DOI: 10.1111/j.1467-8659.2011.01987.x

Eurographics Symposium on Rendering 2011
Ravi Ramamoorthi and Erik Reinhard
(Guest Editors)

Volume 30 (2011), Number 4

A Volumetric Approach to Predictive Rendering of Fabrics

K. Schröder[1], R. Klein[1] and A. Zinke[2]

[1]University of Bonn, Germany
[2]GfaR mbH, Germany

Figure 1: *Several large pieces of cloth rendered with Monte Carlo path tracing in less than 50 minutes (for a resolution of* 1000×500 *pixels with 2048 samples per pixel) using the physically-based volumetric approach presented in this work. Three different materials are shown: A hard looking carpet, a soft looking blanket and a translucent curtain.*

Abstract

Efficient physically accurate modeling and rendering of woven cloth at a yarn level is an inherently complicated task due to the underlying geometrical and optical complexity. In this paper, a novel and general approach to physically accurate cloth rendering is presented. By using a statistical volumetric model approximating the distribution of yarn fibers, a prohibitively costly explicit geometrical representation is avoided. As a result, accurate rendering of even large pieces of fabrics containing orders of magnitudes more fibers becomes practical without sacrifying much generality compared to fiber-based techniques. By employing the concept of local visibility and introducing the effective fiber density, limitations of existing volumetric approaches regarding self-shadowing and fiber density estimation are greatly reduced.

1. Introduction

Predictive rendering of cloth is a challenging task. Both *geometrical complexity*, and *optical complexity* need to be considered: Highly anisotropic single and multiple scattering as well as self-shadowing effects dominate the appearance of cloths. As many types of fibers are highly translucent, multiple scattering significantly influences the observed color.

Two main approaches dominate the literature. Firstly, image-based techniques such as the *Bidirectional Texture Functions* (BTF) that allow one to render reasonably realistic visualizations of cloth. However, BTFs have several limitations by definition: 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. Moreover, applicability to cloth *design* is limited as editing with regard to cloth model parameters (for example the weave pattern or optical properties of fibers) is impossible unless extremely costly virtual gonioreflectometer methods are employed.

A second approach focuses on accurate models of micro-scale geometry described by individual fibers with associated material information e.g. in form of a *Bidirectional Curve Scattering Distribution Function* (BCSDF). Using path tracing methods and measured BCSDFs, this approach allows for highly accurate renderings and has already proven valuable in the context of predictive hair rendering [MJC*03, ZW07]. Due to its accuracy, it has been used to generate all reference renderings presented in this work.

Unfortunately, the method does not scale well – even small patches with 50 warp and 50 weft yarns easily take up to dozens of gigabytes of memory. Although one could imagine the use of geometry instancing as a work around, this idea does not work well in practice. Even if a repeating weave pattern is observed, which is often not to be the case, the geometry of spun yarns usually does not repeat with the same frequency. This is because 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, leading to the specific look and feel of the corresponding cloth.

To overcome the problems of modeling the underlying micro-scale geometry exactly, different solutions have been suggested – among them volumetric methods. While previous volumetric methods for rendering cloth already deliver a very good overall impression of cloth they are not accurate enough for predictive rendering since the anisotropic nature of the micro-geometry is not modeled properly. To overcome this problem, Jakob et al. [JAM*10] recently devised an elegant framework for volumetric modeling and rendering of materials with anisotropic micro-structure. Here, the local aggregate optical behavior of complex materials is modeled as a distribution of well-separated, non-spherical particles approximating the phase function of scattering events on a per-voxel level. The resulting volumetric representation is then rendered by employing a novel anisotropic diffusion method. The approach is very general and can describe many kinds of anisotropic structures well. It integrates perfectly into modern physics-based rendering systems. To be of practical use in the context of predictive cloth rendering, suitable particle models that efficiently pre-integrate to phase functions are required. Unfortunately, deriving such models directly from measured optical properties of the yarn fibers (i.e. the underlying fiber scattering functions) is very challenging – in particular as differently colored yarns and composite materials are common. Furthermore, while desirable from a theoretical standpoint, strictly volumetric approaches generally suffer from issues related to high-frequency volumetric detail. A sufficiently high volumetric resolution is required to describe small-scale structures in global illumination patterns and to avoid shadowing artifacts caused by discontinuities at interfaces between optically dense and sparse regions.

In the context of hair rendering, Moon et al. [MWM08] have developed a hybrid approach to speed up multiple scattering computations. Instead of modeling phase functions in a volume, they approximate geometry volumetrically using a statistical description and use the BCSDF of the individual fibers to model optical material properties. The hair geometry is first voxelized in a pre-processing step. Relevant fiber properties (the scattering coefficient, the average tangent of hair strands intersecting a voxel and the standard deviation with respect to this tangent) are stored in a uniform grid data structure. This information is used in a subsequent light tracing phase to approximate multiply scattered light by spherical harmonics. The statistical model is only used to accelerate the computation of multiple scattering while for single scattering the complete micro-geometry is still required. Therefore, the geometrical complexity remains high.

Since the micro-geometry of cloth is significantly more complex than the one of homogeneous assemblies of fibers, the simple statistical model introduced by Moon et al. is not sufficient for predictive cloth rendering. Yarns are created by twisting several hundreds of fibers into a cohesive thread; cloth is then formed by combining thousands of yarns. The tangent direction of spun fibers varies significantly over small distances and nearby yarns may have completely different material parameters. For the modeling of hairs it is reasonable to assume that the distribution of fiber directions is homogeneous inside a voxel. This has two main implications: Fiber directions can be locally described by a single Gaussian, which does not hold for cloth. Additionally, a rough approximation of the fiber density is sufficient to obtain satisfying results, while for cloth, this density has to be estimated much more accurately as fibers might be highly curved even inside a single voxel.

Nevertheless, the idea of statistically modeling the micro-geometry, while keeping the material properties of the individual fibers, is very appealing. It first offers the possibility to use optical material properties of individual fibers (which are either provided by manufacturers or can be measured efficiently [ZLHA*09]) and second delivers a possibility to deal with the inherent geometrical complexity. Therefore, we have incorporated this general idea into our novel approach and have developed a statistical model for the micro-geometry of cloth that is used for both single and multiple scattering.

Our main contributions are

- the simulation of single as well as multiple scattering using the statistical model
- the approximation of the fiber distribution as a Gaussian mixture model stored in an octree which uses one component for each yarn in a voxel
- the introduction of a novel approach to compute a so called *effective fiber density* allowing for highly accurate renderings
- a solution to a general shadowing problem inherent to any discrete volumetric representation rendered with path tracing by introducing the concept of *local visibility* for self-shadowing and modeling it as a new material parameter - the *Bidirectional Visibility Distribution Function* (BVDF).

In comparison with methods that represent fibers geometrically, our approach is able to handle larger pieces of cloth containing several orders of magnitude more fibers. Besides drastically reducing the memory requirements compared to an explicit representation, the rendering times decrease by a factor. Our results have a visual quality comparable to fiber-based reference solutions (as we show in Sec. 6) while employing a completely generic method for modeling fibers without limitations (including the use of arbitrary state-of-the-art fiber scattering models). We can handle high-frequency detail in both multiple scattering (because of the use of virtual scattering events) and shadowing (supported by the BVDF).

2. Related Work

Surface Reflectance Models. Daubert et al. [DLH01] model yarns using implicit surfaces and generate a BTF-like data representation using hardware rendering. Adabala et al. [AMTF03] attempt to use a simple BRDF model for efficient cloth rendering. For both works, the underlying BRDF is not expressive enough to model fabrics realistically.
A data driven microfacet-based BRDF approach was proposed by Wang et al. [WZT*08]. Here, for each measured surface point a normal density function model best fitting the observed data is used for rendering anisotropic spatially-varying materials, such as cloth. Other image-based methods use measured BTFs for realistic cloth rendering [MTCK*04]. A more sophisticated approach to modeling the appearance of woven cloth using BRDF and BTF at yarn level was presented by Irawan [Ira07]. 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 and cannot be directly inferred from optical fiber properties. However, BTFs have several problems as mentioned in Sec.1.

Explicit cloth modeling methods. Methods for modeling cloth at a fiber level, based on ray-tracing, have been presented in several works [WAT92, VKKK97]. Their main objective was to use micro-scale simulation to derive a BRDF/BSDF model for a small patch of cloth, whereas our objective is being able to render large pieces of cloth directly.

Volumetric Methods. A volumetric approach for modeling knitwear was proposed by Gröller et al. [GRS95]. By measuring the cross-sectional distribution of yarn fibers, a density field is created which is swept along a three-dimensional curve to form the entire yarn. The results look quite impressive, but the question how to deal with realistic fiber scattering models is not addressed. Similarly, [XCL*01] base all computations on a structure called lumislice: a light field of a yarn-cross-section. It remains unclear how a physically-based light field can be derived efficiently from optical properties of real yarn fibers. As already discussed in Sec. 1, [JAM*10] have presented a general framework for volumetric rendering of anisotropic structures.

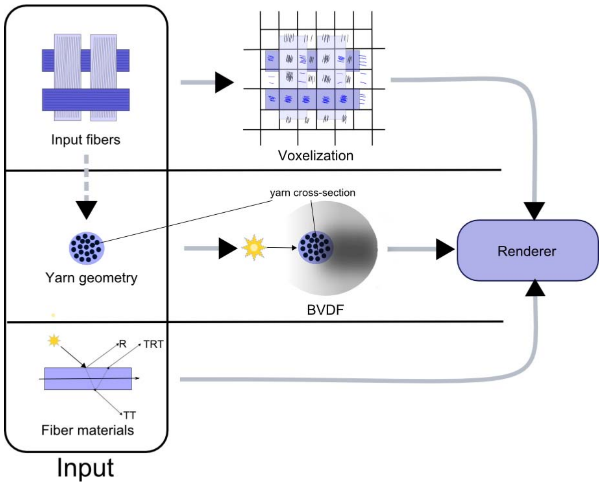

Figure 2: *Illustrating the pipeline, progressing from the input data on the left, over computation of voxelization and BVDF to the renderer on the right.*

Hair Rendering Methods. Besides explicit path tracing methods that require each hair strand to be modeled [ZSW04], very efficient approximations regarding multiple scattering computations have 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 [MM06, ZYWK08, MWM08].

3. Overview

The main advantage of the proposed method arises from avoiding an explicit representation of individual yarn fibers using less extensive volumetric statistical models. Fig. 2 illustrates the pipeline and the input to the Monte Carlo path tracing renderer. We first voxelize the model and store statistics modeling the distribution of yarn fibers in an octree. For each octree cell, we store the main direction and standard deviation of fibers of each single yarn inside a cell, attributed with information about local fiber density and a material index. Scattering at a fiber level is fully described by a BCSDF. Moreover, the BVDF, which stores information about local shadowing, is pre-computed from the yarn geometry. Ray/cloth intersections are computed as random variables via virtual scattering events. The whole process can be seen as sampling a voxel's phase function (which stays implicit) "on the fly".

4. Statistical Volumetric Modeling of Cloth

For comparison purposes, we first generate the cloth geometry explicitly in a representation our fiber-based path tracing reference implementation can render and then directly infer the volumetric representation from this geometry. For each voxel cell we describe the distribution of fibers using a statistical model.

4.1. Input data

In theory, our volumetric representation for a piece of cloth can be generated from several different kinds of input data ranging from ad-hoc line segment models synthesized onto a simulated base mesh, over full mechanical simulations at the yarn [KJM08, KJM10] or fiber level and up to measurements of real cloth based on Micro CT data. The only requirements are that the distribution of fibers can somehow be measured inside a voxel and that local visibility information can be estimated (c.f. Sec. 4.3). This paper focuses on rendering and optical simulation. Therefore, to allow for an easy comparison with state-of-the-art ray tracing of fibers modeled using cylinders, we have chosen an ad-hoc model: The basic path each yarn follows can be described by a spline which deforms with respect to the binary weave matrix [Gla02]. We generate geometry by spinning fibers around this base spline: We sample initial positions of fibers from a uniform distribution inside a circle around one end of the spline forming the cross section shape [SB06]. Then we connect successive positions as the circle is moved and optionally rotated along the spline, forming cylindric segments with radii corresponding to the radii of single fibers. Yarns generated in this way resemble clean filament yarns made up from industrial fibers. The composition of all fiber-cylinders of all yarns forms the cloth geometry which can be synthesized onto any warped grid mesh. For the examples shown in Fig. 6 the generation of geometry took only a second. For all results presented in this paper, we use a fully energy conserving variant of a state-of-the-art scattering model for dielectric fibers [ZW07] with varying, manually adjusted parameters.

4.2. Voxelization

The previously created cloth model is voxelized using an octree data structure that is also used for rendering. We generate the tree top-down by propagating cylinder segments intersecting a voxel to its children. Finally, for each leaf voxel cell containing cloth, we create a Gaussian mixture model representing local fiber distribution according to Sec. 4.3. In our current implementation, the maximum octree-level is set manually and all leaf voxels have that same level. For timings regarding the construction, we refer to Table 1. As can be seen, voxelization is extremely fast and does not constitute a bottleneck compared to the generation of acceleration data structures for explicit geometry. Construction time could be greatly reduced by directly generating statistics from the procedural cloth model instead of first constructing the geometry explicitly.

4.3. The statistical model

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

$$g_i = \{\mathcal{N}_i(m_i, s_i^2), \rho^i, \mathrm{BCSDF}_i\}. \quad (1)$$

Besides the average tangent direction m_i and its associated standard deviation s_i (forming a Gaussian $\mathcal{N}_i(m_i, s_i^2)$), a material index BCSDF_i is stored. In addition, the mean free path length inside a voxel, represented by the *effective fiber 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, in the following, we will refer to the whole tuple g_i as a Gaussian mixture component of a mixture model $G_V = \{g_1, ..., g_n\}$ where n equals the number of yarns inside voxel V. To minimize memory requirements, a compact representation was chosen to encode a Gaussian yarn model: Each mixture component is represented by as few as 13 bytes (16 bits for each of the 4 angles required for tangent direction and normal, 20 bits for the weight, 12 bits for sigma, 1 byte for a material index).

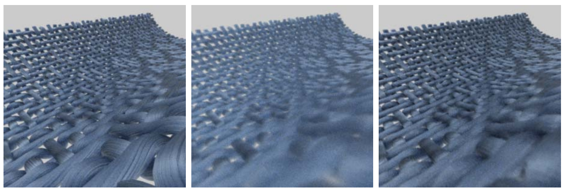

(a) Reference (b) Fiber counting (c) Our method

Figure 3: *Comparison of the effects of our effective fiber density to the fiber counting method presented in Moon et al. [MWM08] (apart from that, the same rendering technique has been used). Note that this is a very challenging setup showing a combination of directional and isotropic illumination with light falling through the cloth. In this setup, our method not only shows less blurring but can also reproduce the color resulting from multiple scattering of light much better. Note the light blue in image (b) compared to the dark blue in images (a) and (c).*

Fiber density. The mean free path length is calculated based on the perpendicular attenuation coefficient $\sigma^{\perp} = 2r_f\rho$ with fiber radius r_f. Existing approaches developed for hair rendering that attempt to estimate the density ρ by ad-hoc methods, being based on counting fibers in a sphere volume enclosing a voxel, are not suitable for cloth. As the length of lines covered inside a voxel is not properly taken into account, they fall short in case of highly curved fibers (such as spun yarn) and in case of fibers which barely intersect the volume of interest.

In the following, we will assume that a set of n line segments

$\{L_j, j = [1..n]\}$ is used to represent the cloth. To avoid the limitations of the fiber counting approach we introduce the concept of *effective fiber density* ρ.

Let's consider a spherical volume (similar to the search distance used by [MWM08]) for density estimation. Then the average length of a line intersecting a sphere with radius r equals $\bar{l}_R = \bar{l}_{sphere} = \frac{4 \cdot r}{3}$. The projected area of the sphere along the tangent direction equals $A_R = A_{sphere} = \pi r^2$. Based on these properties, the effective fiber density for a spherical region then is given as

$$\rho^{sphere} = \frac{l_{total}}{\bar{l}_R A_R} = N / A_R \tag{2}$$

with $l_{total} = \sum_{j=1..n} l_j$ denoting the total summed intersection length l_j of all line segments with the sphere.

Intuitively, this can be seen as computing an effective fiber count N, the number of infinite spatially uniformly distributed fibers intersecting R that would yield exactly the same density.

Of course the concept of effective fiber density may be generalized to arbitrarily shaped regions, i.e. to voxels. As the cloth model is discretized by voxels, undesirable spatial blurring occurs if spheres were used to estimate ρ. This blurring not only smooths out details but also affects the amount of multiply scattered light relevant for the appearance of cloth.

However, for voxels, the angular variations of $\bar{l}_R$ and A_R need to be considered. The effective fiber density can be computed by summing over all line segments intersecting a given voxel:

$$\rho^{voxel} = \sum_j \frac{l_j}{\bar{l}_j(t_j) A_{voxel}(t_j)}. \tag{3}$$

Here, $A_{voxel}(t_j)$ denotes the area of a voxel projected onto a plane perpendicular to the tangent direction t_j of a given line segment and $\bar{l}_j$ is the expectation of intersection length (the average) for t_j.

To speed up computations, tabulated values of $A_{voxel}(t_j)$ and $\bar{l}_j$ are used during voxelization. In Fig. 3, a comparison between our method and the method proposed by Moon et al. [MWM08] is given for a challenging close-up example.

The effective scattering coefficient As very briefly discussed by Moon et al. [MWM08], 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^j_{single} with respect to a given direction d computes as

$$\sigma^j_{single}(\alpha_j) = 2 r_f \rho^j_{single} \sin(\alpha_j) \tag{4}$$

where ρ^j_{single} denotes the fiber density induced by the (single) fiber of radius r_f and $\alpha_j = \angle(d, t_j)$.

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:

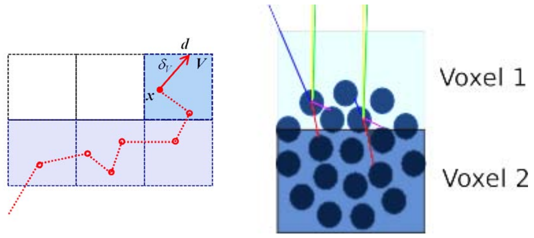

Figure 4: ***Left:*** *Path tracing with virtual scattering events. The last virtual scattering event was created in voxel V at position x. The scattered ray is leaving x in direction d. The probability of being scattered in V along the ray is computed based on* δ_V *and total scattering coefficient* σ_s. ***Right:*** *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. Note that the voxels shown have twice the size (per dimension) of the voxels used for the statistical model (c.f. Sec. 5.2).*

$$\sigma^i = \sum_{j=1..n_i} \sigma^j_{single}(\alpha_j). \tag{5}$$

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

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

where $\beta_i = \angle(d, m_i)$. To avoid computationally costly integration during rendering, we follow Moon et al. [MWM08] and pre-compute the integral. The time required for such pre-computation is negligible.

5. Monte Carlo Path Tracing with Virtual Scattering Events

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.

Because of the volumetric representation of geometry, discrete scattering events along ray paths need to be synthesized stochastically based on the statistical information stored in the octree.

Let x be the position of the last vertex of a ray path and d denote the direction of the associated outgoing ray R. Then the next vertex of the light path is computed based on virtual scattering events by the following steps (c.f. Fig.4 Left):

- **Step 1:** Compute the voxel V that includes x
- **Step 2:** Compute the position of a new potential virtual scattering event along d according to the total scattering coefficient (related to mean free path length) associated to V. This virtual scattering event is rejected if it lies outside the boundaries of V. In this case no scattering occurs and R advances to the next intersecting voxel.
- **Step 3:** Otherwise, a Gaussian mixture component of V is chosen and the ray gets scattered and attenuated according to the BCSDF.

5.1. Virtual Scattering Events

Once the current voxel V including the last scattering event has been computed, a new virtual scattering event needs to be created based on the Gaussian mixture components $G = \{g_1, .., g_n\}$ associated to V. Let δ_V be the distance at which the associated outgoing ray R is leaving (intersecting the boundary of) V. Then the probability p for being scattered inside V is given by Beer-Lambert's law as

$$p(\delta_V) = 1 - T(\delta_V) = 1 - e^{\delta_V \sigma_s}, \tag{6}$$

where T is the transmittance and σ_s is denoting the total scattering coefficient of V. To obtain the combined scattering probability for all mixture components, σ_s is computed by summing over all scattering coefficients $\sigma_s = \sum_i \sigma_s^i$. Based on the above, the virtual scattering event is stochastically computed using four uniformly distributed random numbers $\xi_{1..4} \in (0,1]$.

First, the distance δ of the virtual scattering event along the ray R is found by inverting transmittance T:

$$\delta = -\frac{\log(\xi_1)}{\sigma_s}. \tag{7}$$

If $\delta > \delta_V$, this scattering event is rejected, as it lies outside the voxel V.

Second, if scattering takes place, the new i^*-th Gaussian mixture component g_{i^*} is randomly selected with a probability proportional to its density

$$i^* = \underset{k \in \{1,..,n\}}{\operatorname{argmin}} \left(\xi_2 \leq \sum_{j=1}^{k} w_j\right), \tag{8}$$

with $w_j = \frac{\rho^j}{\sum_i \rho^i}$. Finally, based on ξ_3 and ξ_4, we choose a fiber direction determined by average tangent direction and standard deviation of g_{i^*} and eventually scatter according to the fiber scattering model BCSDF_{i^*}.

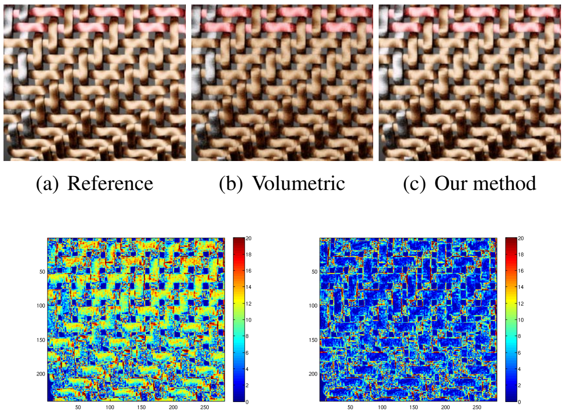

(a) Reference (b) Volumetric (c) Our method

(d) Difference: Volumetric (e) Difference: Our method

Figure 5: *Close-up comparison of the reference solution (a) to naive handling of shadowing by using the volumetric information only (b) and to our method using the BVDF for local visibility estimation (c). The scene is lit using a single point light – the fibers in this image exert a significant amount of absorption and specular highlights dominate the appearance. Incorrect handling of shadowing results in darkening and an incorrect output color for the purely volumetric technique. (d) and (e) show Delta-E color difference images for the volumetric technique and our method. Apart from general variance, our method shows most differences near yarn borders (caused by the discrete voxelization) whereas the purely volumetric method overestimates shadowing everywhere.*

5.2. Direct Lighting and Self-Shadowing

Taking a volumetric approach to model cloth statistically, self-shadowing details below the scale of a single voxel cannot be captured.

In particular, the directionality of local self-shadowing is lost unless a prohibitively fine voxelization is employed. However, especially for direct illumination, in order to obtain satisfactory results, this effect is critical. Consider the following example: A head light is located at the same position as the camera – all points which are visible from the camera should also be lit by the light. As there is a non-vanishing probability for virtual scattering events along the shadow-ray for any voxel containing fibers, shadowing is significantly overestimated. Moreover, as for any volumetric model, discontinuities at interfaces between optically dense and sparse regions are not well approximated (see Fig. 4 Right). This is particularly true for densely woven cloth (the common case).

To compensate for the above limitation of purely density-based representations, statistics for angularly dependent occlusion in case of direct illumination are used for approximating self-shadowing at a local level. Noting that these two kinds of bias affect shadowing within at least two adjacent voxels, statistics — which we call *Bidirectional Visi-*

bility Distribution Function — are computed at a resolution twice the desired size of a leaf voxel.

Bidirectional Visibility Distribution Function (BVDF). We propose 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 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 same order of magnitude as the size of a yarn's cross section, local visibility is modeled by the BVDF V_{bvdf}, a material property of yarns given by averaging V_l inside a voxel cell:

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

Note that V_l and V_g are binary, whereas $V_{bvdf} \in [0,1]$. In contrast to V_g, which 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 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 our case). Due to the small size of the sample (containing just a few yarns) this pre-computation is no bottleneck of our method and typically takes less than two minutes of computation for each yarn type.
To transfer the visibility information to a bent piece of cloth, inside each octree leaf cell we store a surface normal of the underlying locally planar base mesh used for cloth synthesis. w_o and w_i are expressed in a local coordinate frame spun by this normal and the tangent direction of the last virtual scattering event.

The general structure of BVDFs always looks similar: When light and camera position are adjacent (green rays in Fig. 4), $V_{bvdf} \approx 1$, conversely when light and camera are on opposing sides of the cloth (pink and red 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. This observation is supported by the testing results summarized in Fig. 8 (more details can be found in the supplemental material). 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. A valuable side effect of using the BVDF is that rendering times are reduced since fewer costly shadow-rays need to be evaluated. Fig. 5 shows that the BVDF helps significantly in obtaining the correct color and brightness of the piece of cloth.

METHOD	MEM	IBL	POINT LIGHT	CONSTR.
reference (20M lines)	10GB	07:32 min	03:16 min	7:20 min
level 9 (2.5M voxels)	177MB	02:27 min	01:19 min	1:28 min
level 8 (400k voxels)	30MB	02:18 min	01:15 min	0:54 min
level 7 (75k voxels)	5.6MB	02:12 min	01:12 min	0:45 min
level 6 (13k voxels)	1.1MB	02:09 min	01:10 min	0:36 min

Table 1: *Comparing average rendering times for image-based lighting (IBL) using an environment map and for a single point light located at the camera position and comparing memory consumption for different octree resolutions. Times for constructing the kD-tree and for voxelization are noted under the name "CONSTR.". All times were taken on a Core i7 CPU operating at 3.07 Ghz, rendering images of size 300x210 with 512 samples per pixel. Note that although the kD-tree implementation is generally very efficient, its construction process has not been parallelized.*

6. Results

We have tested our approach with synthetic cloth models with varying weave patterns (see Fig. 6), different types of yarn geometry and realistic and at the same time challenging BCSDFs matching glossy dielectric fibers. To better identify potential weaknesses, cloth was coiled around a cylinder. For the same reason, we have selected examples with a very regular fiber structure. The geometric complexity of the models was taken such that a fiber-based reference rendering could be created on a computer with 12GB RAM.
All images have been simulated with full global illumination based on Monte Carlo path tracing. As can be seen from the close-ups in Fig. 6, the colors of certain yarns (i.e. the red ones in the "colorful cloth example") are largely determined by the multiple scattering component – therefore, correct reproduction of this color is far from being trivial. For other yarns like the green ones with higher absorption values, multiple scattering is less prominent. The results given in Fig. 6 indicate that the volumetric approach delivers results which are visually almost identical to the reference solution for the highest octree level with 2 million voxels while already needing much less memory. Levels 8 (400k voxels) and 7 (80k voxels) only show some (expected) blurring caused by the discretization, but the overall appearance is captured faithfully and consistently across different voxel resolutions. For level 7 the cross section of a yarn is cov-

ered by only slightly more than a single voxel. In this figure, voxels are projected to several pixels to make the artifacts which are introduced by the discretization visible – in practice one would choose a level where the voxel size is selected such that voxels are projected to no more than (for example) a single pixel. At Level 6 (10k voxels) the length of a voxel edge is slightly larger than the yarn diameter – artifacts are introduced because the separation of yarns cannot be resolved correctly (for example, if a strand of yarn passes behind another, it shows through and the distributions of the two strands become mixed). Although we have mainly included this level for illustrational purposes in the figure, even this coarse representation might sometimes be sufficient for images of distant clothes as the overall color impression is still reproduced. In addition to the cylinder images, we have simulated BRDFs for the reference technique as well as our approximation based on a small patch of cloth. Two exemplary results for fixed lighting directions are illustrated in Fig. 7: Reference and approximation match very well, regardless of the viewing direction.

Note that, assuming a fixed spatial resolution, the difference in memory and computational costs is even more dramatic for more complex dense models. This is because the costs for the novel method are bounded by the number of voxel cells, whereas the explicit model requires memory at least linear in the size of the fiber primitives as well as costly acceleration data structures (such as a kD-tree) for rendering. For example in the case of staple yarns made from several plies the memory savings could be much higher still. Please note that the comparison of rendering times may be biased in favor of the explicit technique: it already becomes completely unusable due to excessive memory requirements for geometries which can still be regarded as trivial for the volumetric case. It should also be noted that the octree has been optimized for size (i.e. no voxel coordinates are stored explicitly) whereas the kD-tree has been optimized for speed (and e.g. stores bounding box coordinates). Several examples of large pieces of cloth that — in contrast to the volumetric method — are prohibitively costly to render with the fiber based approach are presented in Fig.1. The curtain in the background also serves as an example for a transparent piece of cloth.

7. Limitations

Due to the assumptions made, our approach has 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 which 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.

Our method is not effective in case of extreme close-ups where individual yarn fibers become visible as a prohibitively high spatial resolution would be required to resolve such fine geometric detail. The BVDF, as formulated above, only describes visibility for points which are hit by eye-rays – although this 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 which are shot during multiple scattering that have an origin near the surface might profit from a BVDF representation as well.

8. Conclusion and Future Work

In this paper, a physically-based alternative to fiber-based rendering of fabrics has been presented. As for methods using an explicit representation of fiber geometry, our approach allows a simulation of light scattering based on small scale optical properties of yarn fibers. By taking a statistical model, the memory as well as computational costs were greatly reduced. We are convinced that our approach will be valuable for applications (e.g. virtual prototyping of cloth) where a high degree of realism is desirable but fiber-based simulations are not practical.

We believe that there is still potential for increasing the efficiency of the method in several different ways: First of all, as some types of cloth exhibit a repetitive weave pattern, Gaussian mixture components could be computed only once per pattern and referenced accordingly. Moreover, radiance caching techniques could be applied, that first store average radiance values in the voxel and then attempt to decrease variance by combining results across similar repetitions of a weave pattern. Two straightforward extensions would be the use of an adaptive octree (which tries to keep yarns separate and approximates the silhouette to some specified degree but otherwise uses as few voxels as possible) and utilizing the hierarchical structure for level of detail by not only storing Gaussian mixtures in leaf voxels but also for lower levels. It would be interesting to investigate what other acceleration data-structures could be applied: The reference technique does not perform that much slower in our examples although it has to perform intersection tests which are a lot more costly compared to sampling our statistical model – this is mainly due to the efficiency of the kD-tree compared to the octree.

Cloth can include other effects like dirt, brighteners (these can be modeled to some degree using the BCSDF) or a piece of cloth could simply be wet – future work could go into modeling these effects as well. Due to modeling constraints, most presented images contain repeating structures. However, this is no limitation of our volumetric representation which can handle arbitrary fiber-based input data. For some types of *fancy yarns*, a spatially varying BVDF could actually be required. As the BVDF varies smoothly with respect to yarn properties, interpolation between a few basis BVDFs could be applied.

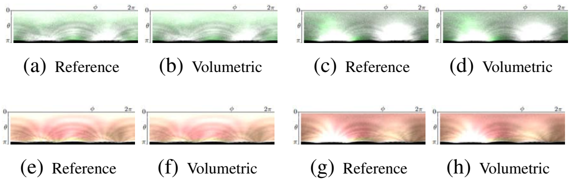

(a) Reference (b) Volumetric (c) Reference (d) Volumetric

(e) Reference (f) Volumetric (g) Reference (h) Volumetric

Figure 7: *Simulated reflectance fields (slices of the BRDF) of a quadratic piece of cloth consisting of 30 warp and 30 weft yarns for two fixed lighting directions showing "green cloth" and "colorful cloth". Direction is parametrized by spherical angles $0 < \phi < 2\pi$ and $0 < \theta < \frac{\pi}{2}$. Incident light direction is $\phi = \theta = 0$ for (a),(b),(e),(f) and $\phi = \theta = \frac{\pi}{4}$ for (c),(d),(g),(h). Reference (7M line segments) on the left of each pair / Volumetric (400K voxels) on the right.*

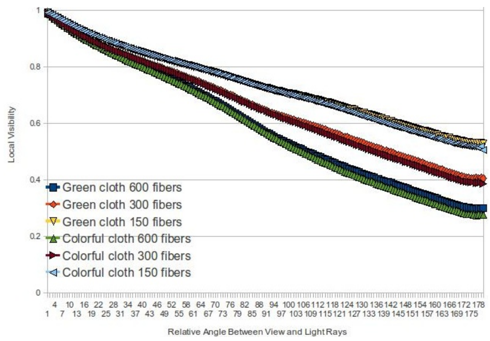

Figure 8: *Local visibility for two different cloth samples (400K voxels examples of Fig.6) with varying numbers of fibers per yarn: For illustrational purposes only the relative angle between light and view direction is considered. Visibility almost only depends on yarn properties, independent of the weave pattern (at this scale). Despite the significant amount of change in fiber density, visibility only changes moderately. Even for solid pieces of cloth, it will not reach 0 (fully occluded) for 180° (where light and view direction oppose each other) as a significant amount of occlusion is modeled by the global factor.*

So far, we have demonstrated the effectiveness of our method only by comparing it to synthetic reference images. However — although we believe that this comparison is fair as it reflects the current state-of-the-art in the field — measuring real cloth samples and using this information for validation will be an interesting topic of future research. One might also consider using the technique for other materials for which a similar statistical model is appropriate.

Acknowledgments: This work was supported in part by NRW State within the B-IT Research School.

References

[AMTF03] Adabala N., Magnenat-Thalmann N., Fei G.: Real-time visualization of woven textiles. In *Publication of EUROSIS* (2003), pp. 502–508. 3

[DLH01] Daubert K., Lensch H., Heidrich W.: Efficient cloth modeling and rendering. In *Rendering techniques 2001, United Kingdom, June 25-27, 2001* (2001), p. 63. 3

[Gla02] Glassner A.: Digital weaving, part 1. *IEEE Computer Graphics and Applications* (2002), 108–118. 4

[GRS95] Groller E., Rau R. T., Strasser W.: Modeling and visualization of knitwear. *IEEE Transaction on Visualization and Computer Graphics 1*, 4 (1995), 302–310. 3

[Ira07] Irawan P.: *Appearance of Woven Cloth*. PhD thesis, Cornell University, 2007. 3

[JAM*10] Jakob W., Arbree A., Moon J. T., Bala K., Marschner S.: A radiative transfer framework for rendering materials with anisotropic structure. *ACM TOG 29* (July 2010), 53:1–53:13. 2, 3

[KJM08] Kaldor J., James D., Marschner S.: Simulating knitted cloth at the yarn level. In *ACM SIGGRAPH 2008 papers* (2008), ACM, pp. 1–9. 4

[KJM10] Kaldor J., James D., Marschner S.: Efficient yarn-based cloth with adaptive contact linearization. In *ACM SIGGRAPH 2010 papers* (2010), ACM, pp. 1–10. 4

[MJC*03] Marschner S. R., Jensen H. W., Cammarano M., Worley S., Hanrahan P.: Light scattering from human hair fibers. *ACM TOG 22*, 3 (2003), 780–791. SIGGRAPH 2003. 2

[MM06] Moon J. T., Marschner S. R.: Simulating multiple scattering in hair using a photon mapping approach. *ACM TOG 25*, 3 (2006), 1067–1074. SIGGRAPH 2006. 3

[MTCK*04] Magnenat-Thalmann N., Cordier F., Keckeisen M., Kimmerle S., Klein R., Meseth J.: Simulation of clothes for real-time applications. In *Eurographics 2004, Tutorials* (2004), INRIA and the Eurographics Association. 3

[MWM08] Moon J. T., Walter B., Marschner S.: Efficient multiple scattering in hair using spherical harmonics. *ACM TOG 27*, 3 (2008). SIGGRAPH 2008. 2, 3, 4, 5

[SB06] Sreprateep K., Bohez E.: Computer Aided Modelling of Fiber Assemblies. *Comput. Aided Des. Appl 3*, 1-4 (2006), 367–376. 4

[VKKK97] Volevich V. L., Kopylov E. A., Khodulev A. B., Karpenko O. A.: An approach to cloth synthesis and visualization. In *The 7-th International Conference on Computer Graphics and Visualization* (1997). 3

[WAT92] Westin S. H., Arvo J. R., Torrance K. E.: Predicting reflectance functions from complex surfaces. *SIGGRAPH Comput. Graph. 26*, 2 (1992), 255–264. 3

[WZT*08] Wang J., Zhao S., Tong X., Snyder J., Guo B.: Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph. (to appear in SIGGRAPH 2008) 27*, 3 (2008), 41:1–41:9. 3

[XCL*01] Xu Y.-Q., Chen Y., Lin S., Zhong H., Wu E., Guo B., Shum H.-Y.: Photorealistic rendering of knitwear using the lumislice. In *SIGGRAPH '01* (New York, NY, USA, 2001), ACM Press/Addison-Wesley Publishing Co., pp. 391–398. 3

[ZLHA*09] Zinke A., Lay Herrrera T., Andriyenko A., Rump M., Weber A., Klein R.: A practical approach for photometric acquisition of hair color. *ACM TOG* (Dec. 2009). 2

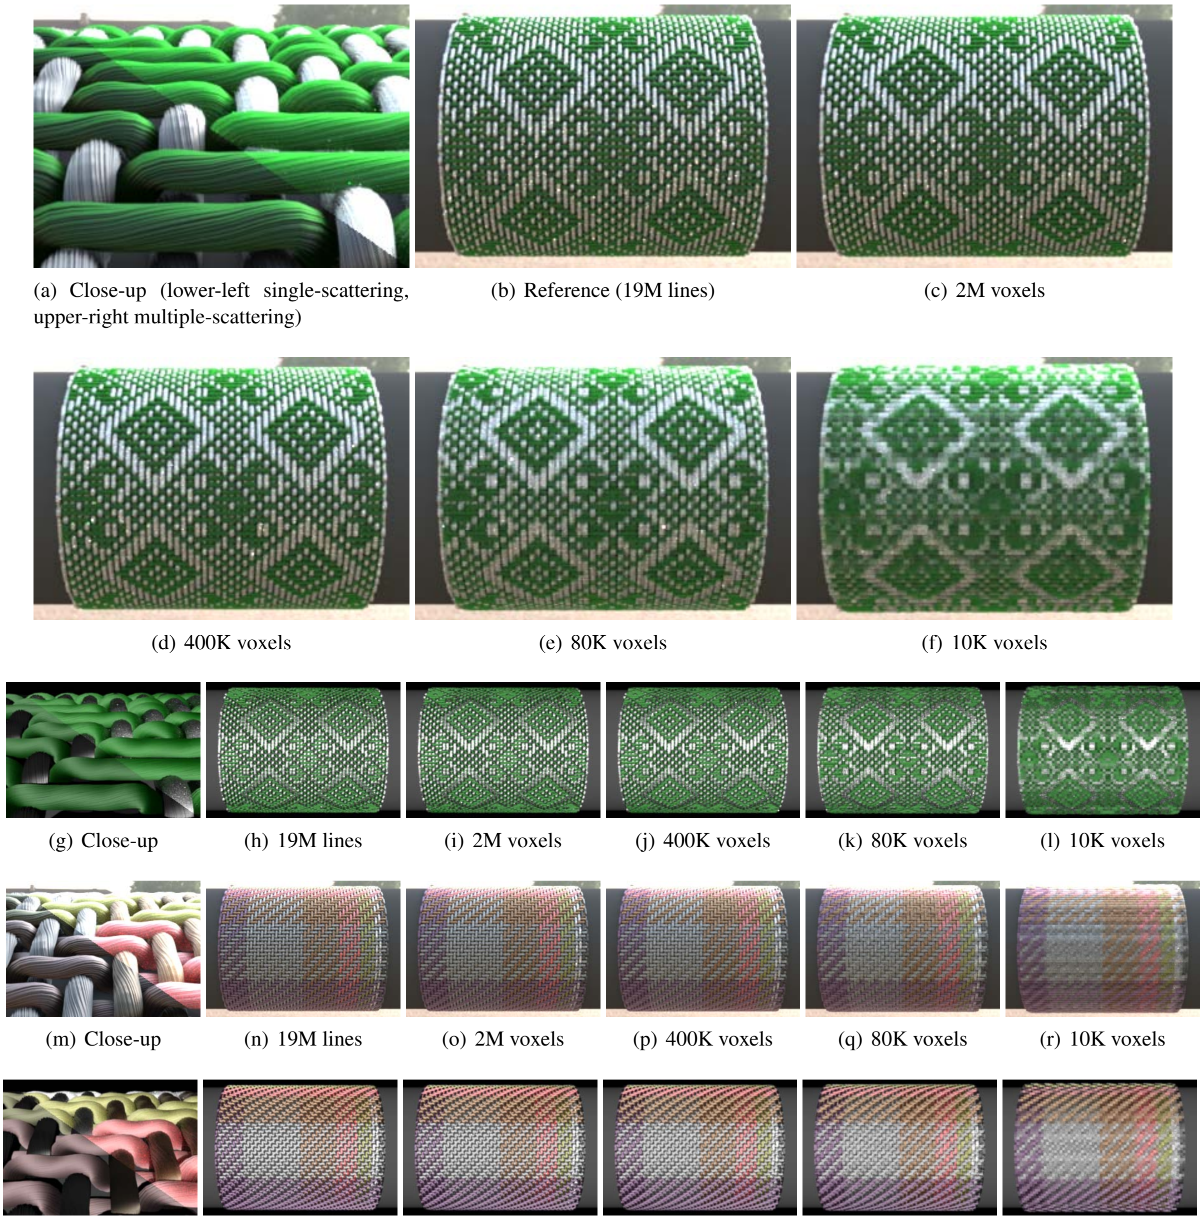

(a) Close-up (lower-left single-scattering, upper-right multiple-scattering) (b) Reference (19M lines) (c) 2M voxels

(d) 400K voxels (e) 80K voxels (f) 10K voxels

(g) Close-up (h) 19M lines (i) 2M voxels (j) 400K voxels (k) 80K voxels (l) 10K voxels

(m) Close-up (n) 19M lines (o) 2M voxels (p) 400K voxels (q) 80K voxels (r) 10K voxels

(s) Close-up (t) 19M lines (u) 2M voxels (v) 400K voxels (w) 80K voxels (x) 10K voxels

Figure 6: *"green cloth" and "colorful cloth" - Illustrating the effect of changing the octree resolution. Each piece of cloth is lit using a sunny outdoor environment map and a point light coming from the camera direction. For a better comparison, the image resolution is kept constant across different resolutions. All presented cloth samples have 400 fibers per yarn with a fixed index of refraction (1.5) and varying absorption coefficients using the BCSDF model of [ZW07]. The variance is comparable for the reference images and the volumetric method.*

[ZSW04] Zinke A., Sobottka G., Weber A.: Photo-realistic rendering of blond hair. In *VMV 2004* (Stanford, U.S.A., Nov. 2004), pp. 191–198. 3

[ZW07] Zinke A., Weber A.: Light scattering from filaments. *IEEE TVCG 13*, 2 (2007), 342–356. 2, 4, 10

[ZYWK08] Zinke A., Yuksel C., Weber A., Keyser J.: Dual scattering approximation for fast multiple scattering in hair. *ACM TOG 27*, 3 (2008). SIGGRAPH 2008. 3