



Recent Advances in Physically-Based Appearance Modeling of Cloth

Kai Schröder
Bonn University

Shuang Zhao
Cornell University

Arno Zinke
TeamUp Technologies



Figure 1: *Two pieces of cloth rendered using state-of-the-art volumetric techniques. Left: [Jakob et al. 2010a], Right: [Schröder et al. 2011b]*

Contents

1	Introduction	5
2	Background	8
2.1	Cloth Structure	8
2.1.1	Fibers	8
2.1.2	Yarns	8
2.1.3	Composition	9
2.2	Related Work	11
2.2.1	Cloth Appearance Models	12
2.2.2	Data-driven and Image-Based Cloth Rendering	13
3	Single-Fiber-Based Cloth	16
3.1	Light Scattering from Fibers	16
3.2	Efficient Single Fiber Scattering: BFSDF and BCSDf	17
3.3	Notation	18
3.4	Bidirectional Fiber Scattering Distribution Function — BFSDF	19
3.5	Far-Field Approximation and BCSDf	21
3.6	Explicit Rendering Methods	24
4	Volumetric Cloth	27
4.1	Volumetric Light Transport for Cloth	27
4.1.1	Transport Theory	27
4.1.2	Anisotropic Light Transport	28
4.2	Volumetric Models of Cloth	29
4.2.1	The Micro-flake Model	29
4.2.2	A Gaussian Mixture Model of Fibers	30
4.2.3	Modeling the Statistics of Self-Shadowing for Volumetric Models	31
4.2.4	Discussion of Statistical Models	33
4.3	Building Volumetric Models of Cloth	34
4.3.1	Using Micro CT Imaging	34
4.3.2	Synthesizing Models with Complex Designs	37
5	Practical Physically-Based Rendering of Cloth	39
5.1	Monte Carlo Path Tracing	39
5.2	Other Approaches	42
5.3	Bidirectional Texture Function Synthesis	45
6	Acknowledgements	47
A	Useful Transformations for Fiber Scattering	52

1 Introduction

Textiles are an essential component of most virtual scenes: The appearance of human avatars relies on realistic virtual clothing; other textiles such as carpets and curtains are common indoor elements. Apart from rendering plausible images for movies, virtual prototyping and design are common applications for cloth rendering. An enormous effort is currently spent on the design of seat cushions for cars. In several iterations, different variations of fabrics are produced until a final design decision is found. To shorten the process by performing some of these iterations only virtually, we want to design rendering algorithms that can faithfully predict the appearance of textiles. In such a system, it is desirable that domain specific parameters of cloth can be edited. The types of fibers should be exchangeable, differently spun yarns should be available and compositional elements such as weave-, and knit-patterns should be editable.

Modeling of materials has a long tradition in computer graphics. Despite much work in this domain, predictive rendering remains a challenging task even if the micro-geometry of a material is well-understood. This is especially true for cloth. Apart from *geometrical complexity*, *optical complexity* presents complications as highly anisotropic single and multiple scattering and self shadowing effects often dominate the appearance. Many types of fibers are highly translucent and multiple scattering significantly influences the observed color (see Figure 2). Since a cloth model may consist of potentially hundreds of millions of fibers, finding a viable level of geometrical abstraction is difficult.

For general material rendering applications, much work has focused on statistical surface-based models such as *Bidirectional Reflectance Distribution Functions* (BRDFs) and *Bidirectional Texture Functions* (BTFs). These offer a simplified representation of a material that is often sufficient for rendering from a certain distance. BTFs can be measured for real materials and rendering can be extremely efficient (even real-time) for simple lighting setups. Data-driven image-based approaches naturally incorporate effects such as correct handling of interface properties (e.g. water or oil) as well as colorizations for a measured sample. However, silhouette effects are commonly ignored, handling of transparency is difficult and correct diffusion of light at shadow boundaries is usually not possible. Apart from this, these approaches also fall short when the target application is the design of virtual materials.

Recent techniques that represent micro-geometry as fiber assemblies or in volumetric ways model more aspects of the material explicitly. This makes them attractive for predictive rendering in the context of *computer aided design* (CAD) of materials. Only when the three inherent scales of cloth (fibers, yarns, compositions) are represented explicitly, they can be edited directly. In contrast, when they are aggregated in a unified representation, editing can be restricted and difficult. In the first part of the course, we give an introduction to the structures that constitute the micro-geometry of typical types of cloth and explain how fibers are combined to form yarns and how yarns can be composited to form textiles (Section 2). By using domain specific knowledge about the structure of a material one can choose an appropriate representation for the different scales.

In the next section, we explain the optical properties of fibers, which constitute the finest scale. These properties can be derived from first principles of physics such as absorp-




Figure 2: *The effect of multiple scattering simulated for a patch of cloth. The lower left part of the image shows only local illumination. The upper right part shows a full global illumination solution. The final color is largely influenced by light scattering inside the yarn. Taken from [Schröder et al. 2011b].*

tion or index of refraction. Understanding light scattering from fibers is essential, when physically-based cloth rendering approaches are designed. Together with models of the micro-geometry, we obtain a full representation of a cloth material and can compute its appearance using global illumination rendering methods. Section 3 will be devoted to this topic. We present different abstractions to accelerate the required computations of multiple scattering within strands for explicitly modeled fiber assemblies. A problem of all of these approaches is, however, that even small patches of cloth easily need dozens of gigabytes in memory.

To overcome difficulties of exactly modeling the underlying micro-scale geometry while retaining most of the expressiveness of explicit techniques, different solutions have been suggested, among them volumetric methods. While early volumetric approaches for rendering cloth already deliver a very good overall impression of the appearance, they are not accurate enough for predictive rendering since the anisotropic nature of the micro-geometry is not modeled properly. Jakob et al. [2010a] recently devised a framework for volumetric modeling and rendering of materials with anisotropic micro-structure. Here, the local optical behavior of complex materials is modeled as a distribution of non-spherical particles that produce desired scattering distributions when considered as an aggregate. The resulting volumetric representation is then rendered using Monte Carlo integration, Finite Elements, or an anisotropic dipole model. This approach is very general and can describe many kinds of anisotropic structures well. It integrates perfectly into modern physics-based rendering systems. [Zhao et al. 2011] model cloth based on micro CT data and photographs and apply this framework for rendering.

Another volumetric approach has been presented by [Schröder et al. 2011b] who render woven materials using *virtual scattering events*: Instead of representing all fibers explicitly, only their statistical distribution is stored in a volume. Explicit intersection tests of rays with fiber geometry are replaced with a technique that computes locations of intersections by sampling from this statistical distribution. The approach permits using the same descriptions for optical properties of individual fibers as one can use for completely modeled micro-geometry. This can be advantageous in the context of editing.

A major part of this course will be devoted to volumetric approaches (Section 4). We first discuss the theory and practice of physically-based rendering of anisotropic media. The discussion begins with a brief review of classic volumetric modeling of isotropic media. Following this, we show how to extend the framework to anisotropic materials that consist of fibers. Relevant implementation details are discussed at each stage, and the result will be the pseudo code of a Monte Carlo path tracer for volumetric cloth representations. Finally, we give an overview of alternative rendering techniques for cloth, with different trade-offs for accuracy and speed.

In Table 1, we have summarized the major differences of several appearance models to aid in choosing the right approach. Combinations of the different algorithms may also be used, especially in the context of level of detail rendering. Although rendering of cloth is a specialized area, most of the approaches that are developed in this context can certainly be extended to describe other materials with complex micro-geometry. Cloth is currently an appealing research topic, since its micro-geometry is relatively well-understood, and powerful volumetric rendering techniques have recently become available. We believe that the versatility of these methods will make them a popular alternative when explicit representations are too costly.

	Surface Reflectance		Volumetric Rendering		Explicit Models
Technique	BSDF	BTF	Micro-Flake	Virtual Scatter.	Explicit
Translucency	✓	✓ ¹	✓	✓	✓
Silhouettes			✓	✓	✓
Light Diffusion			✓	✓	✓
Real-time	✓	✓			
Scalability	✓	✓	✓	✓	
Edit all scales				✓	✓
General	✓	✓	✓		
Integration-scale	Composition	Variable	Yarn / Fiber	Fiber	Fiber
Viewing Distance	far	medium	medium	medium	close

Table 1: Comparing the most common recent approaches that are used to render cloth.

¹Current measurement setups for BTFs often only capture one hemisphere and therefore, translucency is not always supported correctly.

2 Background

In the following, we provide relevant background information. We explain the structure of typical types of textiles and give an overview of different approaches that have been used to render cloth.

2.1 Cloth Structure

Cloth consists of three natural scales: fibers, yarns and yarn-compositions.

2.1.1 Fibers




Figure 3: *Electron micrograph of a hair fiber. Taken from [Robbins 1994]*

The micro-structure of fibers is described by the geometric and optical properties of the small dielectric fibers that comprise yarns – they are mainly determined by *absorption*, *refractive index* (e.g. index of refraction for wool 1.576, silk 1.35, polyester 1.53) and *cross sectional shape* (often close to circular e.g. for wool and many industrial fibers) with diameters of 10–100 μm (e.g. on average 17–42 μm for wool, 15 μm for silk, 13 μm for polyester) – values according to [Morton and Hearle 1962]. The geometry of most types of fibers used for textiles can hardly be distinguished from bent cylinders by the human eye. Their micro-geometry, e.g. the scales of wool hairs (c.f. Figure 3 showing a human hair fiber), however, can have a great effect on a rendered image, as these significantly alter the light transport in a scene. We describe the optical properties of fibers in more detail in Section 3.1.

2.1.2 Yarns

Filament yarns are produced by grouping or twisting a few hundred, long, continuous fibers (e.g. fibers that are taken from cocoons made by the larvae of the silkworm which can be hundreds of meters long).

Spun yarn is created by twisting hundreds of thousands of short staple fibers together to form a cohesive thread (see Figure 4 for close-up and cross-sectional views). This enables the production of yarns of arbitrary length even though individual fibers may be significantly shorter (e.g. cotton only has a typical length of 2–3 cm). Even stronger yarns are made up of a number of plies, where each ply is a single spun thread. Different variations exist concerning the direction in which fibers are twisted to form a ply and the direction in




Figure 4: From left to right: Close up of a typical woven cloth. (taken from [USDA E.R.S. 2004]), cross-sectional view of ring spun stable fibers and cross-sectional view of filament yarn (taken from [Sreprateep and Bohez 2006])




Figure 5: Illustration of several yarns created by different combinations of S- and Z-twist (taken from [Kistler 2011 (accessed February 2, 2012)]). In the right-most image you can see a photograph of a yarn which is made thicker by adding one more layer of plies.

which plies are twisted to form a yarn (Figure 5). These directions are called S- and Z-twist. There also exist yarns with even more levels of a twisting hierarchy (as in plies of plies...).

Usually small fibers protrude from a yarn – a property called *hairiness*. Different techniques exist to reduce this often unwanted effect. One of them is mercerization: Yarns created from cotton or hemp are subjected to chemical treatment to reduce hairiness and increase strength and luster. Additionally, several different colorization procedures and brighteners can be applied to significantly alter the optical appearance of yarns.

2.1.3 Composition

Many techniques have been developed to create pieces of cloth from yarns and fibers. Here we describe three of the most common ones: Woven cloth and two types of non-wovens called knitwear and felt.

¹Woven-image: CC-BY 2.0 Scott Robinson

²Knit-image: CC-BY 2.0 Chris Phan

³Felt-image: CC-BY 2.0 Siona Karen




Figure 6: *Different types of cloth, characterized by different yarn and fiber arrangements*




Figure 7: *Three common weave patterns are shown (illustration taken from [Irawan and Marschner 2006]).*

Woven Cloth: For single layer woven cloth (Figure 6(a)), the layout of yarns is described by a weave pattern. A number of parallel threads called *warp yarns* are orthogonally interlaced by another group of parallel threads called *weft yarns*. The weave pattern can be described by a 2-dimensional binary matrix, which defines for each crossing point whether the horizontally (i.e. warp) or vertically (i.e. weft) running yarn is visible from above. Figure 7 shows three of the most common types of weave patterns. More complex pieces of cloth can be created by using several layers.

Knitwear: In knitwear (Figure 6(b)), yarns are arranged in 3 dimensional loops. Two basic stitch types are combined to create different patterns (Figure 8). Starting with an initial row of loops, the loops of the next row are created by pulling through the loops of the upper row. One distinguishes between two types of stitches: When the yarn is pulled through a loop from the previous row from below, this is called a *knit* stitch. If it is pulled from above, it is called a *purl* stitch. The three most common ways to combine these stitches are

- the *stockinette* which only consists of *knit* stitches
- the *garter* where the two stitch types alternate row-by-row. This is the most common type in hand-knitting – in practice one can always use *knit* stitches but than one has to reverse the cloth sample after each row.

- the *1-1 rib* where the stitch types alternate within a row and the *2-2 rib* which alternates between two *knit* stitches and two *purl* stitches within a row.





Figure 8: Three standard knitting patterns are shown (illustration taken from [Kaldor et al. 2008]).

The direction along the length of a row is called the *course* direction. Orthogonal to that we have the stacking of rows in the *wale* direction.

Felt: This is one of the few types of textiles where fibers are directly combined instead of first creating yarns. As we have mentioned in Section 2.1.1, wool fibers consist of many scales. If many wool fibers are moisturized and subjected to friction, these scales make the fibers stick together. This results in a material of relatively homogeneous density (Figure 6(c)). Other types of fibers can also be mixed into the material as long as there are enough fibers with scales that stick to each other.

Finally, textiles can be created by cutting and sewing of cloth samples of various types. A description of that process is beyond the scope of this work.

2.2 Related Work



(a) [Adabala et al. 2003]

(b) [Daubert et al. 2001]

(c) [Irawan 2007]

Figure 9: Several methods modeling surface reflectance

Recently, Yuen et al. [2011] have presented a short survey of current cloth rendering techniques. Here, we also give a summary of the state-of-the art, to give a complete overview of cloth rendering approaches. A common way to render fabrics is to model the interactions of light with fibers and yarns using surface reflectance models. There are two main

directions: The first one uses heuristic methods of procedural modeling to create plausible images. The second one is data-driven and directly based on optical measurements.

2.2.1 Cloth Appearance Models

The most simple approach to modeling cloth is probably to use a combination of texture and normal mapping. However, this approach fails to model the anisotropic highlights and effects of multiple scattering and is nowadays only employed for very simple real-time applications. A first improvement is to store tangent directions of fibers in additional texture maps and use anisotropic shading models.


The early work by Yasuda et al. [1992] focuses on reproducing the characteristic gloss of cloth. A micro-facet model is proposed to model the appearance at distances where individual yarns cannot be distinguished. Daubert et al. [2001] model yarns using implicit surfaces and generate a BTF-like data representation using hardware rendering. They compute lighting using a geometrical model of a stitch. By sampling the stitch regularly within a plane, a view-dependent texture with per-pixel normals and material properties is generated. Adabala et al. [2003] use a simple BRDF model for efficient cloth rendering. They interpret the industry standard *weave information file format (WIF)* to generate reflectance data for arbitrary woven patterns. The appearance of the micro-geometry of spun yarn is visualized by creating a procedural texture with parameters that capture the tightness of twisting and thickness of the thread. An average bidirectional reflectance distribution function (BRDF) is estimated using a micro-facet model, which incorporates information from the WIF by analyzing the fraction of the surface which is visible as warp or weft threads. The amount of shadowing that a thread can cast onto neighboring thread facets is captured by a horizon map.

Ray-tracing based methods for modeling cloth at a fiber level have been presented in several works [Westin et al. 1992; Volevich et al. 1997]. Their main objective was to use micro-scale simulation to derive a BRDF/BSDF model for a small patch of cloth.


A more sophisticated approach to model the appearance of woven cloth using BRDF and BTF at yarn level was presented by Irawan et al. [2007]. The resulting models which are validated against measurements yield visually plausible results for a wide range of fabrics. Unfortunately, some of the model parameters are based on ad-hoc assumptions that cannot be directly inferred from optical fiber properties.

Sadeghi et al. [2011] present a shading model for cloth. They formulate a light scattering model for threads, which is validated against measurements of individual threads. A weave pattern is used to model the coverage of different thread types, their tangent directions and to obtain a shadowing and masking term. Results are compared to real cloth samples.

Volumetric Approaches Several volumetric techniques have been proposed to represent and render cloth. Early approaches were able to create good looking images of complex pieces of cloth. However, they were not based on physical properties and could therefore not be used for predictive rendering: A volumetric approach for modeling knitwear was proposed by Gröller et al. [1995]. By measuring the cross-sectional distribution of yarn fibers, a density field is created and swept along a three dimensional curve to form the entire



(a) [Groller et al. 1995]



(b) [Xu et al. 2001]

Figure 10: *Two images which have been created using early volumetric rendering techniques. While these methods can process complex pieces of cloth, the realism of renderings is limited, since all scattering by fibers is purely diffuse.*

yarn. A similar idea was presented by [Xu et al. 2001] where computations revolve around a structure called *lumislice*, a light field of a yarn-cross-section. A detailed description of recent volumetric cloth models is given in Section 4.



Figure 11: *Two renderings of BTFs under environment lighting. Left taken from [Meseth et al. 2003]. Right, image courtesy of Sensible Graphics GmbH; Material sample by Volkswagen AG*

2.2.2 Data-driven and Image-Based Cloth Rendering

[Sattler et al. 2003] physically acquire a BTF from a rectangular probe and generate view-dependent texture-maps using a principal component analysis of the original data. These maps can even reproduce mesoscopic detail and can then be illuminated and rendered in real-time on graphics hardware. Several variations of this technique have been investigated for realistic cloth rendering [Magnenat-Thalmann et al. 2004]. A data driven micro-facet-based BRDF approach was taken by Wang et al. [2008]. Here, for each measured surface point a normal density function model best fitting the observation is used for rendering

anisotropic spatially-varying materials, such as cloth. As BTFs are successfully used in many practical applications, we describe them in more detail.




Figure 12: 2-dimensional spatial slices through a BTF showing a material sample under several viewing and lighting directions.

Bidirectional Texture Functions (BTFs). BTFs [Dana et al. 1999] are view- and light dependent textures that describe the appearance of a material. Formally, it's a 6-dimensional function $(u, v, \phi_i, \omega_i, \phi_o, \omega_o) \rightarrow \mathbb{R}$, where the angles (ϕ_i, ω_i) represent the direction of an incident light field, which is transferred to an outgoing light field at the spatial position (u, v) inside the texture, with (ϕ_o, ω_o) being the direction of outgoing radiance. For spectral rendering, the BTF is parameterized by one additional dimension – the wavelength λ . Figure 12 shows spatial slices through a BTF, which show how the appearance of a material varies under different viewing and lighting directions. See [Müller et al. 2005; Haindl and Filip 2011] for excellent surveys on this topic. In the following we give a short overview:

BTF-measurement devices can capture the look of real existing materials. Occlusion and parallax effects, micro-geometry-based self-shadowing and global illumination effects — including local subsurface scattering — are incorporated naturally without the need for any explicit knowledge about these phenomena. This makes BTFs especially attractive if one wants to reproduce complex materials realistically in a virtual setting. They are currently one of the best techniques for reproducing the appearance of measured cloth samples.

One can also synthesize BTFs of new materials by building virtual gonio-reflectometer setups (see Section 5.3).

Additionally, BTFs are well suited for rendering applications. They are known to be efficiently compressible and under certain lighting conditions (i.e. under a few point lights) they can be used for real-time graphics rendering applications even on the web using WebGL [Schwartz et al. 2011].

Unfortunately, BTFs also have inherent limitations: Light diffusion is not modeled properly (especially at shadow boundaries), regions of high curvature are not represented correctly if the BTF is measured with regard to a flat sample, and silhouette information is missing. Handling of transparency is difficult and rarely considered.

Polynomial Texture Maps. Polynomial Texture Maps (PTMs) [Malzbender et al. 2001] represent a subset of a full BTF, fixing the exitant direction (using a single camera view). This makes it possible to store all information in a two dimensional multi-channel texture map. In addition to storing a color value per pixel, coefficients for polynomials are stored, approximating the variation of luminance for different light directions (a common variation of the approach is to store different coefficients for each RGB channel). Although the simple polynomial approximation cannot reproduce all variations of cloth appearance faithfully, the technique may be appealing for some applications due to its simplicity.

3 Single-Fiber-Based Cloth

An alternative to image-based approaches is provided by geometric models of micro-geometry in conjunction with global illumination rendering. Micro-geometry of cloth can for example be synthesized procedurally. Suitable parameters for the optical properties of fibers have to be chosen (for example based on measurements [Zinke et al. 2009]) and the light/surface interaction has to be simulated. This approach offers the advantage that a model-driven architecture can allow for intuitive editing paradigms. Model parameters may be described in such a way that they resemble parameters that experts in the field are used to. In the context of cloth, these could for example be certain descriptions for weaving patterns, twists per centimeter for spun yarn or optical properties of fibers. In addition to purely procedural models, physically based simulation of deformations of yarns and fibers can greatly enhance the realism (this is a huge field on its own and a description of it is beyond the scope of this course).

3.1 Light Scattering from Fibers




Figure 13: *Left:* Light scattering from dielectric fibers for the strongest scattering components: surface reflection (R), transmission (TT) and back side reflection (TRT). Transmission and back side reflection are colored due to absorption inside the fiber. For fibers with a constant cross sectional shape, directional light is scattered to a specular cone (indicated by the red dotted cone), regardless of the component. *Right:* Characteristic anisotropic scattering effects of woven cloth caused by the specular cone (note the strong highlights)

Light scattering from individual filaments is the basis of physically accurate rendering of fiber assemblies such as cloth. If the internal structure of a filament is known a full geometric representation of individual filaments is possible in principle.

The appearance of fibers is typically dominated by two very prominent features: strong, distinctive anisotropic highlights and translucency. Both effects are well understood and can be explained by analyzing (cross sectional) shape and optical material parameters of the fibers interior. For homogeneous dielectric fibers (the common case) scattering can be modeled by considering only a few components (cf. Figure 13 Left). As first shown by Marschner et al. [2003] in his seminal work on hair rendering, light scattering can be well approximated by a direct surface reflection component (R), light that gets transmitted through the fibers (TT) and back-scattered light that got internally reflected (TRT). All these modes are highly anisotropic with respect to the fibers' tangent since directional incident

lighting is scattered to a cone centered around the tangent direction. This effect is causing the characteristic anisotropy of woven cloth shown in Fig. 13 Right.

However, even though thinkable in principle, geometrically modeling the exact shape of a fiber is not practical with regard to memory consumption and computation time. For that reason, simplified formalisms are highly desirable. Ideally, such methods would render the simulation of all internal light transport inside a strand unnecessary.

3.2 Efficient Single Fiber Scattering: BFSDF and BCSDf

In the realm of applied optics, light scattering from straight smooth dielectric fibers with constant, mainly circular or elliptical cross sections have been analyzed in a number of publications [Adler et al. 1996; Adler et al. 1998; Adler et al. 2001; Mount et al. 1998; Mees et al. 1998]. Schuh et al. [2003] also introduced an approximation that is capable of predicting scattering of electro-magnetic waves from curved fibers. All these approaches produce highly accurate results for individual strands, but are either more specialized in predicting peaks of the scattering distributions or impractical for rendering.

Over the last two decades BRDF/BSDF and BSSRDF have proven to be very successful concepts to describe light scattering from surfaces. Comparable concepts have also been presented for fiber optics [Marschner et al. 2003; Zinke and Weber 2007].

For a fiber, the relationship between exitant radiance scattered from a surface patch dA_o in direction ω_o and incident flux at surface patch dA_i from direction ω_i could be basically represented by a BSSRDF (Fig. 14 Left). Such a BSSRDF would in general depend on macroscopic deformations of the fiber, how the fiber is oriented and bent. As a consequence local (microscopic) fiber properties could not be separated from its global shape. Furthermore, all radiometric quantities would need to be parameterized with respect to the actual surface of the filament, which may not be well defined.

To make parameterization independent of shape, local minimum enclosing cylinder (Fig. 14 Right) is used. Essentially, radiance transfer at the fiber is locally approximated by a scattering function on the minimum enclosing cylinder of a straight fiber, which means a first order approximation with respect to the curvature. This function is called *Bidirectional Fiber Scattering Distribution Function* (BFSDF). The BFSDF is giving accurate results, if the effect of the global deformation (i.e. the curvature) on the scattering distribution is negligible, thus if

- the filament's curvature is small compared to the radius of its local minimum enclosing cylinder or
- most of the incoming light scatters only locally.

The latter condition is satisfied for instance for opaque wires, where no internal light transport occurs inside the filament. Here, scattering is a purely local phenomenon and the curvature of the fiber plays no role at all. On the other hand, for example, the first condition holds in case of hair, fur and for many types of synthetic fibers. In this case internal




Figure 14: Modeling scattering from a fiber with a pentagonal cross section by: **Left:** a BSSRDF at the actual surface of the fiber. **Right:** a BFSDF at the local minimum enclosing cylinder.

light transport may not be neglected, but since the curvature is small compared to the radius of the fiber and since the light gets attenuated due to volumetric absorption, the error introduced by this approximation is still acceptable.

3.3 Notation

Before technically defining the BFSDF we will give an overview of our notation. The local filament axis (tangent) is denoted $\vec{u}$ and we refer to the planes perpendicular to $\vec{u}$ as *normal planes*. All local surface normals lie within the local normal plane. All incoming and outgoing radiance is parameterized with respect to the local minimum enclosing cylinder of radius r oriented along $\vec{u}$ and is therefore independent of the actual cross section geometry of the filament (Fig. 15 Left). For a ray with direction $\vec{\omega}$ which is intersecting the minimum enclosing cylinder at a point $\vec{X}$, the infinitesimal surface patch centered around $\vec{X}$ will be denoted dA .

We introduce two sets of variables. The first parameter set describes $\vec{X}$ by its position s along the tangent and its azimuthal position ξ within the normal plane (cylindrical coordinates). The direction $\vec{\omega}$ of a ray intersecting the cylinder is then given by two spherical angles α and β , which span a hemisphere at $\vec{X}$ oriented along the local surface normal of the enclosing cylinder (Fig. 15 Right, Fig. 16 Right).

Parameterizing scattering this way is intuitive but not suitable for further approximations. Therefore, a second set of variables according to Marschner et al. [2003] is used for the course. Here, parameterization of azimuthal scattering within the normal plane (h, φ) is separated from longitudinal scattering along the tangent (s, θ) . The angle θ denotes the inclination of $\vec{\omega}$ with respect to the normal plane, φ its azimuthal angle with respect to a fixed axis $\vec{v}$ and h an offset position of $\vec{X}$ at the minimum enclosing cylinder, measured within the normal plane. The variable s is the same as introduced before and means the position along the axis of the minimum enclosing cylinder (Fig. 16). In Appendix A, we derive transformations between variables of both sets.




Figure 15: Left: An “intuitive” variable set for parameterizing incoming and outgoing radiance at the minimum enclosing cylinder. The vector $\vec{u}$ denotes the local tangent (axis), s the position along this axis. The angles α and β span a hemisphere over the the surface patch at $\vec{X}$, with α being measured with respect to $\vec{u}$ within the tangent plane and β with respect to the local surface normal $\vec{n}$. **Right:** Azimuthal variables of the first and the second parameter set. The parameter h is an offset position at the perimeter. Note that parallel rays have same φ (measured counterclockwise with respect to $\vec{u}$) but different h .

3.4 Bidirectional Fiber Scattering Distribution Function — BFSDF

Similarly to the BSSRDF/BSDF for surfaces, light scattering from fibers can also be approximated using bidirectional scattering functions. To capture near-field effects (e.g. to resolve detail across the width of a strand) the *Bidirectional Fiber Scattering Distribution Function* (BFSDF) was introduced by Zinke et al. [Zinke and Weber 2007]. As discussed earlier this function represents the light transport at the local minimum enclosing cylinder surrounding the fiber. Like the BSSRDF, this six dimensional function — for a fixed wave-length — relates incident flux approaching a point on this cylinder to the outgoing radiance from another point on the cylinder surface.

We now technically define the **Bidirectional Fiber Scattering Distribution Function** (BFSDF) of a fiber. Please note that for further considerations all radiometric quantities at the minimum enclosing cylinder are implicitly transformed by a transformation that unwraps the cylinder to a plane. By applying such a transformation the local surface normals of the cylinder become equally aligned. As a consequence the angles α and β can be expressed with respect to a global coordinate system.

The BFSDF relates incident flux Φ_i at an infinitesimal surface patch dA_i to the outgoing radiance dL_o at another position on the minimum enclosing cylinder:

$$f_{\text{BFSDF}}(s_i, \xi_i, \alpha_i, \beta_i, s_o, \xi_o, \alpha_o, \beta_o) := \frac{dL_o(s_i, \xi_i, \alpha_i, \beta_i, s_o, \xi_o, \alpha_o, \beta_o)}{d\Phi_i(s_i, \xi_i, \alpha_i, \beta_i)}. \quad (1)$$

Based on the BFSDF the local orientation of a fiber and the incoming radiance field dL_i ,




Figure 16: Left: Parameterization with respect to the normal plane and filament axis (tangent) $\vec{u}$. The vector $\vec{w}'$ denotes the projection of a direction $\vec{w}$ onto the normal plane. The angle θ ranges from $-\pi/2$ to $\pi/2$, with $\theta = \pi/2$ if $\vec{u}$ and $\vec{w}$ are pointing in the same direction. **Right:** Parameterization of the BFSDF.

the total outgoing radiance of a particular position can be calculated by integrating the irradiance over all surface patches and all incoming directions as follows:

$$L_o(s_o, \xi_o, \alpha_o, \beta_o) = \int_{-\infty}^{+\infty} \int_0^{2\pi} \int_0^{2\pi} \int_0^{\pi/2} f_{\text{BFSDF}}(s_i, \xi_i, \alpha_i, \beta_i, s_o, \xi_o, \alpha_o, \beta_o) L_i(s_i, \xi_i, \alpha_i, \beta_i) \sin \beta_i \cos \beta_i d\beta_i d\alpha_i d\xi_i ds_i. \quad (2)$$

If the optical properties of a fiber are constant along s , then the BFSDF does no longer depend on s_i and s_o but only on the difference $\Delta s := s_o - s_i$. Thus the dimension of the BFSDF function decreases by one and the scattering integral simplifies to:

$$L_o(s_o, \xi_o, \alpha_o, \beta_o) = \int_{-\infty}^{+\infty} \int_0^{2\pi} \int_0^{2\pi} \int_0^{\pi/2} f_{\text{BFSDF}}(\Delta s, \xi_i, \alpha_i, \beta_i, \xi_o, \alpha_o, \beta_o) L_i(s_i, \xi_i, \alpha_i, \beta_i) \sin \beta_i \cos \beta_i d\beta_i d\alpha_i d\xi_i d\Delta s. \quad (3)$$

This special case is also an appropriate approximation for the general case if the following condition is satisfied:

- The cross section shape and material properties vary slowly along s or
- have only very high frequency detail on a scale that does not need to be resolved.

Good examples for high frequency detail are cuticula tiles of hair fibers or surface roughness due to the microstructure of a filament. Since this assumption is at least locally valid for a wide range of different fibers, we will restrict ourselves to this case in the following.

The BFSDF and its corresponding rendering integral was introduced with respect to the “intuitive variable set” first. We now derive a formulation for the second variable set. Reparameterizing the scattering integral yields:

$$\begin{aligned}
L_o(s_o, h_o, \varphi_o, \theta_o) &= \int_{-\infty}^{+\infty} \int_{-1}^1 \int_0^{2\pi} \int_{-\pi/2}^{\pi/2} \left| \frac{\partial(\xi_i, \alpha_i, \beta_i)}{\partial(h_i, \varphi_i, \theta_i)} \right| f_{\text{BFSDF}}(\Delta s, h_i, \varphi_i, \theta_i, h_o, \varphi_o, \theta_o) \\
&\quad L_i(s_i, h_i, \varphi_i, \theta_i) \sqrt{1 - \cos^2 \theta_i (1 - h_i^2)} \sqrt{1 - h_i^2} \cos \theta_i \\
&\quad d\theta_i d\varphi_i dh_i d\Delta s.
\end{aligned} \tag{4}$$

The above equation can be finally simplified to the following equation:

$$\begin{aligned}
L_o(s_o, h_o, \varphi_o, \theta_o) &= \int_{-\infty}^{+\infty} \int_{-1}^1 \int_0^{2\pi} \int_{-\pi/2}^{\pi/2} f_{\text{BFSDF}}(\Delta s, h_i, \varphi_i, \theta_i, h_o, \varphi_o, \theta_o) \\
&\quad \cos^2 \theta_i L_i(s_i, h_i, \varphi_i, \theta_i) d\theta_i d\varphi_i dh_i d\Delta s.
\end{aligned} \tag{5}$$

This form of the scattering integral is especially suitable when having distant lights, i.e. parallel illumination, since parallel rays of light have same φ_i and θ_i but different h_i . Consider for example directional lighting, which can be expressed by the radiance distribution $L(\varphi'_i, \theta'_i, \varphi_i, \theta_i) = E_{\perp}(\varphi'_i, \theta'_i) \delta(\varphi_i - \varphi'_i) \delta(\theta_i - \theta'_i) / \cos \theta_i$ with normal irradiance $E_{\perp}(\varphi'_i, \theta'_i)$. Substituting the radiance in the scattering integral by this distribution yields the following special case:

$$L_o(h_o, \varphi_o, \theta_o) = E_{\perp}(\varphi'_i, \theta'_i) \cos \theta'_i \int_{-\infty}^{+\infty} \int_{-1}^1 f_{\text{BFSDF}}(\Delta s, h_i, \varphi'_i, \theta'_i, h_o, \varphi_o, \theta_o) dh_i d\Delta s. \tag{6}$$

3.5 Far-Field Approximation and BCSDf

Fibers are very thin and long structures. Hence, compared to its effective diameter both viewing and lighting distances are typically very large. Therefore all surface patches along the normal plane have equal (φ_o, θ_o) and the fibers cross section is locally illuminated by



Figure 17: Close-up of an oval made of bluish plastic. **Left:** BFSDF rendering, **Right:** BCSDF approximation of the BFSDF rendering. Here, the exitant radiance is averaged across the width of the fiber. Note the detail in the BFSDF version which is missing in the BCSDF rendering.

parallel light (of constant radiance) from a fixed direction (φ_i, θ_i) . Furthermore adjacent surface patches of the local minimum enclosing cylinder have nearly the same distance to both light source and observer.

When rendering such a scene the screen space width of a filament is less or in the order of magnitude of a single pixel. In this case fibers can be well approximated by curves having a cross section proportional to the diameter of the minimum enclosing cylinder D .

That allows us to introduce a simpler *far-field* approximation. Using similar quantities as Marschner et al. [Marschner et al. 2003]—curve radiance and curve irradiance—we show how an appropriate scattering function can be derived right from the BFSDF. The *curve radiance* $d\bar{L}_o$ is the averaged outgoing radiance along the width of the fibers minimum enclosing cylinder times its effective diameter (Fig. 18).

Assuming parallel light from direction (φ_i, θ_i) being constant across the width of a fiber this curve radiance can be computed by averaging the outgoing radiance across the width:

$$d\bar{L}_o(\varphi_i, \theta_i, \varphi_o, \theta_o) := \frac{D}{\pi} \int_{\varphi_o - \pi/2}^{\varphi_o + \pi/2} \cos(\xi_o - \varphi_o) dL_o(\varphi_i, \theta_i, \xi_o, \alpha_o(\theta_o, h_o(\xi_o, \varphi_o)), \beta_o(\theta_o, h_o(\xi_o, \varphi_o))) d\xi_o. \quad (7)$$

Here a factor of $\cos(\xi_o - \varphi_o)$ accounts for the projected width of the minimum enclosing

cylinder. Notice that for the integration with respect to ξ_o one has to take into account its transition of from zero to 2π properly. Using the other set of variables this can be transformed to:

$$d\bar{L}_o(\varphi_i, \theta_i, \varphi_o, \theta_o) = \frac{D}{2} \int_{-1}^1 dL_o(h_o, \varphi_i, \theta_i, \varphi_o, \theta_o) dh_o. \quad (8)$$

Since we assume constant incident radiance $L_i(\varphi_i, \theta_i)$, dL_o may be substituted by:

$$\begin{aligned} dL_o(\varphi_i, \theta_i, h_o, \varphi_o, \theta_o) &= \cos^2 \theta_i L_i(\varphi_i, \theta_i) d\varphi_i d\theta_i dh_o \int_{-1}^1 \int_{-\infty}^{\infty} \\ &\quad f_{\text{BFSDF}}(\Delta s, h_i, \varphi_i, \theta_i, h_o, \varphi_o, \theta_o) d\Delta s dh_i \end{aligned} \quad (9)$$

which yields:

$$\begin{aligned} d\bar{L}_o(\varphi_i, \theta_i, \varphi_o, \theta_o) &= \frac{D}{2} L_i(\varphi_i, \theta_i) \cos^2 \theta_i d\varphi_i d\theta_i \int_{-1}^1 \int_{-1}^1 \int_{-\infty}^{\infty} \\ &\quad f_{\text{BFSDF}}(\Delta s, h_i, \varphi_i, \theta_i, h_o, \varphi_o, \theta_o) d\Delta s dh_i dh_o \end{aligned} \quad (10)$$

This scattered radiance $d\bar{L}_o$ is proportional to the incoming *curve irradiance* $d\bar{E}_i$:

$$d\bar{E}_i(\varphi_i, \theta_i) := DL_i(\varphi_i, \theta_i) \cos^2 \theta_i d\varphi_i d\theta_i. \quad (11)$$

Thus, a new bidirectional far-field scattering distribution function for curves f_{BCSDF} that assumes distant observer and distant light sources can be defined. It relates the incoming curve irradiance to the averaged outgoing radiance across the width (curve radiance). We will call this far-field approximation the **Bidirectional Curve Scattering Distribution Function** (BCSDF).

$$f_{\text{BCSDF}}(\varphi_i, \theta_i, \varphi_o, \theta_o) := \frac{d\bar{L}_o(\varphi_i, \theta_i, \varphi_o, \theta_o)}{d\bar{E}_i(\varphi_i, \theta_i)}. \quad (12)$$

With (10) and (11) we obtain:

$$f_{\text{BCSDF}}(\varphi_i, \theta_i, \varphi_o, \theta_o) = \frac{1}{2} \int_{-1}^1 \int_{-1}^1 \int_{-\infty}^{\infty} f_{\text{BFSDF}}(\Delta s, h_i, \varphi_i, \theta_i, h_o, \varphi_o, \theta_o) d\Delta s dh_i dh_o. \quad (13)$$

Finally, in order to compute the total outgoing curve radiance one has to integrate all incoming light over a sphere with respect to φ_i and θ_i :

$$\bar{L}_o(\varphi_o, \theta_o) = D \int_0^{2\pi} \int_{-\pi/2}^{\pi/2} f_{\text{BCSDF}}(\varphi_i, \theta_i, \varphi_o, \theta_o) L_i(\varphi_i, \theta_i) \cos^2 \theta_i d\theta_i d\varphi_i. \quad (14)$$

This equation matches the rendering integral given in [Marschner et al. 2003, equation 1] in the specific context of hair rendering. Thus the BCSDF is identical to the fiber scattering function very briefly introduced in [Marschner et al. 2003]. In contrast to [Marschner et al. 2003] our derivation shows the close connection between the BFSDF and the BCSDF. In our formalism the BCSDF is just one specific approximation of the BFSDF and can be computed directly from (13).




Figure 18: *Far-field assumptions: A slice at the minimum enclosing cylinder illuminated by parallel light and viewed by a distant observer.*

Besides of the advantage of being less complex than the BFSDF the BCSDF can help to drastically reduce aliasing artifacts. Moreover, efficient importance sampling is possible along the lines of BRDF.

3.6 Explicit Rendering Methods

The described fiber and curve scattering models form a solid basis for physically-based rendering techniques for cloth, in particular if combined with classical importance sampling techniques. There exists a strong connection between cloth rendering and rendering of

hair, as yarns consist of many individual fibers (in particular, the fibers are hair in the special case of wool yarns). Therefore, hair rendering algorithms could be applied in principle for rendering explicit models of cloth. Besides explicit path-tracing methods that compute both single and multiple scattering based on individual strands [Zinke et al. 2004], efficient approximations regarding multiple scattering have also been presented. All these approximations rely on the fact that the multiple scattering distribution in hair tends to be smooth with only little high frequency detail [Moon and Marschner 2006; Zinke et al. 2008; Moon et al. 2008]. Although this might be true for hair, this assumption is not valid for woven cloth where the micro structure exhibits much spatial variation that furthermore induces spatially varying single and multiple scattering effects.

For the hybrid approach presented by Moon et al.[2008] single scattering is computed from a single fiber based hair model whereas multiple scattering is estimated by tracing light rays through a volumetric approximation of that model. In a pre-processing step, the hair geometry is first voxelized and relevant fiber properties (i.e. the scattering coefficient, the average tangent of hair strands intersecting a voxel, the standard deviation with respect to tangent) are stored in a uniform grid data structure. Based on this statistical information, scattering events are generated during a subsequent light tracing phase that finally approximates multiple scattered light by using spherical harmonics basis functions.

As we will see in Section 4 taking a volumetric approach is essential to overcome limitations caused by geometrical complexity of larger cloth models.




Figure 19: *Left: Close-up view of a global illumination rendering of explicitly modeled micro-geometry. Right: Modeling each yarn as a single fiber with an appropriate BCSDF*

Limitations One problem of explicitly modeling fiber geometry is that up to billions of fibers are needed to create larger textiles. A straightforward solution to this problem could be instancing – however this only works in a limited amount of situations. Even for the case of a repeating weave pattern, the geometry of spun yarns usually does not repeat with the same frequency. This is due to the fact that cloth is a soft material that can be drastically deformed resulting in stretching and squeezing of yarns. Furthermore, *fancy yarns* may

include spatial variations themselves along their length leading to the specific look and feel of the corresponding cloth.

4 Volumetric Cloth

Surface-based models have produced many nice results for cloth. However, they lack the ability to capture the thickness and fuzziness of the material, which are particularly important when viewing from a medium or close distance. To better capture such effects, we present several volumetric models. They treat a piece of cloth as a volume and model light scattering explicitly.

4.1 Volumetric Light Transport for Cloth

In this section, we discuss the theory and practice of physically-based rendering of anisotropic media. Because it is instructive to do so, we begin our discussion with a review of linear transport theory from hydrologic and atmospheric optics, upon which current methods for rendering volumetric cloth are based.

Following this, we show how to extend the framework to anisotropic materials that consist of fibers. Relevant implementation details are discussed at each stage, and the final result will be the pseudocode of a Monte Carlo path tracer for volumetric cloth representations.

4.1.1 Transport Theory




Figure 20: For light propagating along a direction ω , the radiative transfer equation describes the balance between absorption, out-scattering (combined into a quantity referred to as extinction) and in-scattering.

Linear transport theory [Case and Zweifel 1967] provides a convenient mathematical framework for describing the propagation of light through a *random medium*, a statistical interpretation of the aggregate scattering behavior of small particles that fill a region of space.

The behavior of light within such a medium is governed by absorption and scattering, and these interactions are made precise by the radiative transfer equation

$$(\omega \cdot \nabla)L(\mathbf{x}, \omega) = -\sigma_t L(\mathbf{x}, \omega) + \sigma_s \int_{S^2} f_p(\omega' \rightarrow \omega) L(\mathbf{x}, \omega') d\omega' + Q(\mathbf{x}, \omega) \quad (15)$$

Here, $L(\mathbf{x}, \omega)$ is the radiance traveling along the ray $(\mathbf{x}, \omega)$, f_p is a spherical scattering distribution analogous to a BSDF and referred to as *phase function*⁴. The two constants σ_s

⁴The name “phase function” is historically due to the phases (i.e. the varying brightness) of celestial bodies in astronomy.




Figure 21: *The distinction between isotropic and anisotropic media.*

and σ_t are known as scattering and extinction coefficients, respectively, and Q is a source term that is usually set to zero when the medium in question does not emit light.

Inspecting the terms of (15) yields an intuitive interpretation of their meaning (Figure 20): in particular, the equation relates the directional derivative (i.e. the increase or decrease) of radiance at $\mathbf{x}$ in some direction ω to a sum of several effects.

The first accounts for light that is blocked due to the cross-section of the particles that make up the medium; their density is proportional to σ_t . The particles also scatter a portion of the illumination passing through $\mathbf{x}$, and roughly speaking, some of this light will be scattered in direction ω . This is the second term, which integrates the radiance and phase functions over the sphere to obtain all light that is locally “added” to the ray $(\mathbf{x}, \omega)$.

In the specification of a random medium, the main objects of interest are thus:

- (i) the extinction cross-section σ_t , which can be interpreted as the inverse expected distance between interactions with particles (units of m^{-1}). The higher this number, the more “optically dense” the material will be.
- (ii) the scattering albedo σ_s/σ_t , i.e. the relative amount of scattering (unitless).
- (iii) the phase function $f_p(\omega_i \rightarrow \omega_o)$, a conditional probability distribution over scattered directions ω_o for a given incident direction ω_i (units of $1/sr$)

A straightforward extension of this equation introduces heterogeneity by allowing σ_s , σ_t , and f_p to depend on the position $\mathbf{x}$.

The classical approach for solving this kind of equation is direct numerical integration based on ray marching. Here, the basic idea is to accumulate radiance approaching the eye along rays. However, this method is only suitable for optically thin media with simple isotropic phase functions. As will be discussed in Section 5, an alternative technique, random walk sampling, is more efficient in our case.

4.1.2 Anisotropic Light Transport

Traditional volume light transport is limited to isotropic media and are inherently unable to describe the directional scattering needed to achieve realistic appearance in nontrivial solid materials. To upgrade the underlying framework of participating media to handle anisotropic effects, such as specular reflections from cloth fibers, Jakob et al. [2010a] proposed a theoretical framework for *anisotropic light transport* that allows the machinery of physics-based rendering to be brought to bear on a wider range of volume models.

The standard radiative transfer model is limited to isotropic media – media whose properties are invariant to rotation. For instance, a cloud of steam filled with spherical water droplets fits within the standard model; or a cirrus cloud filled with randomly oriented ice crystals: either medium looks the same to light traveling in different directions. But a volume filled with aligned ice crystals; or hair, wood, or cloth fibers; or fibrous tissue like muscle or collagen; does not fit the usual model because its properties depend on orientation. This is the type of medium we call anisotropic (Figure 21).

In order for (15) to describe an isotropic medium, f_p must be a function of two directions independently: $f_p(\omega' \rightarrow \omega)$. Furthermore σ_s and σ_t also depend on ω , leading to the anisotropic RTE:

$$(\omega \cdot \nabla)L(\omega) = -\sigma_t(\omega)L(\omega) + \sigma_s(\omega) \int_{S^2} f_p(\omega' \rightarrow \omega) L(\omega') d\omega' + Q(\omega) \quad (16)$$

(In this equation all quantities depend on position, so this dependence is left implicit for compactness.) Although this generalization is superficially straightforward, subtle complications with normalization and reciprocity must be resolved. For more details, please refer to [Jakob et al. 2010a]. Accurately solving (16) requires costly Monte Carlo simulations. For better performance, we can take its first order approximation, the *anisotropic diffusion equation*, which turns out to be quite similar with the standard isotropic version and thus can be solved using finite element method or an extended version of the dipole diffusion solution [Jensen et al. 2001a].

4.2 Volumetric Models of Cloth

4.2.1 The Micro-flake Model

Based on anisotropic light transport framework, the micro-flake model represents different volume scattering interactions using a directional *flake distribution* $D(m)$ that describes the orientation m of idealized mirror flakes at every point in space. Similar to micro-facet models, the phase function then involves evaluation $D(m)$ at the half-way direction between the incident and outgoing direction:

$$\begin{aligned} \sigma_t(\omega) &= a\rho \int_{S^2} |\omega \cdot m| D(m) dm \\ \sigma_s(\omega) &= \alpha\sigma_t(\omega) \\ f_p(\omega' \rightarrow \omega) &= \frac{a\rho\alpha}{4\sigma_s(\omega)} \left(D(h(\omega, -\omega')) + D(-h(\omega, -\omega')) \right) \end{aligned}$$

Here, ρ denotes the particle density, a is the area of a single flake, α is the associated albedo, and $h(\omega, \omega') := (\omega + \omega') / \|\omega + \omega'\|$. Note that the above expressions are simplified by assuming the flakes have albedo independent of the scattering angle. This reduces our search space considerably and still leads to a model that can represent scattering interactions with a variety of fibrous materials reasonably well.

To simulate scattering from a rough fiber with direction ω_f , we can set $D(\omega) = c_0 \sin^p(\omega_f, \omega)$, where higher values of p correspond to smoother fibers and c_0 is a normalization constant. This model leads to flake normals concentrated near the plane perpendicular to ω_f .

One serious drawback of the $\sin^p$ -type distribution is that most integrals over it do not have a closed form. This is problematic, since it effectively prevents the use of the inversion method for generating random samples distributed according to D . An alternative flake distribution that is convenient to integrate, while capturing the same key feature of the $\sin^p$ model, is the following truncated Gaussian centered around the great circle perpendicular to ω_f :

$$D(\omega) = \frac{1}{(2\pi)^{3/2} \gamma \operatorname{erf}\left(\frac{1}{\sqrt{2}\gamma}\right)} \exp\left(-\frac{(\omega_f \cdot \omega)^2}{2\gamma^2}\right)$$

where the standard deviation γ determines the roughness of the fiber and ω_f denotes the fiber direction.

4.2.2 A Gaussian Mixture Model of Fibers

An alternative to the micro-flake model is a more direct representation of the distribution of fibers based on fiber orientations and densities. Again, we are interested in the expected distance between scattering events – however in this case these are not interactions between particles but directly between fibers. Instead of computing intersections with actual fiber geometry, as would be the case for rendering explicit models (Section 3.6), virtual fiber locations and scattering events are computed by sampling from this statistical distribution. Light scattering at the fiber-level is described by a preintegrated curve scattering model (such as a BCSDf).

Given a set of n yarns, each consisting of several hundreds of fibers, we compute the statistical volumetric representation as follows: For the i -th yarn intersecting a voxel cell V , a tuple g_i representing Gaussian directionality, density and material properties is created:

$$g_i = \{\mathcal{N}_i(\vec{m}_i, s_i^2), \rho^i, \text{BCSDf}_i\}. \quad (17)$$

Besides the average tangent direction $\vec{m}_i$ and its associated standard deviation s_i (forming a Gaussian $\mathcal{N}_i(\vec{m}_i, s_i^2)$), a material index BCSDf_i is stored. In addition, the mean free path length inside a voxel, represented by the *effective density* ρ^i , is required to model fibers statistically. All $\mathcal{N}_i$ of a voxel cell, constitute components of a Gaussian mixture model. For the sake of simplicity we will in the following refer to the whole tuple g_i as a Gaussian mixture component of a mixture model $G_V = \{g_1, \dots, g_n\}$ where n equals the number of yarns inside voxel V .

The actual scattering coefficient, which is a directionally dependent quantity, is approximated based on ρ^i during rendering along the lines of Moon et al. [2008].

Fiber density. The mean free path length is calculated based on the perpendicular attenuation coefficient $\sigma_s^\perp = 2r\rho$ with fiber radius r . This value describes the inverse expected distance between scattering events for light propagating perpendicular to parallel fibers.

In the following, we will assume that a set of n line segments $\{L_j, j = [1..n]\}$ is used to represent the cloth. Then the *effective fiber density* can be computed by summing contribu-

tions from all line segments intersecting a given voxel:

$$\rho^{\text{voxel}} = \sum_j \frac{l_j}{\tilde{l}_j(\vec{t}_j) A_{\text{voxel}}(\vec{t}_j)}. \quad (18)$$

Here, $A_{\text{voxel}}(\vec{t}_j)$ denotes the area of a voxel projected onto a plane perpendicular to the tangent direction $\vec{t}_j$ of a given line segment, l_j is the intersection length of the j -th line segment, and $\tilde{l}_j$ is the intersection length of an infinite line extending this line segment. To speed up computations, tabulated values of $A_{\text{voxel}}(\vec{t}_j)$ are used during voxelization.

The effective scattering coefficient The effective scattering coefficient not only depends on fiber density but also on tangent direction t_j . For a single fiber indexed by j , the effective scattering coefficient σ_{single}^j with respect to a given direction $\vec{d}$ computes as

$$\sigma_{\text{single}}^j(\alpha_j) = 2r \rho_{\text{single}}^j \sin(\alpha_j) \quad (19)$$

where ρ_{single}^j means the fiber density induced by the (single) fiber of radius r and $\alpha_j = \angle(\vec{d}, \vec{t}_j)$.

Now let g_i be the i -th Gaussian mixture component representing a bundle of n_i fibers. Then the total scattering coefficient σ^i for the mixture component $\mathcal{N}_i$ may be computed by summing the contributions of the individual fibers:

$$\sigma^i = \sum_{j=1..n_i} \sigma_{\text{single}}^j(\alpha_j). \quad (20)$$

Finally, assuming $\mathcal{N}_i$ with average tangent direction $\vec{m}_i$ and a standard deviation s_i , σ^i may be approximated based on the fiber density ρ^i :

$$\sigma^i(\beta_i, s_i) \approx 2r \rho^i \int \sin(\alpha) \mathcal{N}_i(\alpha - \beta_i, s_i^2) d\alpha$$

where $\beta_i = \angle(\vec{d}, \vec{m}_i)$. To avoid computationally costly integration during rendering, the integral is precomputed. The time required for such pre-computation is negligible. To obtain the combined scattering probability for all mixture components, σ_s is computed by summing over all scattering coefficients $\sigma_s = \sum_i \sigma_s^i$.

4.2.3 Modeling the Statistics of Self-Shadowing for Volumetric Models

Taking a volumetric approach to model micro-geometry, self-shadowing details below the size of a single voxel cannot be captured. One might have the intuition that this is only a problem of scale and that finer voxelizations automatically lead to a better representation of self-shadowing, however this intuition is misleading, unless such a fine voxelization is employed that self-shadowing inside voxels can be ignored. A related problem is known in the context of shadow mapping where a common solution is to simply add an offset into the direction of the surface normal to avoid shadowing artifacts. This is not an appropriate




Figure 22: *Illustrating the shadowing issues when representing a yarn cross section volumetrically – each circle represents a single fiber. Yellow lines represent eye-rays – for each eye ray, 4 exemplary configurations of light-rays are shown. Green light-rays are never blocked, blue ones occasionally – pink and red ones illustrate cases of light coming from behind the cloth – the former are blocked locally in Voxel 1 whereas the latter are only blocked globally in Voxel 2 and are therefore regarded as not being blocked for the BVDF. In the volumetric case, all light-rays have a probability > 0 of being blocked inside voxel 1. To approximate shadowing correctly, these directional effects are modeled by the BVDF.*

solution in this case however, as an offset would sometimes have to be as large as the diameter of a single voxel and this would significantly alter the transported energy if the cross section of a yarn is covered by only very few voxels.

In particular, the directionality of local self-shadowing gets lost. However, this effect is critical for the appearance, especially for direct illumination. Consider the following example: A head light is located at the same position as the camera – all points that are visible from the camera should also be lit by the light. If the voxel is half full, there is a high probability for virtual scattering events along shadow rays in any direction and shadowing is significantly overestimated. There actually always is a non-vanishing probability for scattering for every non-empty voxel.

Moreover, as for any volumetric model, discontinuities at interfaces between optically dense and sparse regions are not well approximated (see Fig. 28 Right). This is in particular true for densely woven cloth (the common case) that behaves optically very similar to a solid model with a well-defined boundary.

To compensate for the above limitation of purely density-based representations, statistics for angularly dependent occlusion in case of direct illumination can be used for approximating self-shadowing at a local level.

Bidirectional Visibility Distribution Function (BVDF). Schröder et al [2011b] have proposed to model the correlation between eye-rays from direction ω_o and shadow-rays into direction ω_i by introducing the concept of *local visibility*.

The fiber scattering equation equation for direct lighting is given by:

$$L_{o,d}(x, \omega_o) = \int_{S^2} f(x, \omega_i, \omega_o) L_d(x, \omega_i) V(x, \omega_i, \omega_o) \sin(\alpha) d\omega_i$$

with BCSDf f , scattered radiance $L_{o,d}$, incident radiance distribution L_d due to direct lighting, visibility V and $\alpha = \angle(\omega_i, u)$. V is split into a local part V_l which accounts for shadowing due to occlusion inside a voxel cell and a global term V_g which models occlusion outside that cell:

$$V(x, \omega_i, \omega_o) = V_l(x, \omega_i, \omega_o) V_g(x, \omega_i, \omega_o)$$

Generally, V_l is a spatially varying quantity. However, assuming voxels with an extent in the order of the size of a yarn’s cross section, local visibility is modeled by the *Bidirectional Visibility Distribution Function* V_{BVDF} , a material property of yarns given by averaging V_l inside a voxel cell:

$$V(x, \omega_i, \omega_o) = V_{\text{BVDF}}(\omega_i, \omega_o) V_g(x, \omega_i, \omega_o)$$

Note that V_l and V_g are binary, whereas $V_{\text{BVDF}} \in [0, 1]$. In contrast to V_g , that is estimated using conventional shadow rays with virtual scattering events (see Sec. 5.1), tabulated values are used for V_{BVDF} . To this end the average visibility inside a voxel is computed for multiple viewing and lighting direction pairs in a pre-processing pass: A small planar patch of cloth is generated using the given yarn properties. We divide the space into cells of twice the size (see [Schröder et al. 2011b] for details) we plan to use for octree leaf voxels of the whole piece of cloth. The proportion of local shadowing is then computed using a fiber based rendering system for all cells hit by eye-rays. The resulting smooth function is stored by a 4D table (with $32 * 64^3$ bins in this case). Due to the small size of the sample (containing just a few yarns) this pre-computation is no bottleneck of the method and typically takes less than two minutes to compute for each yarn type.

The general structure of BVDFs always looks similar: When light and camera position are next to each other (green rays in Fig. 28), $V_{\text{BVDF}} \approx 1$, conversely when light and camera are on opposing sides of the cloth (black rays), V_{BVDF} is small. Moderate deviations in cross section shapes and densities of yarns have hardly any impact on these general directional effects – they only result in subtle changes.

In [Schröder et al. 2011b], the BVDF is used only in case of direct illumination to avoid distracting shadowing artifacts that otherwise would occur. For all other shadow-rays, which are calculated during multiple-scattering, occlusion is computed based on virtual scattering events along the ray according to Sec. 5.1. Once rays have entered the cloth, the actual shadowing is less critical and may be computed using the statistical model. Although describing visibility only for points hit by eye-rays is not unreasonable, as mainly effects at the surface are described, a bias (which mostly results in darkening) is introduced in case of multiple scattering. Shadow-rays that are shot during multiple scattering having an origin near the surface might profit from a BVDF representation as well.

4.2.4 Discussion of Statistical Models

Limitations Due to the assumptions made, the approaches have limitations compared to methods explicitly modeling individual yarn fibers. First of all, only effects on a scale bigger than the size of a voxel can be resolved properly. Thus, if the voxel’s extent is not chosen appropriately, bias occurs that causes spatial and angular blurring as well as shadowing artifacts. Naturally, artifacts become more evident in case of extremely curved

yarns, very directional lighting and for very specular fibers. However, it is important to recognize that these are the practical limitations of any approach not explicitly modeling fine scale geometry. The methods are less effective in case of extreme closeups where individual yarn fibers become visible as a high spatial resolution is required to resolve fine geometric detail.

The micro-flake model is a very general description for anisotropic media and can capture the basic phenomena of surface-like, fiber-like, and point-like scattering. However, there is no generic way of representing arbitrary scattering functions by micro-flake distributions and many relevant effects, such as Fresnel reflectance, cannot be properly modeled. In addition, such volumetric models consume significant memory since they are augmented with spatially varying vectors for anisotropy. Finally, evaluating anisotropic models consumes more computation time than for example looking up scalar parameters (as is the case for virtual scattering).

The micro-flake model describes both scattering within a fiber and between fibers using the same framework, whereas virtual scattering handles both kinds of interactions differently and thus follows a more domain specific approach. An advantage, however, is that one can directly make use of conventional fiber scattering functions, such as the BCSDF. This offers the possibility to use optical material properties of real individual fibers, which are either provided by manufacturers or can be measured efficiently [Zinke et al. 2009]. Therefore, one can clearly separate the modeling of yarns geometry from the modeling of material properties of fibers. The whole process of computing virtual scattering can also be seen as sampling a voxels phase function (which stays implicit) “on the fly”. A major disadvantage of this model is that we have to design new statistical geometry descriptions if something other than a fiber assembly should be represented. Also, the micro-flake model can easily scale to resolve the exact shapes of individual fibers, whereas the Gaussian mixture model and BCSDF descriptions only really make sense for scales larger than a fiber.

4.3 Building Volumetric Models of Cloth

In this section, we introduce some recent research that builds high-quality volumetric appearance models for cloth with minimal manual efforts.

4.3.1 Using Micro CT Imaging

In the previous sections, we have seen that the widely varying appearance of cloths are mainly driven by different small-scale structures. Furthermore, the volumetric appearance models can be derived directly from the structural information which used to be obtained by writing programs that create synthetic curves reproducing the weaving or knitting structures.

Zhao et al. [2011] introduced a technique that takes micro CT scans of cloth and produces micro-flake models. This technique begins with a scan of a small area of material, showing detail at the level of individual fibers over a fraction of a square centimeter. In a sequence of three stages (Figure 23), this data is processed and augmented, ending with a volume that defines the required scattering model parameters using density and orientation fields




Figure 23: *The volume appearance modeling pipeline by [Zhao et al. 2011]. (a) CT images are acquired, (b) the density field and orientation field of the volume are created, and (c) optical parameters of the volumetric model are assigned by matching statistics of photographs with rendered images. (d) Larger models are rendered using the acquired volumetric appearance and geometry models.*

derived from the CT data, plus three global parameters: the albedo, the lobe width, and a density multiplier that scales the density field.

CT Image Processing Micro CT devices, which use X-ray CT methods to examine small to microscopic structures, are increasing in availability, and this imaging modality is suited to a wide range of materials from which a small sample can be extracted for scanning. Next we describe the process of extracting fiber orientation from the CT density volume using a special fiber-detecting filter. Following this, we explain the processing steps needed to obtain orientation and density fields suitable for rendering.

CT images provide a voxelized density field with no direction information. Since the fiber scattering model requires an orientation for the phase function, it is necessary to reconstruct an orientation for every non-empty voxel. This approach uses oriented filters to detect fibers, based on similar filters used by Shinohara et al. [2010] to locate fibers in CT data.

To detect a fiber with orientation $\mathbf{d}$ at location $\mathbf{p}$, Shinohara proposes a cylindrically symmetric filter oriented with the axis $\mathbf{d}$, consisting of a difference of Gaussians in distance from the axis:

$$q(\mathbf{d}; \mathbf{p}) := -2 \exp(-sr^2) + \exp(-tr^2)$$

where $r = \|\mathbf{p} - (\mathbf{p} \cdot \mathbf{d})\mathbf{d}\|$ is the distance from the filter's axis and the parameters s and t (normally $s < t$) are empirically adjusted based on the size of the fibers present in the sample (see Figure 24).

The raw CT volume is thresholded at a value ϵ_d , resulting in a binary volume f :

$$f(\mathbf{x}) := \begin{cases} 0 & \text{CT}_{\text{raw}}(\mathbf{x}) \geq \epsilon_d \\ 1 & \text{CT}_{\text{raw}}(\mathbf{x}) < \epsilon_d \end{cases}$$

Then f is convolved with the filter q for each of a fixed set of orientations:

$$J(\mathbf{x}, \mathbf{d}) := \sum_{\mathbf{p} \in V} q(\mathbf{d}; \mathbf{p}) f(\mathbf{x} + \mathbf{p})$$

where V is a cubic volume of edge length h .




Figure 24: Computing function J in 2D: (a) shape of the filter q ; (b) when q is aligned to the fiber; (c) when q is unaligned.




Figure 25: (a) Renderings of a cylinder tiled with the satin volume, with fixed albedo and varying lobe width γ and density multiplier d . (b) The corresponding standard deviation of pixel values for the satin sample: sharper lobes provide shinier appearance and result in greater standard deviation. The role of d is more complicated.

Appearance Matching Processing the CT data yields the spatially varying density and orientation for the volume. But the optical appearance parameters of the model remain to be determined. Since the CT scan does not give the material's optical properties, Zhao et al. [2011] made use of a photograph of the material to compute the appearance parameters.

They assumed that the volume contains the same material, with differences only in density and orientation. Thus, the appearance parameters that must be determined are the same across the whole volume. They are: the standard deviation of the flake distribution γ (corresponding to fiber roughness), the scattering albedo α (corresponding to material color), and the density scale d (corresponding to opacity). Figure 25-(a) illustrates the effects of these parameters.

Since the density multiplier plays a fairly complicated role with respect to both measures, Zhao et al. [2011] chose to pre-determine the density multiplier in their implementation, which greatly simplifies the inverse problem and leads to a practical solution. With a fixed density multiplier, we can solve for the values of albedo (α , estimated separately in red, green, and blue) and lobe width (γ , a single scalar value) using an iterative algorithm. Note that the mean and standard deviation of pixel values change monotonically with changes in α and γ respectively. Thus, a binary search can be used to significantly improve performance as follows: first, an initial guess of γ is assumed, and search for the α to match the mean pixel value. Then, fixing α , perform a search for the γ to match the standard

deviation. These iterations are repeated until a match is found.

Finally, physically-based rendering algorithms, such as Monte Carlo path tracing, can be performed to create high-quality images of the modeled cloth. Figure 26 shows a rendered silk satin in a draped configuration. The orientation information in the volume automatically causes the characteristic appearance of this fabric to emerge.




Figure 26: *Silk satin rendered with volumetric model in [Zhao et al. 2011].*

4.3.2 Synthesizing Models with Complex Designs

Recall that the technique introduced in Section 4.3.1 has a major limitation: it can only model the exact scanned material. In addition, recent micro CT scanners can only take a small sample (about $0.5 \times 0.5 \text{ cm}^2$). So the measured information cannot capture yarn-level textures. Zhao et al. [2012] extended their previous work and introduced a structure-aware synthesis algorithm to overcome this limitation. Next, we briefly introduce this approach.

The goal is to produce volume models of woven materials, suitable for realistic close-up renderings, from two inputs: a description of the material to be simulated, and a few examples of similar but simpler fabrics. This can be accomplished in two phases: in the first phase, which only needs to be done once for a whole class of materials, CT scans of the example fabrics are used to build exemplars that contain all the information needed to synthesize large areas of complex fabrics; in the second phase, which is done once per material to be simulated, the exemplars are used in a new structure-aware volumetric texture synthesis method to synthesize a volume model according to the colors and weave pattern of the target material.

Exemplar creation phase. The purpose of the exemplar creation phase is to turn raw volume data into a database that can be used to synthesize volume models of a range of fabrics that are made from materials similar to the example materials, but have different




Figure 27: Pillows rendered using volumetric models created by [Zhao et al. 2012]. Each pillow contains 3×10^{12} effective voxels.

structure. Normally the example materials are a set of simple weaves made using particular types of yarns.

The input to this phase is the raw volume containing density information on a fine voxel grid covering a small patch of a fabric. High resolution scans are required to resolve fiber orientation and flyaway fibers, so each scan observes an area on the order of 5mm across, which, after cropping, typically produces exemplars with about 6×9 yarn crossings. A processing pipeline takes this data and produces output by denoising the input density data, automatically tracking yarns in the data to detect the yarn trajectories, segmenting the voxels to match them to the appropriate yarns, and then automatically detecting the pattern of yarn crossings. Each exemplar in the resulting database includes a voxel grid (containing density, fiber orientation, and yarn ID) and a small binary image representing the weave pattern.

Synthesis phase. The input describing a new fabric to be simulated includes a 2D binary array giving the weave pattern for the whole cloth, and 2D arrays specifying the type of warp and weft yarn present at each yarn crossing. The synthesis phase creates an output volume that respects the input specification while displaying local structure and details that match the exemplars. More details are available in their paper.

Figure 27 shows a rendered image using this approach.

5 Practical Physically-Based Rendering of Cloth

Once we have a volumetric representation of cloth, we want to render it efficiently. For all recent approaches dealing with volumetric cloth rendering, a voxelized model based on octrees or grids is used. Essentially, this gives a thin layer of voxels representing the cloth. In contrast to classical media such as smoke, cloth is (optically) very thick meaning that there is a very high probability for a ray of being scattered or absorbed inside the cloth model. In addition, multiple scattering tends to be an important factor affecting appearance. For that reason conventional ray marching is not suitable for rendering. On the other hand random walk sampling based on ray tracing is quite effective in this case, either based on Delta-tracking [Woodcock et al. 1965] or even simpler based on analytical sampling, assuming constant density inside a voxel. This approach is surprisingly fast for moderately fine voxelizations. It comprises two main steps: finding a new location where a scattering function is evaluated and computing a new direction for continuing the random walk. Here, the scattering function may be either a dedicated volumetric phase function (Section 4.2.1) or a fiber scattering function (Section 4.2.2); both found by sampling a yarn's statistics.

Volumetric rendering methods take its advantage mainly from avoiding an explicit representation of individual yarn fibers using less extensive statistical models. This way they can handle large pieces of cloth with orders of magnitude more fibers than methods representing fibers geometrically without any restrictions with respect to self similarities (as in repeating textures). Besides drastically reducing the memory requirements compared to an explicit representation, the rendering times also decrease by a factor. Using this approach, we are able to render pieces of cloth with a visual quality comparable to fiber-based reference solutions while employing a completely generic method for modeling fibers without limitations.

We take a Monte Carlo path tracing approach to render the voxelized cloth. In contrast to participating media or highly scattering materials (such as skin), path tracing is effective in our case as, due to absorption inside the fibers, energy quickly decays to zero after a few scattering events.

5.1 Monte Carlo Path Tracing

The basic idea of unidirectional backward Monte Carlo path tracing is computing pixel intensities by performing stochastic sampling in path space starting from the eye. Essentially, this means, tracing rays through the virtual scene and finding contributions due to surface and volumetric scattering.

In the following we present pseudocode ⁵ for rendering volumetric descriptions of cloth. Starting with a standard backward path tracer for surfaces as reference (Algorithms 1, 2), we discuss volumetric random walk sampling (3), which is finally extended to support both the micro-flake model as well as virtual scattering events.

The path tracing algorithm is comprising two main steps:

⁵The pseudo code for algorithms 1, 2, 3 and 4 is loosely based on lecture notes from Steve Marschner

- Generating primary camera rays by sampling the image plane and iterating over individual pixels (Algorithm 1).
- Estimating radiance along light paths (Algorithms 2, 3 and 6).

For surfaces, a rays throughput is computed by considering emmitance and BRDF at scattering events according to Algorithm 2. *Russian roulette* is used to terminate paths without introducing bias.

Algorithm 1 Iterating over all pixels

```

1: function RENDERIMAGE(camera)
2:   for all pixel indices i,j do
3:     result  $\leftarrow$  0
4:     for k=1 to #spp do ▷ #spp = samples per pixel
5:        $x, \omega \leftarrow$  camera.chooseRay(i,j) ▷ proportional to pixel filter
6:        $y \leftarrow$  traceRay( $x, \omega$ )
7:       result += radianceEstimator( $y, -\omega$ ) / #spp
8:     image(i,j)  $\leftarrow$  result

```

Algorithm 2 Pseudo code for surface- / BRDF-based radiance estimation

```

1: function RADIANCEESTIMATOR( $x, \omega$ )
2:   if random() < survivalProbability then ▷ russian roulette
3:      $\omega', \text{pdf} \leftarrow$  chooseDirection( normal( $x$ ),  $\omega$  ) ▷ compute new direction and
     corresponding probability density
4:      $y \leftarrow$  traceRay(  $x, \omega'$  )
5:     return (emittedRadiance( $x, \omega$ ) +  $\frac{\rho(x, \omega, \omega') * \text{radianceEstimator}(y, -\omega')}{\text{pdf}}$ ) / survivalProbability ▷ with BRDF
     bility
      $\rho$ 
6:   else
7:     return 0

```

For volumes the general structure is very similar (Algorithm 3). However, inside the medium, scattering needs to be computed from the statistical information stored in the voxels. The key function is given by Algorithm 4. Here, points of interaction with the medium are computed based on extinction cross section σ_t . Let δ be the distance to the next light source. Then the probability p for being scattered before approaching the respective point on the light source is given by:

$$p(\delta) = 1 - T(\delta) = 1 - e^{-\delta \sigma_t}, \quad (21)$$

where T is the transmittance.

First, the distance s to the next (hypothetical) point of interaction along the ray is found by inverting transmittance T :

$$s = -\frac{\log(\xi_1)}{\sigma_t}. \quad (22)$$

If $s > \delta$ no scattering occurs.

Algorithm 3 Pseudo code for random walk sampling in a volume

```
1: function RADIANCEESTIMATOR( $x, \omega$ )
2:   if random() < survivalProbability then                                ▷ russian roulette
3:      $s \leftarrow \text{chooseDistance}(\sigma_t)$                                 ▷ in  $[0, \text{inf}]$  prop. to  $\exp(-\sigma_t s)$ 
4:      $y \leftarrow \text{traceRay}(x, -\omega)$ 
5:      $\delta = |x - y|$                                                         ▷ distance to next light source
6:     if  $s > \delta$  then
7:       return emittedRadiance( $y, \omega$ ) / survivalProbability
8:     else if random() < albedo then                                       ▷ albedo =  $\frac{\sigma_s}{\sigma_t}$ 
9:        $x' \leftarrow x - s * \omega$ 
10:       $\omega', \text{pdf} \leftarrow \text{chooseDirection}(x, \omega)$                     ▷ compute new direction and
corresponding probability density
11:      return  $\frac{\text{radianceEstimator}(x', \omega')}{\text{pdf}}$  / survivalProbability
12:    else
13:      return 0
14:  else
15:    return 0
```

Algorithm 4 Choose distance to next point of interaction in case of a constant medium

```
1: function CHOOSDISTANCE( $\sigma_t$ )
2:    $\xi = \text{random}()$                                                         ▷ uniform in  $[0, 1]$ 
3:   return  $-\log(\xi) / \sigma_t$                                            ▷ natural log
```

Monte Carlo Path Tracing with Anisotropic Media Algorithm 3 also works for anisotropic media. However, Algorithm 4 needs to be extended since σ_t is no longer constant and depends on both position x and direction ω . Assume

$$\sigma_t^M = \max_{x, \omega} \sigma_t(x, \omega),$$

the following algorithm implements Woodcock tracking [1965] and samples a distance d for ray (x, ω) in an unbiased manner.

Algorithm 5 Choose distance to next point of interaction in case of an anisotropic medium

```
1: function CHOOSDISTANCE( $x, \omega, \sigma_t^M$ )
2:    $d = -\log(\text{random}()) / \sigma_t^M$ 
3:   while  $\sigma_t(x - d\omega, \omega) / \sigma_t^M < \text{random}()$  do
4:      $d = d - \log(\text{random}()) / \sigma_t^M$ 
5:   return  $d$ 
```

Monte Carlo Path Tracing with Virtual Scattering Events using an Octree Now we extend the random walk path tracer to support virtual scattering events (Algorithm 6). The major difference compared to Algorithm 3 is that we step through the voxels of the octree during ray traversal, assuming constant optical properties inside a voxel. Moreover,




Figure 28: Path tracing with virtual scattering events. The last virtual scattering event was created in voxel V at position $\vec{x}$. The scattered ray is leaving $\vec{x}$ to direction $\vec{d}$. The probability of being scattered in V along the ray is computed based on δ_V and total scattering coefficient σ_s .

in contrast to Algorithm 3, virtual scattering events are computed according to the statistics about fiber orientation and fiber scattering distribution functions (BCSDF) stored in a voxel. Please note that all attenuation is completely handled by the BCSDF and therefore, in this case, only σ_s needs to be considered for choosing the next virtual scattering event.

Recalling Section 4.2.2 each yarn in a voxel V is represented by one of $\#GMM(V)$ components of a Gaussian mixture model. Therefore, to obtain the total scattering coefficient σ_s^j in voxel V for all yarns, we have to sum over all individual scattering coefficients of the mixture components (Algorithm 6, Line 10) computed from ρ^j . If we scatter inside a voxel (Algorithm 6, Line 18), we choose a mixture component and select a fiber direction accordingly. Figure 28 illustrates the process.

As for conventional path tracing rendering efficiency can be improved by employing (multiple) importance sampling based on BCSDF as well as other optimizations such as explicit computation of direct lighting by using shadow rays (next event estimation).

5.2 Other Approaches

A variety of different methods have been proposed that try to avoid some of the problems of pure Monte Carlo path tracing. Although being general rendering approaches, many of them have been successfully used for the special case of cloth rendering. Instead of using a unidirectional path tracer, bidirectional path tracing [Veach and Guibas 1994], Metropolis light transport [Veach and Guibas 1997; Pauly et al. 2000] or energy redistribution path tracing [Cline et al. 2005] could directly be applied as well.

Instant Radiosity [Keller 1997] attempts to lower the computational overhead of global illumination rendering by reducing the problem to direct illumination computations for a large number of virtual point lights (VPLs). These VPLs are placed in a particle tracing step

Algorithm 6 Pseudo code for virtual scattering

```
1: function RADIANCEESTIMATOR( $x, \omega$ )
2:   if random() < survivalProbability then ▷ russian roulette
3:      $V \leftarrow \text{voxelThatIncludes}(x)$ 
4:      $y \leftarrow \text{traceRay}(x, -\omega)$  ▷ next hit point with light source along ray
5:      $L_e \leftarrow \text{emittedRadiance}(y, \omega) / \text{survivalProbability}$  ▷ emitted radiance, return
   zero if no light gets hit
6:   while insideVolume( $x$ ) do
7:      $\delta \leftarrow |x - y|$ 
8:      $b \leftarrow \text{nextIntersectionWithVoxelBoundary}(x, \omega)$ 
9:     ▷ Recall mixture component:  $g_j = \{\mathcal{N}_j(\vec{m}_j, s_j^2), \rho^j, \text{BCSDF}_j\}$ 
10:     $\sigma_s \leftarrow \sum_j^{\text{\#GMM}(V)} \text{getEffectiveCrossSection}(\rho^j)$  ▷ Equation 21
11:     $s \leftarrow \text{chooseDistance}(\sigma_s)$  ▷ in  $[0, \text{inf}]$  prop. to  $\exp(-\sigma_s s)$ 
12:     $\delta_V \leftarrow |x - b|$ 
13:    if  $s > \delta_V$  then
14:       $V \leftarrow \text{advanceToNextIntersectingVoxel}(x, \omega)$ 
15:       $x \leftarrow b$ 
16:    else if  $s > \delta$  then
17:      return  $L_e$ 
18:    else
19:       $x' \leftarrow x - s * \omega$ 
20:       $j \leftarrow \text{chooseMixtureComponent}(V)$ 
21:       $\psi \leftarrow \text{chooseFiberDirection}(s_j^2, \vec{m}_j)$ 
22:       $\omega', \text{pdf} \leftarrow \text{chooseDirection}(x, \omega, \psi, \text{BCSDF}_j)$  ▷ compute new
   direction and corresponding probability density
23:      return  $\frac{\text{radianceEstimator}(x', \omega')}{\text{pdf}} * \text{BCSDF}_j(\omega, \psi, \omega') / \text{survivalProbability}$ 
24:    return  $L_e$ 
25:  else
26:    return 0

1: function CHOOSEMIXTURECOMPONENT( $V$ )
2:    $\xi \leftarrow \text{random}()$ 
3:   ▷  $\rho_i$  is the density of the  $i$ -th mixture component of voxel  $V$ 
4:   return  $\underset{k \in \{1, \dots, \text{\#GMM}(V)\}}{\text{argmin}} \left( \xi \leq \sum_{j=1}^k \frac{\rho^j}{\sum_i \rho^i} \right)$ 

1: function CHOOSEFIBERDIRECTION( $\text{variance}, d$ ) ▷ formally accurate only if
    $\text{variance}$  is small
2:    $\theta \leftarrow \text{gaussRandom}(\text{variance})$ 
3:    $\phi \leftarrow \text{random}() * 2\pi$ 
4:   return sphericalToVector( $d, \theta, \phi$ )
```



Figure 29: Several large pieces of cloth rendered with Monte Carlo path-tracing and virtual scattering in less than 50 minutes (for a resolution of 1000×500 pixels with 2048 samples per pixel) using a physically-based volumetric approach. Material properties are linked directly to optical properties of fibers, yarns and the weave pattern. Three different materials are shown: A hard looking carpet, a soft looking blanket and a translucent curtain. Figure taken from [Schröder et al. 2011b].

similarly to photon mapping. Combined with Lightcuts [Walter et al. 2005], a technique to render scenes with large amounts of point lights efficiently, this leads to a powerful global illumination approach. Lightcuts work by building a hierarchy of lights, where for each level, a representative light is selected from a cluster of individual lights; an appropriate level is selected during rendering per pixel based on perception guided error bounds. The error that is introduced by this approach has a different visual quality than that of a typical Monte Carlo rendering. High-frequency noise is replaced with lower frequency errors and banding artifacts. For material perception, these types of errors may actually be more visually pleasant than noise: The type of high frequency noise is responsible for cloth being perceived as soft or hard. If the noise at fine scales is dominated by Monte Carlo noise, the look and feel of a material may change drastically. Especially in the context of computer aided design, one may rather prefer small errors in global illumination if high frequency detail is reproduced correctly.

In more recent work (Bidirectional Lightcuts [Walter et al. 2012]), not only virtual point lights are placed but also virtual sensors are distributed in the scene to obtain even more of the relevant light paths (in the spirit of bidirectional path tracing). The paper successfully demonstrates rendering of a volumetric cloth dataset. While the approach can be very effective in many cases, problems may arise when highlights for very specular fibers should be reproduced. Clustering of lights works well for diffuse materials, however, for mirror-like surfaces many virtual point lights are required to faithfully contribute to narrow reflection lobes. Especially during animation rendering, this can lead to serious artifacts as temporal coherence of sampling patterns cannot be guaranteed. Also, one has to develop an equation for every new type of material that conservatively bounds the error introduced

by light clustering.

Another common approach for global illumination computations is the photon mapping algorithm originally proposed by Jensen [1996]. Here, particles carrying flux (“photons”) are traced from the light sources and at each interaction with a surface, a photon is stored in a spatial data structure, the so called photon map.

After the particle tracing process, rays are shot from the camera and for each intersection point a set of k nearest photons is queried from the photon map. Based on these photons incident irradiance is approximated using density estimation. This idea was adopted to the realm of volume rendering as well.

Recently, to avoid practical problems of photon mapping, such as memory requirements for high quality results, progressive photon mapping has been proposed by Hachisuka et al. [2008]. Here, the process is reversed: Instead of storing photons, the locations of the first hit with a diffuse surface are stored in a spatial search data structure for eye paths. In a second pass, photons are shot and their contributions to stored locations, reachable within some search radius, are computed. The search radius is reduced over time and this allows us to obtain a consistent estimator.

In our context, however, radii larger than the thickness of cloth may easily result in mixing photons from the front and back sides. Similarly, for an exact reproduction of the material appearance, search radii larger than the cross section of yarns may be problematic when differently colored yarns are next to each other.

5.3 Bidirectional Texture Function Synthesis

For both explicit representations and volumetric models, all light transport effects including expensive computations of multiple-scattering have to be computed online. An alternative is to precompute radiance transfer for cloth in an offline process; this information can then be used for efficient rendering of the material afterwards. A natural representation for storing light transport of a material is the BTF (Section 2.2.2). This approach is generally interesting for predictive rendering of complex materials that cannot be efficiently represented by simple BRDF or texture models.

To compute a BTF, a virtual gonio-reflectometer setup can be constructed, which computes the angular reflectance data for a virtual material sample. Essentially, light transport has to be simulated for thousands of viewing and lighting directions. A straight forward solution is to build a virtual half-dome multi-view multi-light measurement device and render e.g. 70×70 images (“BTF slices”) of a flat material sample for 70 viewing and 70 lighting directions. As this setup shares similarities with existing measurement devices, tool support is available. While cloth tends to look relatively diffuse when light is scattered many times between individual fibers of yarns, the individual fibers themselves are highly specular. The angular resolution has to be carefully adapted to capture all effects at the scale one is interested in. As materials often consist of repeating structures, recent work [Schröder et al. 2011a] suggests to use redundancy of information in the spatial domain to synthesize BTFs faster.

For material design, one can imagine a system where an explicit representation of micro-geometry is used in a close-up view during a design phase. However, as explicit represen-

tations do not scale and may easily become too costly in terms of both memory size and rendering times, when applied to whole objects and large scenes, one can then synthesize BTFs based on the created micro-structure for interactive inspection.

6 Acknowledgements

We would like to thank Wenzel Jakob for proof reading parts of the course notes and writing an initial version of Section 4.1.1. This work is supported in part by NRW State within the B-IT Research School.

References

- ADABALA, N., MAGNENAT-THALMANN, N., AND FEI, G. 2003. Real-time visualization of woven textiles. In *Publication of EUROSIS*, 502–508.
- ADLER, C. L., LOCK, J. A., STONE, B. R., AND GARCIA, C. J. 1996. High-order interior caustics produced in scattering of a diagonally incident plane wave by a circular cylinder. *Optical Society of America* 14.
- ADLER, C. L., LOCK, J. A., AND STONE, B. R. 1998. Rainbow scattering by a cylinder with a nearly elliptical cross section. *Optical Society of America*.
- ADLER, C. L., PHIPPS, D., SAUNDERS, K. W., NASH, J. K., AND LOCK, J. A. 2001. Supernumerary spacing of rainbows produced by an elliptical-cross-section cylinder. *Optical Society of America*.
- ARBREE, A., WALTER, B., AND BALA, K. 2009. Heterogeneous subsurface scattering using the finite element method. To appear in *IEEE Transactions on Visualization and Computer Graphics*.
- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000*, 65–74.
- CASE, K., AND ZWEIFEL, P. 1967. *Linear transport theory*, vol. 29. Addison-Wesley Reading, Mass.
- CEREZO, E., PÉREZ, F., PUEYO, X., SERON, F. J., AND SILLION, F. X. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21, 5, 303–328.
- CLINE, D., TALBOT, J., AND EGBERT, P. 2005. Energy redistribution path tracing. In *ACM Transactions on Graphics (TOG)*, vol. 24, ACM, 1186–1195.
- CRASSIN, C., NEYRET, F., LEFEBVRE, S., AND EISEMANN, E. 2009. Gigavoxels : Ray-guided streaming for efficient and detailed voxel rendering. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D)*.
- DANA, K., VAN GINNEKEN, B., NAYAR, S., AND KOENDERINK, J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics (TOG)* 18, 1, 1–34.
- DAUBERT, K., LENSCH, H., AND HEIDRICH, W. 2001. Efficient cloth modeling and rendering. In *Rendering techniques 2001, United Kingdom, June 25-27, 2001*, 63.
- DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, ACM, New York, NY, USA, 1032–1039.

- DREBIN, R. A., CARPENTER, L., AND HANRAHAN, P. 1988. Volume rendering. In *Computer Graphics (Proceedings of SIGGRAPH 88)*, 65–74.
- GROLLER, E., RAU, R. T., AND STRASSER, W. 1995. Modeling and visualization of knitwear. *IEEE Transaction on Visualization and Computer Graphics* 1, 4, 302–310.
- HACHISUKA, T., OGAKI, S., AND JENSEN, H. 2008. Progressive photon mapping. In *ACM Transactions on Graphics (TOG)*, vol. 27, ACM, 130.
- HAINDL, M., AND FILIP, J. 2011. Advanced textural representation of materials appearance. In *SIGGRAPH Asia 2011 Courses*, ACM, 1.
- HEINO, J., ARRIDGE, S., SIKORA, J., AND SOMERSALO, E. 2003. Anisotropic effects in highly scattering media. *Phys. Rev. E* 68, 3.
- IRAWAN, P., AND MARSCHNER, S. 2006. A simple, accurate texture model for woven cotton cloth. Tech. rep., Tech. Rep. PCG-06-01, Cornell University, Department of Computer Science.
- IRAWAN, P. 2007. *Appearance of Woven Cloth*. PhD thesis, Cornell University.
- ISHIMARU, A. 1978. *Wave Propagation and Scattering in Random Media*. Academic Press, New York, USA.
- JAKOB, W., ARBREE, A., MOON, J. T., BALA, K., AND MARSCHNER, S. 2010. A radiative transfer framework for rendering materials with anisotropic structure. *ACM TOG* 29 (July), 53:1–53:13.
- JAKOB, W., ARBREE, A., MOON, J. T., BALA, K., AND MARSCHNER, S. 2010. A radiative transfer framework for rendering materials with anisotropic structure. Tech. rep., Cornell University. (Expanded version), <http://hdl.handle.net/1813/14982>.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001*, 511–518.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 01*, 511–518.
- JENSEN, H. 1996. Global illumination using photon maps. *Rendering Techniques* 96, 21–30.
- JOHNSON, P. M., AND LAGENDIJK, A. 2009. Optical anisotropic diffusion: new model systems and theoretical modeling. *Journal of Biomedical Optics* 14, 5.
- KAJIYA, J. T., AND HERZEN, B. P. V. 1984. Ray tracing volume densities. In *Computer Graphics (Proceedings of SIGGRAPH 84)*, 165–174.
- KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with three dimensional textures. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, 271–280.
- KALDOR, J., JAMES, D., AND MARSCHNER, S. 2008. Simulating knitted cloth at the yarn level. In *ACM Transactions on Graphics (TOG)*, vol. 27, ACM, 65.

- KELLER, A. 1997. Instant radiosity. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 49–56.
- KISTLER, R. 2011 (accessed February 2, 2012). *Kleines Zwirn Lexikon*.
- LEVOY, M. 1988. Display of surfaces from volume data. *IEEE Computer Graphics & Applications* 8, 3 (May), 29–37.
- MAGNENAT-THALMANN, N., CORDIER, F., KECKEISEN, M., KIMMERLE, S., KLEIN, R., AND MESETH, J. 2004. Simulation of clothes for real-time applications. In *Eurographics 2004, Tutorials*, INRIA and the Eurographics Association.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, 519–528.
- MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. *ACM TOG* 22, 3, 780–791. SIGGRAPH 2003.
- MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Trans. Graph.* 24 (July), 727–734.
- MAX, N. L. 1994. Efficient light propagation for multiple anisotropic volume scattering. In *Fifth Eurographics Workshop on Rendering*, 87–104.
- MEES, L., REN, K. F., GREHAN, G., AND GOUESBET, G. 1998. Scattering of a Gaussian beam by an infinite cylinder with arbitrary location and arbitrary orientation: numerical results. *Optical Society of America*.
- MESETH, J., MÜLLER, G., SATTTLER, M., AND KLEIN, R. 2003. Btf rendering for virtual environments. In *Virtual Concepts 2003*, 356–363.
- MOON, J. T., AND MARSCHNER, S. R. 2006. Simulating multiple scattering in hair using a photon mapping approach. *ACM TOG* 25, 3, 1067–1074. SIGGRAPH 2006.
- MOON, J. T., WALTER, B., AND MARSCHNER, S. 2008. Efficient multiple scattering in hair using spherical harmonics. *ACM TOG* 27, 3. SIGGRAPH 2008.
- MORTON, W. E., AND HEARLE, J. W. S. 1962. *Physical properties of textile fibres / W.E. Morton and J.W.S. Hearle*. Textile Institute, Manchester, England :.
- MOUNT, C. M., THIESSEN, D. B., AND MARSTON, P. L. 1998. Scattering observations for tilted transparent fibers: evolution of airy caustics with cylinder tilt and the caustic merging transition. *Applied Optics* 37, 9, 243–249.
- MÜLLER, G., MESETH, J., SATTTLER, M., SARLETTE, R., AND KLEIN, R. 2005. Acquisition, synthesis, and rendering of bidirectional texture functions. In *Computer Graphics Forum*, vol. 24, Wiley Online Library, 83–109.
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. *Metropolis light transport for participating media*. Fachbereich Informatik, Univ.

- PERLIN, K., AND HOFFERT, E. M. 1989. Hypertexture. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, 253–262.
- PREMOZE, S., ASHIKHMIN, M., TESSENDORF, J., RAMAMOORTHY, R., AND NAYAR, S. 2004. Practical rendering of multiple scattering effects in participating media. In *Rendering Techniques 2004: 15th Eurographics Workshop on Rendering*, 363–374.
- ROBBINS, C. 1994. Chemical and physical behavior of human hair. *third ed.* Springer-Verlag, New York.
- SADEGHI, I. 2011. *Controlling the Appearance of Specular Microstructures*. Jacobs School of Engineering, UC San Diego.
- SATTLER, M., SARLETTE, R., AND KLEIN, R. 2003. Efficient and realistic visualization of cloth. In *Proceedings of the 14th Eurographics workshop on Rendering*, Eurographics Association Aire-la-Ville, Switzerland, Switzerland, 167–177.
- SCHRÖDER, K., MÖLLER, D., KLEIN, R., AND ZINKE, A. 2011. Non-local image reconstruction for efficient btf synthesis. In *SIGGRAPH Asia 2011 Sketches*, ACM, 30.
- SCHRÖDER, K., KLEIN, R., AND ZINKE, A. 2011. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum (Proceedings of EGSR 2011)* 30, 4 (July), 1277–1286.
- SCHUH, R., AND WRIEDT, T. 2003. Light scattering from bent cylindrical fibers for fiber length and diameter characterization. *Particle & Particle Systems Characterization* 20, 243–249.
- SCHWARTZ, C., RUITERS, R., WEINMANN, M., AND KLEIN, R. 2011. WebGL-based streaming and presentation framework for bidirectional texture functions. In *The 12th International Symposium on Virtual Reality, Archeology and Cultural Heritage VAST 2011*, Eurographics Association, Eurographics Association, 113–120. Best Paper Award.
- SHINOHARA, T., TAKAYAMA, J., OHYAMA, S., AND KOBAYASHI, A. 2010. Extraction of yarn positional information from a three-dimensional CT image of textile fabric using yarn tracing with a filament model for structure analysis. *Textile Research Journal* 80, 7, 623–630.
- SREPRATEEP, K., AND BOHEZ, E. 2006. Computer Aided Modelling of Fiber Assemblies. *Comput. Aided Des. Appl* 3, 1-4, 367–376.
- USDA E.R.S., 2004. <http://www.ers.usda.gov/briefing/cotton/>.
- VEACH, E., AND GUIBAS, L. 1994. Bidirectional estimators for light transport. *surfaces* 18, 19, 20.
- VEACH, E., AND GUIBAS, L. 1997. Metropolis light transport. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 65–76.
- VOLEVICH, V. L., KOPYLOV, E. A., KHODULEV, A. B., AND KARPENKO, O. A. 1997. An approach to cloth synthesis and visualization. In *The 7-th International Conference on Computer Graphics and Visualization*.

- WALTER, B., FERNANDEZ, S., ARBREE, A., BALA, K., DONIKIAN, M., AND GREENBERG, D. 2005. Lightcuts: a scalable approach to illumination. In *ACM Transactions on Graphics (TOG)*, vol. 24, ACM, 1098–1107.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Eurographics Workshop on Rendering 2007*, 195–206.
- WALTER, B., KHUNGURN, P., AND BALA, K. 2012. Bidirectional lightcuts. *ACM Transactions on Graphics (TOG)* 31, 4, 59.
- WALTER, B. 2005. Notes on the Ward BRDF. Tech. Rep. PCG-05-06, Program of Computer Graphics, Cornell University.
- WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph. (to appear in SIGGRAPH 2008)* 27, 3, 41:1–41:9.
- WESTERMANN, R., AND ERTL, T. 1998. Efficiently using graphics hardware in volume rendering applications. In *Proceedings of SIGGRAPH 98*, 169–178.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. *SIGGRAPH Comput. Graph.* 26, 2, 255–264.
- WOODCOCK, E., MURPHY, T., HEMMINGS, P., AND LONGWORTH, S. 1965. Techniques used in the gem code for monte carlo neutronics calculations in reactors and other systems of complex geometry. In *Proc. Conf. Applications of Computing Methods to Reactor Problems*, 557.
- XU, Y.-Q., CHEN, Y., LIN, S., ZHONG, H., WU, E., GUO, B., AND SHUM, H.-Y. 2001. Photorealistic rendering of knitwear using the lumislice. In *SIGGRAPH '01*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 391–398.
- YASUDA, T., YOKOI, S., TORIWAKI, J., AND INAGAKI, K. 1992. A shading model for cloth objects. *Computer Graphics and Applications, IEEE* 12, 6 (nov.), 15–24.
2011. An evaluation on woven cloth rendering techniques. 7–12.
- ZHAO, S., JAKOB, W., MARSCHNER, S., AND BALA, K. 2011. Building volumetric appearance models of fabric using micro CT imaging. *ACM Trans. Graph.* 30 (Aug.), 44:1–44:10.
- ZHAO, S., JAKOB, W., MARSCHNER, S., AND BALA, K. 2012. Structure-aware synthesis for predictive woven fabric appearance. *ACM Trans. Graph.* 31 (Aug.), XX:1–XX:10.
- ZINKE, A., AND WEBER, A. 2007. Light scattering from filaments. *IEEE TVCG* 13, 2, 342–356.
- ZINKE, A., SOBOTTKA, G., AND WEBER, A. 2004. Photo-realistic rendering of blond hair. In *VMV 2004*, 191–198.
- ZINKE, A., YUKSEL, C., WEBER, A., AND KEYSER, J. 2008. Dual scattering approximation for fast multiple scattering in hair. *ACM TOG* 27, 3. SIGGRAPH 2008.

A Useful Transformations for Fiber Scattering

We will now derive some useful transformations for variables of the two sets described in Section 3.3. The projection of the angle of incidence at the enclosing cylinder onto the normal plane γ' is given by the following two equations:

$$\cos \gamma' = \frac{\cos \beta}{\cos \theta} \quad (23)$$

$$\gamma' = |\arcsin h| \quad (24)$$

Thus, for the spherical angle β we have:

$$\beta = \arccos (\cos \theta \sqrt{1 - h^2}) . \quad (25)$$

The following relation holds between ξ , φ and h , cf. Fig. 15 Right:

$$\xi = \begin{cases} \varphi + \arcsin h & 0 \leq \varphi + \arcsin h < 2\pi \\ 2\pi + \varphi + \arcsin h & \varphi + \arcsin h < 0 \\ \varphi + \arcsin h - 2\pi & \varphi + \arcsin h \geq 2\pi \end{cases} \quad (26)$$

The spherical coordinate α equals the angle between the projection of $\vec{\omega}$ onto the local tangent plane and $\vec{v}$. Since the angle between $\vec{\omega}$ and $\vec{v}$ is given by $\pi/2 - \theta$ and the inclination of $\vec{D}$ with respect to the tangent plane is $\pi/2 - \beta$ one obtains:

$$\begin{aligned} \cos \alpha &= \cos (\pi/2 - \theta) / \cos (\pi/2 - \beta) \\ &= \sin \theta / \sin \beta \\ &= \frac{\sin \theta}{\sqrt{1 - (1 - h^2) \cos^2 \theta}} . \end{aligned} \quad (27)$$

with

$$\alpha = \begin{cases} \arccos \left(\frac{\sin \theta}{\sqrt{1 - (1 - h^2) \cos^2 \theta}} \right), & h \leq 0, \\ 2\pi - \arccos \left(\frac{\sin \theta}{\sqrt{1 - (1 - h^2) \cos^2 \theta}} \right), & h > 0. \end{cases} \quad (28)$$