattering phenomenon (high-med-low), and the second word describes its scattering lobe shape (shiny-glossy-blurry). Plots in each column of the figure correspond to a light incidence angle θ_i as labeled at the very top.

Thread Type	Back Scatter		Forward Scatter	
	Surface	Internal	Surface	Internal
polyester (f)	high / shiny	low / glossy	high / shiny	low / shiny
polyester (s)	med / blurry	med / blurry	med / shiny	med / blurry
silk (f)	med / shiny	med / glossy	high / shiny	med / shiny
silk (s)	low / glossy	med / blurry	med / glossy	med / glossy
rayon (s)	med / glossy	med / blurry	med / shiny	low / glossy
wool (s)	med / blurry	high / blurry	high / glossy	med / blurry
linen (s)	low / glossy	high / blurry	high / shiny	med / blurry

3.3.3 Scattering Features

We can reason about the data in Figure 3.14 by identifying important features in it.




Figure 3.14: BRDF visualization sets of backward scattering hemisphere (top row) and forward scattering hemisphere (bottom row). The labels on the left correspond to the yarn type and the labels on the right correspond to fiber type.

Color

The color and brightness of each hemisphere is related to the magnitude of light absorption. For example, when we compare a spun rayon (i) to a spun wool (k), we observe that (i) is much dimmer which implies that less light is being scattered, i.e. it is being absorbed. Since in our thread samples this absorption is not uniform across different wavelengths, different threads have different colors. Similarly, the forward and back scattering hemispheres for any one thread have nearly identical colors.

Every hemisphere in our data set can be classified as being composed of two color components. One is a background color and the other is a white foreground color added on top of it. This corresponds to a model of subsurface scattering (colored) and surface reflection (white). For example in (c) we observe a green band with a slightly narrower white band added on top of it. Furthermore, all the forward scattering hemispheres (b, d, f, h, etc.) display a brighter and narrower white band than their back scattering counterpart. This is due to Fresnel reflection which increases that amount of energy reflected as the light becomes more grazing. All of the forward scattering hemispheres observe light at a more grazing angle than their equivalent back scattering hemisphere.

Anisotropy

The shapes of the color distribution in the BSDF hemispheres gives us a hint as to how we might model the light scattering behavior. For example, the white band is very apparent in (a), (c), and even (i). The fact that this band extends from the left side of the hemisphere all the way to the right, implies a cone scattering shape. Figure 3.15 demonstrates the relation of a cone to our hemisphere BSDF measurements.

Another important observation is that subsurface scattering (color) seems to also be constrained to this cone shaped distribution in (a, b c, d, g, h). This implies that despite entering into the bulk of the thread, the light path remains correlated to the ideal surface reflection direction.




Figure 3.15: Schematic relating a bright band in a hemisphere plot to a cone shaped reflectance from a cylinder.

Roughness

By observing the width or diffuseness or the BSDF components, we can assess both the surface roughness and the internal roughness. A rough surface results in light being scattered in a greater variety of directions. For example in (a) we notice a narrower white and color scattering band than we do in (c). This tells us that both the surface and subsurface scattering is less varied in (a) than it is in (c). On the other hand, the wool sample (k) shows a full sphere of color, which implies that internal scattering diffuses light almost uniformly in all directions. Spun silk displays an interesting behavior where backscattering (g) spreads light more diffusely than forward scattering (h). This implies that due to absorption, diffusely scattering rays disappear in the forward scattering direction leaving only weakly scattered rays.

Chapter 4

Light Scattering From Threads

This chapter provides deals with models that describe light behavior when it intersects a surface, i.e. encounters a boundary with a different index of refraction. Due to the complexity of the physical interaction of light and matter, we focus on topics relevant to our study of threads.

4.1 Surface Scattering

Analytical models of surface reflection fall into two categories: phenomenological - models that approximate reflectance without analyzing underlying principles (ad-hoc), and physically based - models that use physical properties and simplifying assumptions (e.g. dielectrics) about a surface to approximate reflectance. We focus our discussion on physical modeling of reflectance as our goal is to rely on physical phenomena instead of ad-hoc parameters.

Cook-Torrance [CT82] introduced a popular reflectance model into computer graphics, based on the work of [TS67]. It models the surface as a collection of tiny mirror-like facets oriented in controllable directions. The model takes into account shadowing between micro-facets as well as masking of facets from the view perspective. To provide physical plausibility, the model includes an averaged (ignoring polarization) Fresnel term for accurately modeling reflectance of individual

facets. The model can be expressed as:

$$f_r = \frac{1}{\pi} \left(k_d + k_s \frac{FDG}{\cos(\theta_i) \cos(\theta_o)} \right) \quad (4.1)$$

where k_d and k_s are diffuse and specular coefficients respectively; F is the Fresnel factor; D is the micro-facet distribution; and G is the shadowing/masking term. The key components that we are interested in are the Fresnel term F and the micro-facet distribution D . These two components are the foundation of describing surface roughness and therefore surface reflectance.

4.2 Smooth Cylinder Model

Based on our measurements of individual threads, we observed an optical behavior that is similar to hair and more generally, smooth dielectric cylinders. We observed that the reflection of a collimated light beam from a taut thread sample forms a cone centered on the thread axis. Additionally, the surface reflection is framed by a subtle color reflection that is also centered on the cone. This type of scattering has been previously studied in the context of hair by [KK89, Kim02, MJC⁺03]. In these previous approaches, the normal plane around the tangent is used as the coordinate frame for computing light scattering behavior (see Figure 4.1).

Notation

Here, the plane perpendicular to the direction of the cylinder is referred to as the normal plane. The direction to the light source is ω_i , and the direction from which scattering is observed is ω_r . The inclination of the two vectors with respect to the normal plane is denoted as θ_i and θ_r . The azimuthal angles around the cylinder are computed by projecting the incident and outgoing vectors onto the normal plane and are referred to as ϕ_i and ϕ_r .

This coordinate frame turns the cylinder into a one-dimensional entity, assuming that it's perfectly circular and eliminating the perception of surface area based on its thickness. Defining a BSDF requires a unit of projected area and




Figure 4.1: Angular notation for cylinder based models.

solid angle. Alternatively, if we replace area with length, then our BSDF definition becomes: the ratio of intensity per unit projected length to flux per unit length.

$$f_s(\omega_i, \omega_r) = \frac{dL(\omega_r)}{dE(\omega_i)} = \frac{\text{Watts/m/sr}}{\text{Watts/m}} \quad (4.2)$$

This curve radiance concept and its integral was introduced by [MJC⁺03]:

$$L_r = D \int f_s(t, \omega_i, \omega_r) L_i(\omega_i) \cos \theta_i \, d\omega_i \quad (4.3)$$

Here D represents the diameter of the cylinder. It is necessary due to the fact that in our BSDF definition the denominator is flux per unit length, which doesn't take into account that a thicker thread will intercept more light than a thin thread.

Scattering

When a beam from ω_i , consisting of parallel rays of light, strikes a cylinder running along the vector t , each ray in the beam reflects at the surface according to the surface normal of the cylinder. These surface normals are all perpendicular to the thread tangent vector t and lie in the normal plane. For a smooth specular cylinder, a beam incident at θ_i will be reflected in the ideal specular direction $-\theta_i$

across the normal plane and, due to the circular cross section of the cylinder, will be spread into a cone [KK89]. The refracted light will enter the cylinder and after any number of internal reflections and refraction will emit into the same cone as the surface specular reflection [MJC⁺03]. This physically based result was derived by [MJC⁺03] using Bravais’ Law.

Table 4.1: Description of important symbols.

f_s	Thread scattering function
F_r	Fresnel reflectance
F_t	Fresnel transmittance
γ_s	Surface reflectance Gaussian width
γ_v	Volume scattering Gaussian width
k_d	Isotropic scattering coefficient
A	Colored albedo coefficient

4.3 A Light Scattering Model for Thread

In our BSDF model we abstract the thread geometry as a cylinder. By using the idealized cylinder formulation described in the previous section, we are able to greatly simplify the analytical description of light scattering from threads. As in previous treatments of BSDFs [HK93], we separate our scattering function $f_s(t, \omega_i, \omega_r)$ into a surface scattering component f_r , and volume scattering component f_v . In addition to the angles in Figure 4.1, we introduce $\phi_d = \phi_i - \phi_r$ and $\theta_h = (\theta_i + \theta_r)/2$ to define the two scattering functions.

Surface Reflection

We model surface reflection similarly to [MJC⁺03], except we do not decompose our computation into longitudinal and azimuthal planes.

$$f_r(t, \omega_i, \omega_r) = F_r(\eta, \vec{w}_i) \cos(\phi_d/2) g(\gamma_s, \theta_h) \quad (4.4)$$

The $\cos(\phi_d/2)$ term arises due to projection of the circular cylinder cross-section, as demonstrated by [Kim02], and previously used by [SPJT10] for hair rendering. To break away from the idealized smooth cylinder representation, we employ a

unit area Gaussian g with width γ_s to simulate surface roughness. Since the Gaussian is dependent on only the θ_h angle, the BSDF takes on a symmetric sweep around the ϕ dimension forming a cone. Finally, we add physical accuracy to the model by attenuating the power by a Fresnel term. The actual angle used to compute the Fresnel term is based on the reflection normal on the cylinder as well as a half-angle between the light and the eye, yielding an exact expression $F_r(\eta, \cos^{-1}(\cos(\theta_h) \cos(\phi_d/2)))$. This surface reflectance approach is nearly identical to the Cook-Torrance micro-facet model, where D and F are equivalent to our g and F_r . This model produces a glossy reflection on a cone around the thread with physical and geometric attenuation. We considered using the full micro-facet specular formulation, but found that it did not improve matching to our measured results.

Volume Scattering

Real threads are composed of fibers that are either twisted together or lay flat next to each other. We make a unifying assumption that all fiber types are cylindrical with minimal eccentricity. This is generally true with the exception of cotton, which resembles a flat ribbon. To summarize, our model is a large thread cylinder composed of tiny fiber subcylinders. This enables us to use the fact that smooth cylinders emit light due to internal scattering into the ideal reflection cone. Therefore, light that enters the thread volume and undergoes any type of scattering interaction with the fiber subcylinders will result in a surface emission into the same cone as the surface reflection. One thing to note is that the orientation of the fiber subcylinders differs from that of the thread cylinder. We model this deviation as a normal distribution centered on the thread tangent. Therefore, a filament thread will have a much smaller variance than a spun thread. We model these behaviors while maintaining physical constraints in the following equation:

$$f_v(t, \omega_i, \omega_r) = F \frac{(1 - k_d) g(\gamma_v, \theta_h) + k_d}{\cos \theta_i + \cos \theta_r} A \quad (4.5)$$

Here $F = F_t(\eta, \vec{w}_i) F_t(\eta', \vec{w}_r')$ is the product of two transmission Fresnel terms. This double Fresnel term was shown by [WNO98] to yield a superior match to measurements of subsurface scattering than the Lambertian reflection model. We define

the subcylinder tangent deviation with a Gaussian lobe g with width γ_v . The Gaussian lobe controls the width the forward scattering cone. For spun threads, which consist of fibers that deviate from the thread tangent direction, this Gaussian is wider than for filament threads, which mainly consist of parallel filaments. Additionally, we define a tunable isotropic scattering term k_d and a color albedo term A . We added the isotropic scattering term to account for cellulose based fibers such as cotton and linen, which predominantly yield isotropic volume scattering instead of a forward scattering cone. The division by the sum of projected cosines comes from [Cha60], in his derivation for diffuse reflectance due to multiple scattering in a semi-infinite medium. Adding this normalization term gave us better matches with our measured results. The complete thread scattering model is a sum of the surface and volume components:

$$f_s(t, \omega_i, \omega_r) = (f_r(t, \omega_i, \omega_r) + f_v(t, \omega_i, \omega_r)) / \cos^2 \theta_d \quad (4.6)$$

Note that the complete scattering formulation contains a division by $\cos^2 \theta_d$, which is necessary to account for the solid angle attenuation of the specular cone [MJC⁺03]. Previous work has addressed volume scattering in threads with a cylindrical phase function in [Ira08] as well as the Henyey-Greenstein phase function in [AMTF03]. We experimented with various phase functions as well, but found them inadequate due to their decoupled behavior from the direction of the thread. Our approach is similar in spirit to [JAM⁺10], which defines phase functions oriented to the direction of fibers to achieve highly anisotropic volume scattering.

We have defined a complete BSDF for individual threads, which matches well to our measured results. It was our goal to define as few non-physical control parameters as possible to enable the physical and geometric scattering constraints to drive its behavior. We note that the model is only suitable for distant viewing of threads since it assumes that the rays of light incident on the thread cylinder are parallel and that the thread is locally straight.

4.3.1 Model Validation

In this section we provide evidence for the validity of our model by comparing it to measured BRDFs of thread. We accomplish this by manually fitting our model parameters to measured results. We did not consider automatic fitting approaches due to the minimal number of control parameters and their predictable nature.

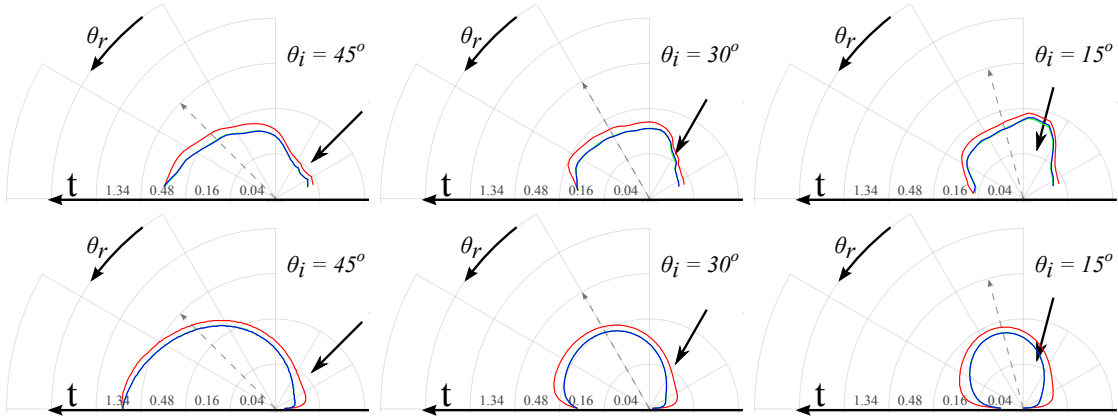


Figure 4.2: Incidence plane thread BRDF measurements in the top row matched by thread model in the bottom row. The thread sample is twisted and was extracted from the *polyester satin charmeuse* cloth. The plots show scattering as a function of view angle.

In Figure 4.2, we demonstrate the performance of our model by comparing the measured BRDF results in the top row with results generated our model in the bottom row. Each row shows BRDF measurements of a staple polyester thread for three incident light angles. Our model results in the bottom row are able to closely match the measured results in the top row. To achieve this close match, we tune the parameters of our model by observing a wide surface reflectance supplemented by an even wider red tinted volume scattering in the measured BRDF.

In Figure 4.3 we present our model for a filament polyester thread. The model again closely simulates the scattering profile of the filament thread by setting a very narrow surface reflectance Gaussian and a small red tinted albedo. Thus, our thread model is validated by being able to closely simulate the scattering behavior of different thread types under varying incident light angles.

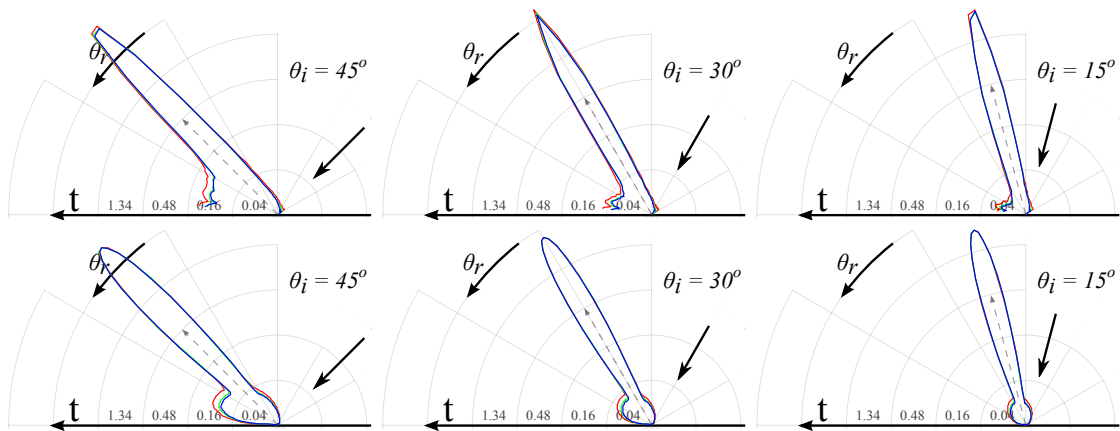


Figure 4.3: Incidence plane thread BRDF measurements in the top row matched by thread model in the bottom row. The thread sample is filament type (un-twisted) and was extracted from the *polyester satin charmeuse* cloth. The plots show scattering as a function of view angle.

Chapter 5

Modeling the Appearance of Cloth

This chapter presents a use-case, fabric visualization, for the thread reflectance model developed in the previous chapter. To reproduce the appearance of fabric, we identify the weave pattern as a critical component of the surface description. We consider fabric as a mesh of interwoven threads oriented in two orthogonal directions. By providing a weave pattern representation and evaluating our analytical thread model, we are able to visualize virtually any fabric.

5.1 Previous Approaches

Rendering cloth has been an active area of research for more than 25 years. The earliest approaches as well as more recent work are based on empirical shading models [Wei86, DLHS01, GD04]. Glumac and Doepp [GD04] developed an anisotropic shader that extends a general lighting model to use thread direction parameters to achieve believable looks for several silk-like fabrics. Daubert et al. [DLHS01] proposed a material representation based on the Lafortune reflectance function coupled with an occlusion term to handle several view dependent effects of coarsely woven and knitted fabrics. Not concerned with physical accuracy, these models only provide an imitation of cloth appearance and typically lack the ability to replicate some of its more complex features.

Microfacet models have been used by Ashikhmin et al. to model satin and velvet [APS00]. Adabala et al. continued this work by including support for weave patterns [AMTF03]. Wang et al. [WZT⁺08] introduced their own microfacet-based BRDF for modeling spatially-varying anisotropic reflectance using data captured from a single view. While microfacet models are effective at simulating complex materials, these models are difficult to control as they strongly depend on the right normal distribution function. Since cloth is often anisotropic it is difficult to obtain this distribution from measured data.

Another approach for simulating cloth is based on modeling the structure of the cloth [XCL⁺01, CLZ⁺03, DC04]. While these methods can reproduce a wide range of appearances they can be difficult to control. Yasuda et al. [YYTI92] modeled the peculiar gloss seen in cloth by accounting for the internal structure, but assumed a very simplified model of the structure and the results lacked verification. Westin et al. [WAT92] computed BRDFs for velvet and plain weave nylon fabrics through statistical ray tracing of a geometric model of the small-scale cloth structure. Zhao et al. [ZJMB11] presented a volumetric rendering approach using CT scanning of fabrics. Their model produces high quality renderings, but it is limited to reproducing specific cloth samples.

Irawan et al. developed a comprehensive model for reproducing both the small-scale (BTF) and large-scale (BRDF) appearance of woven fabric [IM06, Ira08]. This state of the art model is capable of reproducing a wide range of appearances by simulating light scattering at the thread fiber level.

5.2 Weave Pattern Representation

A weave pattern representation needs to be able to describe the topology of fabric and be flexible enough to handle many varieties. As mentioned earlier, the unifying characteristic of all fabrics is that they are composed of two dominant thread directions, warp and weft. In addition to describing occupancy, which thread is where, we need to describe the trajectory of the thread (what angle relative to the surface is the thread). We first describe a simple approach to

defining thread trajectory followed by a description of weave pattern occupancy encoding.

5.2.1 Thread Curve Definition

The warp and weft threads in a weaving pattern follow a repeating path across the fabric. We define a *tangent curve* that describes this repeating tangent distribution for a smallest patch (see Figure 5.1 right). We specify the tangent curve by setting the tangent values at discrete control points. With this definition we can evaluate any point on the curve using linear interpolation.




Figure 5.1: The weaving pattern and a sample tangent curve for the *polyester satin charmeuse* fabric: (left top) the weaving pattern, (left bottom) a smallest patch, (right) the tangent curve for the two types of threads.

5.2.2 Weave Definition

The simplest way to define a weave pattern is using an m by n grid where each grid cell (i, j) represents an intersection between a warp and a weft thread. This binary grid encodes 0 for weft on top and 1 for warp on top. We can come up with a number of (not necessarily intuitive) boolean expressions to generate a variety of weaving patterns. In Figure 5.2 we show the results of a simple encoding for a plain and satin weave. To produce our final results, we extended this basic framework with shift and scale operators to give more control for generating the binary pattern.


```

for i=1:n
  for j=1:m
    A(i,j) = mod((i + j), N) == 0
  end
end
end

```




Figure 5.2: Weave pattern encoding pseudocode (left), and the resultant patterns of satin (middle) and plain (right).

Alternatively, we can specify this grid as a texture that is authored in an image editing program. A thorough treatment of weave pattern synthesis was developed by [Gla02], where warp and weft visibility is encoded more robustly to enable embroidery like patterns.

5.2.3 Local Shadowing

Thread scale shadowing is very important to take into account in order to synthesize a realistic fabric image. Poulin and Fournier [PF90] derived a shadowing and masking term for grooved surfaces composed of cylinders. However, their approach is not applicable to our model since they assumed that the cylinders have a surface patch BRDF and integrated all of the reflected light scattered toward a viewer. Since our formulation treats cylinders as one-dimensional entities, we do not compute the explicit reflectance variation across their circumference.

We adopted an approximate thread-to-thread shadowing approach which leverages our procedural weave definition to produce a robust shadowing solution. The idea stems from a height map representation of our weave pattern. When computing the thread tangent direction, we can also compute a height value (in tangent space) for a position on the thread. Thus, a simple shadowing solution is to check if the vector to the light from our current height is occluded by the horizon.

The horizon point can be computed in several different ways. We take advantage of the repeating nature of a weave pattern and find the highest point on either thread (warp or weft), claiming that it is the highest point in any direction.

This assumes that the occluding horizon thread is not affected by the overall mesh curvature, which is a fair assumption considering how small the threads are relative to the geometric cloth detail. While this doesn't produce an exact shadowing solution, it does provide an approximation for shadowing of low lying threads vs. top threads. The effect of our approximate horizon shadowing is demonstrated in Figure 5.3.




Figure 5.3: Approximate thread-to-thread shadowing is shown (middle), compared to no shadowing (left), and the shadow contribution mask (right). The light is hitting the fabric at a glancing angle and causing higher lying threads to block their lower lying neighbors.

Here, the fabric pattern for *silk crepe de Chine* is rendered when viewed from above. The weave pattern consists of a weft thread (blue) that goes up and down, and a warp thread which lays flat in between the weft threads. With lighting coming from a grazing angle at left, we can see in the unshadowed image that the left weft side is brighter than the right. This is due to the thread BSDF alone, which assumes that that thread segment is in isolation and receiving a full contribution from the light. Additionally, observe that the warp thread which lies lower than most of the weft thread is also fully lit. The shadowed image corrects both of these cases by reducing the amount of light incident on those low lying portions. The image on the right shows the shadow mask that is computed by our horizon shadowing algorithm. As the light moves towards normal with the fabric, the shadow mask gets lighter thereby reducing shadowing. Since the shadow mask is computed from the weave pattern and the thread curve definition, it is dynamic and produces plausible results for a variety of fabrics.

5.3 Image Synthesis

We generate an image of fabric by tracing rays from a virtual camera to a smooth triangulated mesh. The goal is to compute the integral of light transport through each pixel. By creating paths from the camera to the light source(s), with an intermediate bounce at the fabric surface, we are able to reconstruct the appearance of fabric.

At the intersection point of the ray against the mesh, we compute a UV coordinate, and three vectors that form the tangent space basis (tangent, normal, bitangent). Using our weave pattern definition described in the previous section, we determine whether we have hit a warp thread or a weft thread, as well as the offset on that thread. We use the offset to compute the thread tangent direction at the intersection point, using our thread curve definition.

At this stage we know the camera direction, the light direction, and the thread tangent direction, which is everything we need to evaluate our thread BSDF. The solution to the BSDF tells us how much energy is transferred from the light ray to the camera ray by scattering from the fabric surface.

The introduction of spatial variation into our appearance model, requires us to extend the four variable BRDF concept. Adding two more variables (u, v) to the standard BRDF $f_r(\omega_i, \omega_r)$, we arrive at a bidirectional texture function (BTF), which has all the same properties as a BRDF with the addition of spatial variation. A similar concept to the BTF is the SVBRDF, however it does not represent non-local effects such as shadowing which our appearance model does. We formalize our appearance model BTF as the following:

$$T(\omega_i, \omega_r, u, v) = f_r(\omega_i, \omega_r, t(u, v))g(\omega_i, u, v) \quad (5.1)$$

where for a position (u, v) on the fabric, we evaluate the thread BRDF f_r with respect to a tangent direction t , and multiply the result by the local shadowing term g .

Figure 5.4 demonstrates our results for six fabrics, four of which we observed directly under a microscope. All images are rendered using path tracing with 128 samples per pixel, with a maximum of 4 indirect bounces. Note the




Figure 5.4: Path traced images of physically based cloth shading. Each image synthesizes a real world cloth sample. Both the thread light-scattering properties and the exact weave structure are reproduced in our models, yielding a high fidelity simulation of the appearance of the real cloth samples. The fabrics pictured are: (a) *polyester satin charmeuse*, (b) *silk crepe de Chine*, (c) *shot fabric*, (d) *velvet*, (e) *linen*, and (f) *wool*.

variation in the appearance and location of the specular highlight as well as the variation in internal scattering (e.g. shot fabric vs. wool). The appearance of each fabric is entirely driven by intuitive physical parameters of roughness, index of refraction, albedo, and weave structure.

5.4 Results Validation

To demonstrate the effectiveness of our cloth appearance model, we compare our synthetic results against photographs. To capture the anisotropic behavior of different fabrics, we wrapped them around a cylinder in three different directions. We label each mode based on the orientation of the filament threads as vertical, horizontal, and diagonal (see Figure 5.5). For the *plain linen* fabric, the vertical and horizontal modes are identical due to the symmetric weave pattern. For comparison, we present our rendered results of different fabrics in the same setup.

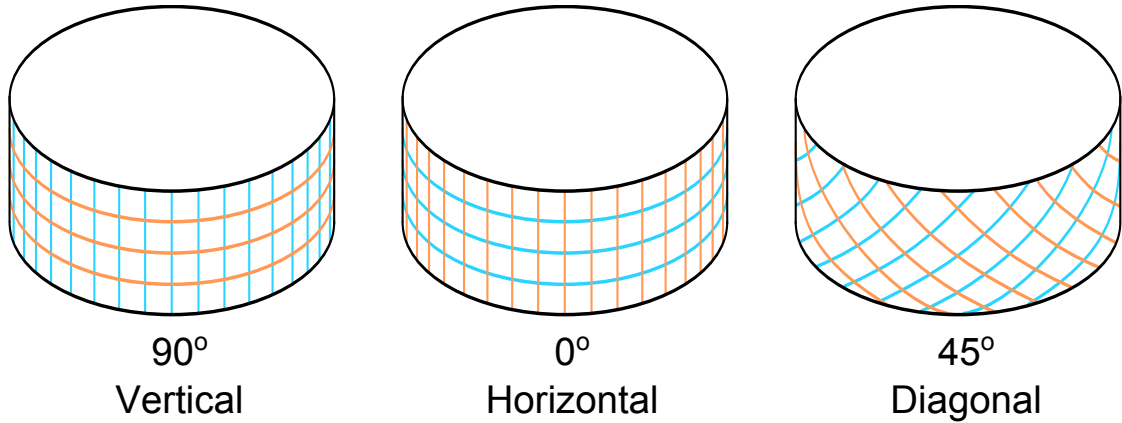


Figure 5.5: Result validation setup, show the direction in which the cloth samples are wrapped.

Figure 5.6 shows our results comparing photographs of *plain linen*, *silk crepe de Chine*, and *polyester satin charmeuse* fabrics against our rendered results. Note, how our model is able to capture the variation in the highlights across the different fabrics including the grazing angle highlight seen on the silk, and the triple-highlight seen on the satin. In selecting our model parameters, we focused on matching the appearance of only the center cylinder. The middle columns of all

three cloth samples show anisotropic behavior (even slightly visible for the linen), and our model is able to not only capture this, but also accurately predict the appearance when the cloth sample is rotated at 90° and 45° .




Figure 5.6: Comparison of reference photographs and the rendered results for (top) *silk crepe de Chine*, (middle) *polyester satin charmeuse*, and (bottom) *plain linen*. The three cylinders for each fabric show three orientations as illustrated in Figure 5.5.

Chapter 6

Conclusion

We have presented a practical appearance model for cloth, which is able to reproduce the complex appearance of a variety of fabrics. It is based on our novel analytical thread BSDF model, which was developed from observations of thread reflectance measurements. To acquire the measurements, we developed an accurate acquisition system for thin cylindrical material samples.

Our appearance model is similar to the model of Ashikhmin et al. [APS00] in the sense that they can both reproduce a specular reflection in any direction by providing a vector distribution in the appropriate direction. We reproduce specular peaks by orienting the tangents so that their reflection cone lies in the desired direction, whereas their model uses a normal distribution function.

The main advantage of our model is that it takes intuitive parameters that are fabric specific. The parameters are derived from observation of the weave pattern, as well as the constrained parametric space of the thread BSDF. This enables the model to describe fabric BRDFs with or without measurements of a real cloth sample. A microfacet model hides the distinction between weave pattern and thread BSDF, by abstracting both as a distribution of normals. While producing reasonable approximations, it does not allow for intuitive modifications of the BRDF, e.g. changing the smoothness of one of the threads, or making the weave tighter/looser. The thread BSDF and the weave pattern are well defined and orthogonal problems. By keeping them separate and enforcing physically plausible constraints for each, we ensure that the resultant fabric appearance is

also plausible. An added benefit of our model is the precise control of spatial variation in fabric color and weave by setting the color of each thread and each tangent curve. This allows our model to realistically mimic the anisotropic color shifts found in shot fabrics as well as the non-uniform nature of coarse wool.

Comprehensive state of the art research in cloth rendering has been carried out by Irawan [Ira08]. They present a rigorous model for computing the light reflection off of yarn threads, which are simulated as an assembly of specular fibers. His model incorporates costly numerical integrations (with no performance reports) and a fitting process to estimate the value of different control parameters. However, they do not present a thread-to-thread shadowing term, which results in having to use a non-physical smoothstep function to adjust grazing angle views and lighting (which is not presented in their measurements). Another limitation of their model is that it approximates the tangent distribution of threads with conic sections, which in turn limits the ability of the model to match complex or irregular tangent distributions. Due to this limitation, their model cannot represent BRDFs with more than two specular highlights (e.g. the *polyester satin charmeuse* studied in this paper) or discontinuous distributions such as Velvet. Irawan [Ira08] does not provide such a complex appearance in rendered results.

One limitation of our approach is that for medium range visualization it produces a distracting moire pattern. This arises due to the high frequency detail in the cloth not being properly resolved by our finite sampling algorithm. Another limitation of our model in its current form is that it cannot accurately produce extreme close-up renderings. It does not reproduce the reflectance variation across each thread’s circumference, instead, assuming that the detail across a single thread is smaller than a pixel in the image plane. It’s also limited by the fact that it only simulates single scattering and thereby loses energy that could otherwise be scattered by multiple light-thread interactions.

6.1 Future Work

An important task in future research is to come up with a better sampling or filtering approach for the fabric appearance model. Given the structured nature of fabric and our analytical thread BSDF, we should be able to reduce the number of samples and improve antialiasing by preconvolving our BSDF with a patch tangent distribution. Another direction for future research is investigating the shadowing and masking between threads with a full horizon mapping [Max88] implementation. In addition, we are interested in testing an automated fitting processes to estimate the control parameters of our model based on photographs of fabric wrapped around a cylinder. Furthermore, we aim to simulate forward scattering through fabric as well as approximate the multiple scattering of light between different threads. Lastly, it would be interesting to investigate different ways of importance sampling our fabric appearance model.

Appendix A

Additional Renders




Figure A.1: Closeup rendered image demonstrating the procedural texture found in fabrics. The texture arises from the weave pattern as well as the local thread shadowing.




Figure A.2: Extreme closeup render (top) showing limitations of our approach where each thread is uniformly colored and lacking detail from the fibers that make it up. The second image (bottom) shows the moiré pattern that arises due to the high frequency texture detail found in fabrics.




Figure A.3: Clothing rendered using our appearance model. The precise control of thread parameters allows us to define multi-color fabrics (right).

Bibliography

- [AMTF03] Neeharika Adabala, Nadia Magnenat-Thalmann, and Guangzheng Fei. Visualization of woven cloth. In *Proceedings of the 14th Eurographics workshop on Rendering*, EGRW '03, pages 178–185, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.
- [APS00] Michael Ashikhmin, Simon Premože, and Peter Shirley. A microfacet-based brdf generator. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '00, pages 65–74, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.
- [Bay76] E Bayer, Bryce. Color imaging array, July 1976.
- [Cha60] Subrahmanyan Chandrasekhar. *Radiative Transfer*. Dover Publications, 1960.
- [CLZ⁺03] Yanyun Chen, Stephen Lin, Hua Zhong, Ying-Qing Xu, Baining Guo, and Heung-Yeung Shum. Realistic rendering and animation of knitwear. *IEEE Transactions on Visualization and Computer Graphics*, 9(1):43–55, 2003.
- [Con10] L. Condat. A simple, fast and efficient approach to denoising: Joint demosaicking and denoising. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pages 905–908, sept. 2010.
- [CT82] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. Graph.*, 1:7–24, January 1982.
- [DC04] Frédéric Drago and Norishige Chiba. Painting canvas synthesis. *Vis. Comput.*, 20(5):314–328, 2004.
- [DLHS01] Katja Daubert, Hendrik P. A. Lensch, Wolfgang Heidrich, and Hans-Peter Seidel. Efficient cloth modeling and rendering. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 63–70, London, UK, 2001. Springer-Verlag.

- [DM97] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques, SIGGRAPH '97*, pages 369–378, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [DvGNK99] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 18:1–34, January 1999.
- [GD04] Rick Glumac and David Doepp. Generalized approach to rendering fabric. In *ACM SIGGRAPH 2004 Sketches*, SIGGRAPH '04, pages 120–, New York, NY, USA, 2004. ACM.
- [Gla02] A. Glassner. Digital weaving. 1. *Computer Graphics and Applications, IEEE*, 22(6):108 – 118, nov/dec 2002.
- [HK93] Pat Hanrahan and Wolfgang Krueger. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques, SIGGRAPH '93*, pages 165–174, New York, NY, USA, 1993. ACM.
- [HP05] K. Hirakawa and T.W. Parks. Joint demosaicing and denoising. In *Image Processing, 2005. IICIP 2005. IEEE International Conference on*, volume 3, pages III – 309–12, sept. 2005.
- [IM06] P. Irawan and S. Marschner. A simple, accurate texture model for woven cotton cloth. Technical report, Cornell University, Department of Computer Science, 2006.
- [Ira08] Piti Irawan. *Appearance of woven cloth*. PhD thesis, Ithaca, NY, USA, 2008. AAI3295837.
- [JAM⁺10] Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Trans. Graph.*, 29:53:1–53:13, July 2010.
- [Kim02] Tae-Yong Kim. *Modeling, Rendering and Animating Human Hair*. PhD thesis, University of Southern California, Los Angeles, CA, USA, 2002.
- [KK89] J. T. Kajiya and T. L. Kay. Rendering Fur With Three Dimensional Textures. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, pages 271–280, New York, NY, USA, 1989. ACM.

- [LA98] Tianyi Liao and Sabit Adanur. A novel approach to three-dimensional modeling of interlaced fabric structures. *Textile Research Journal*, 68(11):841–847, 1998.
- [Lab98] LabsphereE. Color imaging array, September 1998.
- [Max88] Nelson L. Max. Horizon mapping: shadows for bump-mapped surfaces. *The Visual Computer*, 4:109–117, 1988. 10.1007/BF01905562.
- [MJC⁺03] Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light Scattering from Human Hair Fibers. *ACM Transactions on Graphics*, 22(3):780–791, 2003.
- [MWL⁺99] Stephen R. Marschner, Stephen H. Westin, Eric P. Lafortune, Kenneth E. Torrance, and Donald P. Greenberg. Image-based brdf measurement including human skin. In *Rendering Techniques*, pages 131–144, 1999.
- [NRH⁺92] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Radiometry. chapter Geometrical considerations and nomenclature for reflectance, pages 94–145. Jones and Bartlett Publishers, Inc., , USA, 1992.
- [PF90] Pierre Poulin and Alain Fournier. A model for anisotropic reflection. In *Proceedings of the 17th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '90, pages 273–282, New York, NY, USA, 1990. ACM.
- [RBS99] M.A. Robertson, S. Borman, and R.L. Stevenson. Dynamic range improvement through multiple exposures. In *Image Processing, 1999. ICIP 99. Proceedings. 1999 International Conference on*, volume 3, pages 159 –163 vol.3, 1999.
- [Sho24] S. A. Shorter. The physical properties of textile fibres in relation to technical processes and to general colloid theory. *Trans. Faraday Soc.*, 20:228–235, 1924.
- [SPJT10] Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. An artist friendly hair shading system. *ACM Trans. Graph.*, 29:56:1–56:10, July 2010.
- [TS67] K. E. TORRANCE and E. M. SPARROW. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.*, 57(9):1105–1112, Sep 1967.
- [War92] Gregory J. Ward. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.*, 26:265–272, July 1992.

- [WAT92] Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. Predicting reflectance functions from complex surfaces. *SIGGRAPH Comput. Graph.*, 26(2):255–264, 1992.
- [Wei86] Jerry Weil. The synthesis of cloth objects. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '86, pages 49–54, New York, NY, USA, 1986. ACM.
- [WNO98] Lawrence B. Wolff, Shree K. Nayar, and Michael Oren. Improved diffuse reflection models for computer vision. *International Journal of Computer Vision*, 30:55–71, 1998. 10.1023/A:1008017513536.
- [WZT⁺08] Jiaping Wang, Shuang Zhao, Xin Tong, John Snyder, and Baining Guo. Modeling anisotropic surface reflectance with example-based microfacet synthesis. In *ACM SIGGRAPH 2008 papers*, SIGGRAPH '08, pages 41:1–41:9, New York, NY, USA, 2008. ACM.
- [XCL⁺01] Ying-Qing Xu, Yanyun Chen, Stephen Lin, Hua Zhong, Enhua Wu, Baining Guo, and Heung-Yeung Shum. Photorealistic rendering of knitwear using the lumislice. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 391–398, New York, NY, USA, 2001. ACM.
- [YYTI92] Takami Yasuda, Shigeki Yokoi, Jun-ichiro Toriwaki, and Katsuhiko Inagaki. A shading model for cloth objects. *IEEE Comput. Graph. Appl.*, 12:15–24, November 1992.
- [ZJMB11] Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph.*, July 2011.