

Building Volumetric Appearance Models of Fabric using Micro CT Imaging

Shuang Zhao

Wenzel Jakob

Steve Marschner

Kavita Bala

Cornell University*




Figure 1: We build volumetric appearance models of complex materials like velvet using CT imaging: (left) CT data gives scalar density over a small volume; (center) we extract fiber orientation (shown in false color) and tile larger surfaces; and (right) we match appearance parameters to photographs to create a complete appearance model. Both fine detail and the characteristic highlights of velvet are reproduced.

Abstract

The appearance of complex, thick materials like textiles is determined by their 3D structure, and they are incompletely described by surface reflection models alone. While volume scattering can produce highly realistic images of such materials, creating the required volume density models is difficult. Procedural approaches require significant programmer effort and intuition to design special-purpose algorithms for each material. Further, the resulting models lack the visual complexity of real materials with their naturally-arising irregularities.

This paper proposes a new approach to acquiring volume models, based on density data from X-ray computed tomography (CT) scans and appearance data from photographs under uncontrolled illumination. To model a material, a CT scan is made, resulting in a scalar density volume. This 3D data is processed to extract orientation information and remove noise. The resulting density and orientation fields are used in an appearance matching procedure to define scattering properties in the volume that, when rendered, produce images with texture statistics that match the photographs. As our results show, this approach can easily produce volume appearance models with extreme detail, and at larger scales the distinctive textures and highlights of a range of very different fabrics like satin and velvet emerge automatically—all based simply on having accurate mesoscale geometry.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—Three-Dimensional Graphics and Realism

Keywords: appearance modeling, volume rendering, cloth

Links: [DL](#) [PDF](#) [WEB](#)

*E-mail: {szhao, wenzel, srm, kb}@cs.cornell.edu

1 Introduction

The appearance of materials like cloth is determined by 3D structure. Volume rendering has been explored for decades as an approach for rendering such materials, for which the usual surface-based models are inappropriate [Kajiya and Kay 1989; Perlin and Hoffert 1989; Xu et al. 2001]. Recent developments [Jakob et al. 2010] have brought enough generality to volume scattering that we can begin to render fully physically-based *volumetric appearance models* for cloth, fur, and other thick, non-surface-like materials. However, a fundamental problem remains: creating these volumetric models themselves. For surfaces, texture maps derived from photographs are simple and effective, but volumes are not so easy. Previous work has primarily relied on procedural methods for modeling volume density, but this has limited generality: significant creative effort is needed to design special algorithms for each new material. Further, these models often miss the subtle irregularities that appear in real materials.

This paper explores an entirely different approach to building volume appearance models, focusing particularly on cloth. Since cloth’s detailed geometric structure is so difficult to model well, we use volume imaging to measure structure directly, then fill in optical properties using a reference photograph. We do this by solving an inverse problem that statistically matches the texture between photographs and physically based renderings (which include global illumination and multiple scattering). We focus on textiles because they exhibit a wide range of appearance, but share a common basic structure of long, shiny fibers. Textile rendering is important for many applications, but is challenging because cloth is structured, causing complicated textures and reflectance functions, yet irregular, causing difficult-to-model randomness. The thick, fuzzy nature of cloth makes volume models a good fit, if only there were a general solution for constructing them.

Many volume imaging technologies have been developed, including computed tomography (CT), magnetic resonance, ultrasound, and others, but unlike photographs, the resulting data does not directly relate to the optical appearance of the material; only to its structure. As a result, volume renderings of these images are useful for illustrating hidden internal geometry, but not directly for rendering realistic images. For instance, a micro CT scan of woven cotton cloth gives a detailed view of the interlaced yarns and their component fibers, showing exactly how the fibers are oriented and how the yarns are positioned, but no information about how they in-




Figure 2: Our volume appearance modeling pipeline. (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 our acquired volumetric appearance and geometry models.

interact with light: there is no way to tell whether the fabric is black or white or any color in between.

We show in this paper that remarkably little additional information is required to extend CT data to a realistic appearance model. The value of knowing 3D structure is obvious for rendering close-up views where these details are visible. But equally importantly, the shape and arrangement of fibers in the material also determines the overall appearance of the material—the shape and quality of specular highlights, and how the visual texture varies with illumination and view. When coupled with the right rendering technology, a simple local model of reflection from fibers automatically predicts the characteristic appearance of very different materials like velvet and satin, simply by knowing the 3D structure of the material.

The contribution of this paper is to show how to enhance the structural information from a CT scan of a small sample of fabric by combining it with appearance information from a photograph of the material to construct plausible and consistent optical properties that produce realistic appearance when rendered using a physically based volume renderer. We describe our end-to-end volume appearance modeling pipeline and demonstrate it by acquiring models of cloth with very different appearance, ranging from matte to shiny and textured to smooth, capturing their characteristic highlights, textures, and fuzziness.

2 Related Work

We categorize realistic volumetric rendering and modeling research in the related areas of surface appearance modeling, cloth reflectance modeling, and cloth structure modeling.

Appearance modeling: Because standard surface-oriented models are inadequate for complex thick materials, researchers and practitioners have had to fall back on image-based rendering methods like Bidirectional Texture Functions (BTF), which essentially consist of an exhaustive set of photographs of the surface under all possible illumination and viewing directions [Dana et al. 1999; Furukawa et al. 2002; Wang et al. 2005]. Although BTFs produce realistic results for many otherwise difficult materials, the image-based approach requires a significant amount of storage, and is often not of high enough resolution for sharp BRDF features, and generally fails to capture or predict grazing angles, making silhouettes and edges unrealistic.

Two prominent early volume appearance models are Kajjiya and Kay’s [1989] fur rendering, and Perlin and Hoffert’s [1989] “hyper-texture.” Although it has since become more common to render hair and fur using discrete curves, their results demonstrate the value of volumetric models for complex, barely resolved detail. A similar approach is the Lumislice representation [Xu et al. 2001; Chen et al.

2003] which focused on modeling and rendering knitwear. Magda and Kriegman [2006] describe a method for acquiring *volumetric textures* which combine a volumetric normal field, local reflectance functions, and occupancy information. All these approaches need significant modeling effort. Recently Jakob et al. [2010] introduced a principled formulation for rendering anisotropic, oriented volumetric media, which opens possibilities for more physically based volume appearance models.

Cloth reflectance models: Cloth has perennially appeared in graphics as a source of difficult BRDFs. Westin et al. [1992] computed cloth BRDFs by raytracing mesostructure models, which is related to the way cloth highlights emerge in our system. Ashikhmin et al. [2000] rendered velvet and satin using hand-designed microfacet distributions. Adabala et al. [2003] proposed a rendering method for woven cloth based on microfacet models, and Irawan [2008] presented an elaborate model, based on the analysis of fiber tangent directions in a range of woven fabrics, and validated it against BRDF measurements. Each of these methods achieved good appearance relative to the then-current state of the art, but they are all specially hand-designed models for individual materials or specific classes. Lu et al. [1998] measured and analyzed reflections from velvet, and Ngan et al. [2005] measured some fabrics, including satins, but neither proposed models suitable for rendering.

Since our approach is based on a completely general system that only has a volume with fibers as its underlying assumption, we have few fundamental limitations on what textile or textile-like materials can be handled. Further, by importing volumetric detail from the real world, we can achieve good appearance in closeups, and at silhouettes, edges, and corners, where surface models appear unrealistically smooth and flat.

Cloth structure: The geometry of cloth structure has been studied for decades [Pierce 1937; Kawabata et al. 1973]. More recently X-ray tomography, using synchrotron facilities [Thibault and Bloch 2002; Gong et al. 2009] or the rapidly improving micro-CT scanners [Lomov et al. 2002; Shinohara et al. 2010], has been used to examine the structure of textiles in several applications. These studies focus on extracting geometric information related to the material’s mechanical properties, but have produced some analysis tools [Shinohara et al. 2010] that we use.

3 Overview

The goal of our system is to create realistic volumetric appearance models of cloth. We need to generate a sampled 3D volume that describes the optical properties of the material at each voxel so that, when rendered with a physically based rendering system, it realistically reproduces the appearance of real cloth.

Because cloth is made of fibers, we need a volume scattering model that can handle the anisotropy of fibers; we chose a modified version of the model proposed by Jakob et al. [2010] (detailed in Section 4) for this purpose. This model requires an optical density, an albedo, and two phase function parameters: an orientation vector and a specular lobe width.

Our technique begins with a micro CT scan of a small area of material, showing detail at the level of individual fibers over a fraction of a square centimeter. Such scans can readily be ordered at moderate cost (a few hundred US dollars) from a number of facilities, and suitable desktop CT scanners are becoming available. In a sequence of three stages (Figure 2) we process and augment this data, ending with a volume that defines the required scattering model parameters using density and orientation fields derived from the CT data, plus three global parameters: the albedo, the lobe width, and a density multiplier that scales the density field.

The first stage (Section 5) processes the density volume to augment it with orientation information and to remove noise by convolving the data with 3D oriented filters to detect oriented structures, and thresholding to separate meaningful structure from noise. This stage produces the density and orientation fields.

This volume can be rendered only after the global optical parameters are determined. The second stage (Section 6) makes use of a single photograph of the material under known (but not controlled) lighting, and associates optical properties with the oriented volume from the first stage by matching the texture of the rendered volume to the texture of the photograph.

The resulting volume model is good for rendering small samples; the third stage takes this small patch and maps it over a large surface of cloth, using randomized tiling to replicate the material and shell mapping [Porumbescu et al. 2005] to warp it.

The resulting renderings (Sections 7 and 8) show that this unique approach to appearance modeling, leveraging direct information about mesoscale geometry, produces excellent appearance from the small scale, where the geometry itself is visible, to the large scale, where the directional scattering properties naturally emerge from the measured 3D structure. The characteristic appearance of difficult materials like velvet and satin is predicted by our rather minimal volume scattering model, even though we use no light scattering measurements that could tell these materials apart, because accurate geometric information is available.

4 Fiber Scattering Model

We model light transport using the anisotropic radiative transfer equation (RTE) from Jakob et al. [2010]

$$(\omega \cdot \nabla)L(\omega) + \sigma_t(\omega)L(\omega) = \sigma_s(\omega) \int_{S^2} f_p(\omega' \rightarrow \omega)L(\omega') d\omega' + Q(\omega) \quad (1)$$

where σ_s and $\sigma_t : S^2 \rightarrow \mathbb{R}$ are the anisotropic scattering and extinction coefficients, f_p is the phase function, and Q is the source term. Spatial dependence has been omitted for readability.

This equation can be understood as a generalization of the isotropic RTE that adds support for a directionally varying amount of “interaction” with a medium. For instance, the directional dependence of $\sigma_t(\omega)$ is necessary to model the effect that light traveling parallel to coherently aligned fibers faces less obstruction than light traveling perpendicular to the fibers.

To specify the problem to be solved, we must choose a compatible scattering model that will supply internally consistent definitions

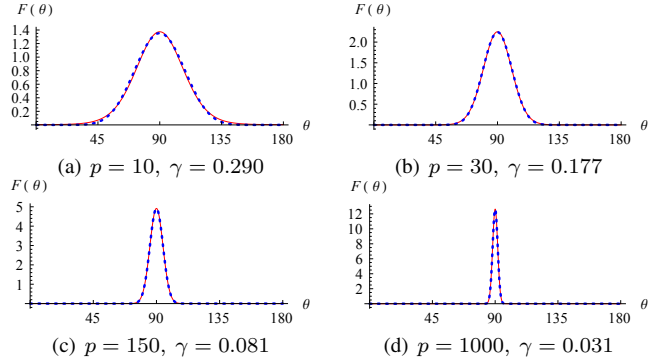


Figure 3: Comparison between the $\sin^p \theta$ -type distribution with exponent p (blue) and our Gaussian-type distribution (red) with standard deviation γ .

of σ_t , σ_s , and f_p . For this purpose, we use the *micro-flake* model proposed in the same work. This volume analogue of microfacet models represents different kinds of volume scattering interactions using a directional *flake distribution* $D(m)$ that describes the orientation m of (unresolved) idealized mirror flakes at every point in space. Similar to microfacet models, the phase function then involves evaluating $D(m)$ at the half-way direction between the incident and outgoing direction. For completeness, we reproduce the model’s definition below:

$$\sigma_t(\omega) = a \rho \int_{S^2} |\omega \cdot m| D(m) dm$$

$$\sigma_s(\omega) = \alpha \sigma_t(\omega)$$

$$f_p(\omega' \rightarrow \omega) = \frac{a \rho \alpha}{4 \sigma_s(\omega)} (D(h(\omega, -\omega')) + D(-h(\omega, -\omega')))$$

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 , Jakob et al. propose the flake distribution $D(\omega) = c_0 \sin^p(\omega_f, \omega)$, 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 ; the underlying motivation is to represent the normal directions observed on the original fiber’s surface, which predominantly point in these directions.

4.1 Alternative flake distribution

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 . Since our rendering pipeline crucially depends on this ability (see Section 7), we propose an alternative flake distribution that is convenient to integrate, while capturing the same key feature of the \sin^p model, namely that it is primarily concentrated perpendicular to the fiber direction.

We use the following density function, which specifies a truncated

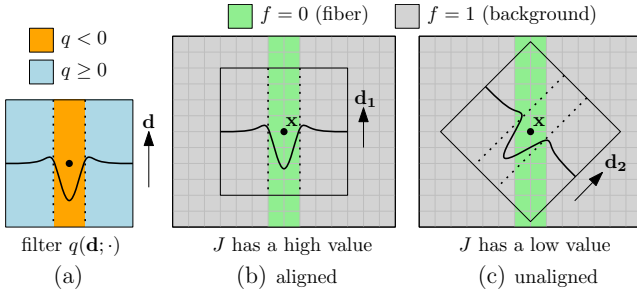


Figure 4: Computing function J in 2D: (a) shape of the filter q ; (b) when q is aligned to the fiber; (c) when q is unaligned.

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. The model captures the same qualitative behavior as the \sin^p model over a large range of parameter values (Figure 3).

To summarize, the parameters required to create renderings are:

- ω_f , the local fiber orientation,
- γ , the standard deviation of the flake distribution,
- α , the single scattering albedo of the flakes,
- a and ρ , the area and density of micro-flakes. Their product roughly corresponds to the interaction coefficient σ_t in traditional isotropic volume rendering, and we therefore set them to a multiple of the processed CT densities, i.e. $a\rho(\mathbf{x}) := d \cdot \text{CT}(\mathbf{x})$, where d is a constant of proportionality.

Section 5 discusses the steps needed to obtain $\text{CT}(\mathbf{x})$ and $\omega_f(\mathbf{x})$. In Section 6, we describe how to find α , γ , and d , and Section 7 explains how to use our scattering model in Monte Carlo rendering.

5 CT Image Processing

Micro CT (computed tomography) 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.

In this section 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.

5.1 Recovering the Orientation Field

CT images provide a voxelized density field with no direction information. Since our optical model requires an orientation for the phase function, it is necessary to reconstruct an orientation for every non-empty voxel. Our approach uses oriented filters to detect fibers, based on similar filters used by Shinohara et al. [2010] to locate fibers in CT data. We chose this approach because of its demonstrated application to fiber detection in CT data, though alternatives [Axelsson 2008] are possible.

To detect a fiber with orientation \mathbf{d} at location \mathbf{p} , Shinohara proposes a cylindrically symmetric filter oriented with the axis \mathbf{d} , con-




Figure 5: Computed orientation field for a piece of gabardine with each direction (x, y, z) mapped to RGB color $(|x|, |y|, |z|)$. Left: without thresholding on J ; right: with thresholding on J .

sisting 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 4).

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:

$$J(\mathbf{x}, \mathbf{d}) := \sum_{\mathbf{p} \in V} q(\mathbf{d}; \mathbf{p}) f(\mathbf{x} + \mathbf{p}) \quad (2)$$

where V is a cubic volume of edge length h . For parameter values, refer to Table 1.

As shown in Figure 4, the function J reaches a maximum value when \mathbf{d} equals the fiber's orientation. So the orientation field is computed by finding, for each voxel \mathbf{x} , the \mathbf{d}' that maximizes $J(\mathbf{x}, \mathbf{d}')$ and setting $\omega_f(\mathbf{x}) = \mathbf{d}'$. In our implementation, we pre-compute q on a set of directions $\{\mathbf{d}_i\}$ picked from a $32 \times 32 \times 6$ cubemap. Then for each non-empty voxel \mathbf{x} , we set $\omega_f(\mathbf{x}) = \mathbf{d}_j$ where $j = \arg \max_i J(\mathbf{x}, \mathbf{d}_i)$.

5.2 Denoising CT Images

The CT images usually contain considerable amounts of noise, particularly for low-density materials like our cloth samples, and removing the noise is critical for obtaining good quality data for rendering. Since cloth structure is always oriented, and the noise is generally fairly isotropic, the value of J is useful in noise removal.

In our system we use two thresholds to remove noise. The first threshold ϵ_d is on the voxel values themselves, and is used to remove faint background noise that would otherwise cloud the model. This thresholding creates the binary volume f . The second threshold ϵ_J is on the value of J and is used to remove isotropic noise that has density values that are too high to remove by the first threshold. We set

$$\text{CT}(\mathbf{x}) := \begin{cases} \text{CT}_{\text{raw}}(\mathbf{x}) & \text{CT}_{\text{raw}}(\mathbf{x}) \geq \epsilon_d \text{ and } J(\mathbf{x}, \omega_f(\mathbf{x})) \geq \epsilon_J; \\ 0 & \text{otherwise.} \end{cases}$$

Figure 5 shows the significance of adding this second threshold.

5.3 Data Replication

The volume data needs to be replicated for rendering since our samples are very small. Texture synthesis methods for images provide sophisticated tools that could be extended to do this, but are beyond the scope of this paper, and an area of future work. Instead we used two simple randomized tiling methods to cover the surfaces with tiles of volume data drawn from our models without introducing distracting regular structures. In both methods the surface is simply covered by a rectangular array of tiles copied from the volume, without continuity at the tile boundaries.

For materials without visible regularity, such as velvet and felt, each tile on the surface is copied from a rectangular region centered in the volume. To provide variation in local structure, for each tile this source rectangle is rotated by a different random angle. For materials with woven structure, like silk and gabardine, we use a similar approach, but use random translations of the source tile instead of rotations. The weave pattern in each sample is manually identified and a rectangular area is marked that contains an integer number of repeats. Then each (smaller) surface tile is chosen from a sub-rectangle that contains a matching section of the weave. The result is a tiling that reproduces the correct weave pattern and avoids obvious repeating of texture. We then map the tiled data to arbitrary surfaces using shell mapping [Porumbescu et al. 2005].

6 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 us the material’s optical properties, we make use of a photograph of the material to compute the appearance parameters.

To make the problem tractable, we assume that the volume contains the same material, with differences only in density and orientation. This is appropriate for fabrics made from a single type of fiber, which encompasses many important examples. Fabrics containing yarns of different materials are future work. 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 6-(a) illustrates the effects of these parameters.

To match the material’s optical properties, we must use photographs of the sample. One approach is to photograph the same sample that was scanned, calibrating the camera to the scan and associating pixels in the image with rays in the volume. This calibration and acquisition is non-trivial; the fine resolution of the scans poses practical difficulties. Further, we found that this level of detail is not required to determine the small number of parameter values we seek. Instead, we assume that the fabric is statistically similar across different patches. Thus, our approach is to statistically match the texture of rendered images with a photograph of a different section of the same cloth under uncontrolled but known lighting.

We now describe the metrics we use to match the optical parameters to the photograph, and then describe our matching algorithm.

6.1 Metrics for matching

Appearance matching is not a straightforward process of mapping colors from the photos into the volume, because the volume model describes local scattering properties, but the appearance is defined by a global volumetric multiple scattering process. Our approach




Figure 6: (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.

is to repeatedly render the volume using our physically based renderer, and adjust the optical parameters to match certain texture statistics of the rendered images to statistics of the photograph.

We match two simple statistical measures: the mean pixel value and the standard deviation of pixel values, computed over corresponding regions of a photograph and a rendering of approximately similar geometry. This approach effectively matches the image brightness and texture contrast in the matching region. We tried measures such as the CDF of intensities [Heeger and Bergen 1995] and skewness [Motoyoshi et al. 2007], but found that the mean and standard deviation measures were simpler and robust. Thus, the only information that flows from the photograph to the volume model is the mean and standard deviation of pixels in a single rectangle.

The appearance matching process involves choosing the geometry, camera position, lighting, and matching region. These are inherently manual choices, and we used the principle of choosing a setup that shows the distinctive features of the cloth’s appearance. For instance, we made sure to use a configuration where the highlight was visible on the satin. Beyond this we did not take any special care in arranging the appearance matching inputs, and the results do not appear to be sensitive to the details.

6.2 Optimization procedure

As shown in Figure 6, the density multiplier plays a fairly complicated role with respect to both measures. Given that our forward process, which is essentially Monte Carlo path tracing, is quite expensive, we chose to pre-determine the density multiplier in our implementation by rendering such a matrix. Fixing the density multiplier simplifies the inverse problem and leads to a practical solution. We found that the algorithm is not particularly sensitive to the choice of density multiplier; our results use two main settings which differ by an order of magnitude (see Table 1).

With a fixed density multiplier, we solve for the values of albedo (α , estimated separately in red, green, and blue) and lobe width (γ ,




Figure 7: Appearance matching results for (from top to bottom) (1) silk, (2) gabardine, (3) velvet, and (4) felt. Columns (a) and (c) show photographs of the materials, and (b) and (d) show rendered images. The left two columns form the appearance matching pair, in which the blue boxes indicate manually selected regions for performing our matching algorithm. The right two columns, the validation pair, validate our matches qualitatively under different configurations.

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 respectively¹. Thus, a binary search can be used to significantly improve performance as follows: first, an initial guess of γ is assumed, and we search for the α to match the mean pixel value. Then, fixing α , we perform a search for the γ to match the standard deviation. These iterations are repeated until a match is found. In practice, this approach converges quickly, usually in 2 or 3 iterations.

Finally, we take another photo under a different setup and render a corresponding image as a qualitative validation (see Section 8). Figure 7 shows the appearance matching results for four different materials.

7 Rendering

We render all our scenes using a basic Monte Carlo path tracer, which handles the directionally varying properties of the medium described in Section 4. One important part of this process entails generating samples from the phase function $f_p(\omega' \rightarrow \omega)$. In this section, we adopt the notation of the integral form of the RTE, i.e. ω is held fixed, and we are interested in sampling the direction ω' , from which to gather illumination.

In prior work, [Jakob et al. 2010] used a spherical harmonics representation for this purpose. However, this approach has several undesirable properties. First, a sample weight is needed to account for the fact that the sampling routine is only approximate, which

¹This holds as long as γ exceeds a minimum value ($\gamma = 0.01$ for all our experiments); below this value the variance of fiber orientations limits glossiness.

increases variance. Second, computing the spherical harmonics coefficients is a time-consuming process, which needs to be repeated for any change in the model parameters. This is problematic, since our fitting stage explores many different parameter sets. Most importantly, the spherical harmonics approach suffers from ringing and therefore cannot handle some of the highly specular material configurations in our parameter space.

In the following section, we first propose a naïve sampling method, which is not directly usable due to its high variance. We then demonstrate how rejection sampling can be used to turn the naïve method into an exact sampling scheme.

7.1 Alternative sampling strategy

In the surface case, importance sampling for microfacet models often takes the approach of sampling a microfacet normal, then using it to compute an outgoing direction [Walter et al. 2007]. Additional factors, such as the Fresnel reflectance and the Jacobian of the direction mapping must be accounted for in a weight associated with the sample. If we apply this approach to the micro-flake model, we obtain a sampling strategy with the following density:

$$f_1(\omega' \rightarrow \omega) = \frac{D(h(\omega, -\omega'))}{2|\omega' \cdot h(\omega, -\omega')|}$$

where the denominator is the aforementioned Jacobian (an extra factor of 2 is required in comparison to the surface case, since micro-flakes reflect from both sides). The sample must be assigned the weight:

$$w_1(\omega' \rightarrow \omega) = \frac{f_p(\omega' \rightarrow \omega)}{f_1(\omega' \rightarrow \omega)} = \frac{a\rho}{\sigma_t(\omega)} |\omega' \cdot h(\omega, -\omega')| \quad (3)$$

Material	Data Size	s	t	h	ϵ_d	ϵ_J	d	γ	α
Gabardine	$992 \times 1012 \times 181$	1	2	16	0.45	-10	5000	0.1	(0.892, 0.063, 0.048)
Silk	$992 \times 1013 \times 46$	3	4	12	0.4	-6	5000	0.01	(0.699, 0.030, 0.080)
Velvet	$992 \times 1012 \times 311$	3	4	12	0.4	-1	500	0.1	(0.555, 0.040, 0.074)
Felt	$992 \times 1012 \times 485$	1	2	16	0.4	-30	500	0.125	(0.518, 0.915, 0.365)

Table 1: Fiber filter and scattering model parameter values for our material samples: s and t are the shape parameters, and h is the volume size of the filter (Section 5); ϵ_d and ϵ_J are the noise thresholds. The optical parameters include d , the density multiplier, and the parameters found by our appearance matching algorithm: γ , the standard deviation of the flake distribution, and α , the single-scattering albedo.

where we have assumed without loss of generality that $D(\omega) = D(-\omega)$. Equation 3 indicates a fundamental problem of this approach, namely that w_1 can become large when $\sigma_t(\omega) \approx 0$. This becomes a fundamental limitation as specularity increases, particularly when there are many long transport paths (weights are multiplicative along paths and can build up).

To deal with this problem, we use rejection sampling to derive a sampling strategy that generates samples exactly according to f_p . Since $a\rho/\sigma_t(\omega)$ is an upper bound on f_p/f_1 , the following algorithm produces samples of the desired density:

```

SAMPLE- $f_p(\omega)$ 
1  repeat
2     $m = \text{SAMPLE-FLAKE-DISTRIBUTION}()$ 
3     $\xi = \text{UNIFORM-RANDOM}(0, 1)$ 
4    if  $\xi < |\omega \cdot m|$ 
5      return  $2(\omega \cdot m)m - \omega$ 

```

The added cost of rejection sampling is very small — in our example scenes, only 2-3 iterations were required on average. The approach described in this section is general and might also be useful to improve sampling techniques that are traditionally used for microfacet reflectance models.

7.2 Sampling the flake distribution

The algorithm above assumes the availability of a routine SAMPLE-FLAKE-DISTRIBUTION that can draw samples distributed according to $D(\omega)$. We apply the inversion method in spherical coordinates to the distribution described in Section 4, and find that the latitude (θ) component integrates to:

$$F(\theta) := \frac{1}{2} \left(1 - \operatorname{erf} \left(\frac{\cos \theta}{\sqrt{2}\gamma} \right) \right) / \operatorname{erf} \left(\frac{1}{\sqrt{2}\gamma} \right)$$

where, $F(0) = 0$, $F(\pi) = 1$, and we have temporarily assumed that $\omega_f = (0, 0, 1)^T$. To sample θ , we find $F^{-1}(\xi_1)$ numerically using Brent’s method, where ξ_1 is uniformly distributed on $[0, 1]$. About 10-18 iterations are required to arrive at machine precision. For the longitude (φ) component, we set $\varphi = 2\pi\xi_2$ (where ξ_2 is another uniform variate). To handle general fiber directions, we rely on the same sampling code and simply apply the appropriate rotations to the incident and outgoing directions.

8 Results

Our results are based on samples of silk satin, velvet, felt, and wool gabardine, which were sent to the High-Resolution X-ray Computed Tomography Facility at The University of Texas at Austin. All fabrics were scanned in an XRadia MicroXCT scanner using 1024^3 volumes with a $5 \mu\text{m}$ voxel size, which observed circular areas of approximately 5 mm diameter.

As necessary, our initial data cleanup included corrections to equalize density and contrast between the center and edge of the volume

(vignetting). Further, we straightened the slightly non-planar cloth samples using geometric warping, by fitting a second-order polynomial $p(x, y)$ to points distributed proportional to the CT densities and then resampling the whole volume using the mapping $f(x, y, z) = (x, y, z - p(x, y))$. We then ran the CT image processing pipeline (see Section 5), with the parameters reported in Table 1. Depending on the thickness of the sample, processing took between 1 and 8 hours on a QSSC-S4R Intel Server with 4 Xeon X7560 8-core processors and 32 GB of memory.

Our rendering implementation is based on the open source rendering system Mitsuba [Jakob 2010], which was extended to handle the new micro-flake distribution (Section 4). The rendering itself was done on the Amazon Elastic Compute Cloud (EC2), where we used between 8 and 32 `c1.xlarge` instances (each having 8 cores and 7GB of memory) to jointly render the individual images at a resolution of 2.6 megapixels. With 32 instances, rendering times range between 3.7 (velvet) and 7.4 (satin) hours per image. The supplementary material has 10M-pixel images and the video.

Figure 7 shows results obtained by our appearance matching scheme; the left two columns, the *appearance matching pair*, show the image pairs used for appearance matching (the blue rectangle is used for matching), and the right two columns, the *validation pair*, show a different image pair, also with known and matched lighting, to test how well the model generalizes to other configurations. The sizes of all samples we used for appearance matching are roughly 10×10 cm.

Figure 8 shows the resulting models shell-mapped onto draped fabric geometry and rendered under environment lighting.

The silk satin (charmeuse) has a structure of mainly parallel fibers on the surface, resulting in a strong anisotropic highlight. In Figure 7-(1), the appearance matching pair uses a cylindrically curved piece of material, and the matching region was chosen to include a highlight to allow the matching process to tune γ appropriately. Good results are obtained despite the mismatch between the ideal cylinder in the rendering and the flatter shape of the real material, illustrating that a casual setup suffices. Using the parameters obtained from this view, the validation pair shows the fabric rotated 90 degrees and draped over the same cylinder. At this angle the fabric exhibits almost no highlight; this anisotropic appearance is correctly predicted by our model.

The satin is shown in a draped configuration in Figure 8-(a) and in the accompanying video and supplementary images. No reflectance model, BRDF, BTF, or other multi-view image data is used for these renderings—the orientation information in the volume automatically causes the characteristic appearance of this fabric to emerge when the model is rendered.

For gabardine, a wool twill fabric, the variation in texture with illumination direction is an important appearance characteristic. In Figure 7-(2), the appearance matching pair is lit with a low-frequency environment map. The validation pair accurately predicts the texture under a different lighting condition, which involves a strong luminaire at the top. In the draped configuration in Figure 8-




Figure 8: Fabrics in draped configurations with our volumetric appearance model (a) silk satin, (b) gabardine, (c) velvet, and (d) felt.

(b), the volume model captures subtle foreshortening effects and the silhouette appearance, as well as the subtle variations in texture across the surface. The appearance at the cut edge gives the proper impression of the thickness of the fabric (compare to the very thin satin material), which is a perennial difficulty with surface models.

Velvet, a material with a cut pile (like a carpet), has a visible surface composed of fibers that stick up from the base material. It has a very distinctive appearance, with a characteristic grazing-angle highlight. Appearance matching for velvet (see Figure 7-(3)) was done using a curved configuration and the same harsh lighting as used for gabardine’s validation, producing distinct highlights on both sides of the cylinder. The validation pair shows a different, softer lighting, which results in a less distinct highlight; our model agrees qualitatively with the photograph. The appearance of velvet depends on how the fibers are brushed, and our random tile rotation method produces randomly brushed velvet. In Figure 8-(c) and in the video, we demonstrate how our model reproduces the characteristic velvet highlights. Further, the edges and silhouettes convey the considerable thickness and weight of this material.

Felt is a nonwoven textile consisting of a disorganized layer of matted fibers. The thickness and fuzziness of this material are important appearance attributes that are generally difficult to model and render. Since felt does not exhibit an overall specular highlight, we used a flat patch for appearance matching; because of limited depth of field we limited the matching region to a thin rectangle where the photograph is in good focus. The illumination conditions for the appearance matching and the validation are the same as those for the gabardine. The color and the contrast due to self-shadowing attributes are matched nicely and generalize well to the second illumination condition. One limitation for this material is that it has substantial low-frequency content in its texture, which our small

sample area did not capture in the CT imaging, leading to a slightly more uniform appearance in our tiled material. Figure 8-(d) demonstrates the ability of our volumetric appearance model to capture the material’s thick, fuzzy appearance.

A 3D, physically based model also allows more meaningful editing than image-based methods. Figure 9 shows renderings created after performing simple edits to the underlying volume representation. In the top row, we reflect lookups into the satin volume data across the central plane of the fabric, conditioned on a binary texture map that covers the surface. This edit reveals the back face of the satin weave, which exhibits softer reflections due to less coherent fiber directions, much as these weaves are interchanged in jacquard-woven satin. Two different lighting conditions are shown.

In the bottom row of Figure 9, we extend the gabardine model with a spatially varying albedo value. The albedo is computed as a function of orientation, so that fibers in the warp and weft are assigned different colors. With blue warp and white weft a fabric similar to denim is produced, though made of wool rather than cotton.

Finally, we compare our method to the surface-based BRDF and texture model introduced by Irawan [2008] (Figure 10). For these two examples, Irawan fit his model to BRDF measurements of exactly the same materials we measured. The renderings for Irawan’s and our methods took roughly 8 and 64 core hours, respectively.

At the large scale, the BRDFs of the fabrics match reasonably well. Irawan showed that his model matches the measured BRDFs of these materials to a similar degree of fidelity, so this confirms that our method predicts large-scale reflectance from the structure and a single image. For yarn-scale texture, the two models produce generally similar results, though Irawan’s model is lower in contrast for the gabardine because it does not account for shadowing. It also




Figure 9: Renderings obtained by editing the volumetric representation; Top row: the material is flipped using a binary texture map (two lighting conditions are shown). Bottom row: the gabardine sample is rendered with a blue hue (left); we then detect weft fibers based on their orientation and color them white, which produces a material resembling denim (right).

produces a more uniform appearance. At the small scale, as seen in the insets, and at silhouettes and edges, our detailed volumetric model produces dramatically more realistic results.

9 Conclusions

We have demonstrated a new, multimodal approach to making realistic volume models of cloth that capture both the 3D structure evident in close-up renderings and the BRDF evident in farther-away views. Unlike previous methods for capturing cloth appearance using BTFs, our method explicitly models the 3D structure of the material and, interestingly, is able to capture the directional reflectance of the material automatically because of this structure.

Our modeling approach uses CT imaging where it is strongest, in measuring 3D structure, and it uses photographs where they are strongest, in measuring color and texture. By matching texture statistics we merge these two sources of information, resulting in a volume model that can produce both close-up views with rich detail of fuzz and fiber structure *and* the characteristic BRDFs (highlights) of these materials that emerge naturally from rendering the measured structure. No BRDF measurements are made, and only a few parameters are adjusted in the optical model. The appearance of the cloth is created by a simple anisotropic phase function model together with the occlusion and orientation information extracted from the volume. This paper shows that since geometric structure is what creates the complex appearance of textiles, once we acquire the structure, we are most of the way to modeling the appearance.

Aside from its implications regarding how material appearance can be modeled from structure, this is also quite a practical method for appearance modeling. All that is required to model a material is a CT scan, which can be obtained at reasonable cost from a number of facilities (or in the future from the rapidly improving technology of desktop CT scanning) and a few photographs under known illumination, which takes only a few minutes with a camera and a mirror sphere. The resulting models are volumetric in nature, and physically based, which makes them easier to edit than image-based data. It is easy to adjust color, glossiness, opacity, and material thick-




Figure 10: Silk satin (1) and gabardine (2) rendered with Irawan’s surface-based representation (a) and with our model (b).

ness by scaling parameters of the volume geometry; and a range of more fundamental changes to the material’s structure can be made by editing the volume data.

This paper has demonstrated the usefulness of the CT modeling approach for textiles, but the approach does have some limitations. Particularly, it requires that changes in optical properties correlate with changes in density, and this requirement could limit the kinds of materials that can be captured using this imaging modality. Further, the scanner can only image small samples, less than a centimeter across, at the resolution needed to produce clear fiber orientation maps. Thick materials that do not fit fully in the volume (e.g., materials with very long flyaway fibers) cannot be handled well. Some unusual materials, such as metallic fibers, may be problematic for CT because of limited dynamic range. Also, texture content

at larger scales will be missed. These problems will decrease as CT scanners improve in resolution and dynamic range. CT is very well suited to textiles, and it remains to be seen what other materials it performs well for, and how other volume imaging methods work in this technique. Further, materials with differently colored yarns cannot be currently captured by our method.

There are many areas of future work. For better models, volume texture synthesis methods are needed that can work with this type of structured geometry and produce large volumes of high quality seamless texture while maintaining rendering efficiency. This work was done using extremely small samples, and with larger samples, which should be possible as CT technology improves and becomes cheaper, better texture could be produced. To extend the range of materials that can be handled, new parameter estimation methods are needed that can identify and fit multiple materials within a single volume. To improve accuracy, more photographs under varying conditions can be used, allowing more parameters (for instance, more complex phase functions) to be fit. Ultimately, this method can be extended to work for a wide range of types of materials whose appearance is difficult to capture using surface models.

Acknowledgements

The authors would like to thank Jessie Maisano at the UTCT scanning facility, and Piti Irawan for providing the gabardine, satin, and velvet samples and his cloth model code. Funding was provided by the National Science Foundation under awards CCF-0644175, CCF-0702490, CCF-0811680 and IIS-1011919, an Intel PhD Fellowship, and Autodesk. We thank Intel Corporation and Amazon for their donations of computing equipment and computer time.

References

- ADABALA, N., MAGNENAT-THALMANN, N., AND FEI, G. 2003. Visualization of woven cloth. In *14th Eurographics Workshop on Rendering*, 180–185.
- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based brdf generator. In *Proceedings of ACM SIGGRAPH 2000*, 65–74.
- AXELSSON, M. 2008. Estimating 3d fibre orientation in volume images. In *International Conference on Pattern Recognition, 2008*, 1–4.
- CHEN, Y., AN HUA ZHONG, S. L., XU, Y.-Q., GUO, B., AND SHUM, H.-Y. 2003. Realistic rendering and animation of knitwear. *IEEE Transactions on Visualization and Computer Graphics* 9, 1, 43–55.
- DANA, K., VAN GINNEKEN, B., NAYAR, S., AND KOENDERINK, J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics* 18, 1, 1–34.
- FURUKAWA, R., KAWASAKI, H., IKEUCHI, K., AND SAKAUCHI, M. 2002. Appearance based object modeling using texture database: acquisition, compression and rendering. In *Eurographics Workshop on Rendering*, 257–266.
- GONG, R., OZGEN, B., AND SOLEIMANI, M. 2009. Modeling of yarn cross-section in plain woven fabric. *Textile Research Journal* 79, 11, 1014–1020.
- HEEGER, D. J., AND BERGEN, J. R. 1995. Pyramid-based texture analysis/synthesis. *SIGGRAPH '95*, 229–238.
- IRAWAN, P. 2008. *Appearance of woven cloth*. PhD thesis, Cornell University, Ithaca, NY, USA. AAI3295837.
- JAKOB, W., ARBREE, A., MOON, J. T., BALA, K., AND MARSCHNER, S. 2010. A radiative transfer framework for rendering materials with anisotropic structure. In *SIGGRAPH 2010*, ACM, New York, NY, USA, 53:1–53:13.
- JAKOB, W., 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with three dimensional textures. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, 271–280.
- KAWABATA, S., NIWA, M., AND KAWAI, H. 1973. The finite deformation theory of plain weave fabrics. part i: The biaxial deformation theory. *Journal of Textile Institute* 64, 1, 21–46.
- LOMOV, S., PARNAS, R., GHOSH, S. B., VERPOEST, I., AND NAKAI, A. 2002. Experimental and theoretical characterization of the geometry of two-dimensional braided fabrics. *Textile Research Journal* 72, 8, 706–712.
- LU, R., KOENDERINK, J. J., AND KAPPERS, A. M. L. 1998. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics* 37, 25, 5974–5984.
- MAGDA, S., AND KRIEGMAN, D. 2006. Reconstruction of volumetric surface textures for real-time rendering. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)*, 19–29.
- MOTOYOSHI, I., NISHIDA, S., SHARAN, L., AND ADELSON, E. 2007. Image statistics and the perception of surface qualities. *Nature*, 206–209.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, 117–126.
- PERLIN, K., AND HOFFERT, E. M. 1989. Hypertexture. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, 253–262.
- PIERCE, F. T. 1937. The geometry of cloth structure. *Journal of the Textile Institute* 28, 3, 45–96.
- PORUMBESCU, S., BUDGE, B., FENG, L., AND JOY, K. 2005. Shell maps. *ACM Transactions on Graphics* 24, 3, 626–633.
- 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.
- THIBAUT, X., AND BLOCH, J. 2002. Structural analysis by X-ray microtomography of a strained nonwoven papermaker felt. *Textile Research Journal* 72, 6, 480–485.
- WALTER, B., MARSCHNER, S., LI, H., AND TORRANCE, K. 2007. Microfacet models for refraction through rough surfaces. In *Eurographics Symposium on Rendering*, 195–206.
- WANG, H., WU, Q., SHI, L., YU, Y., AND AHUJA, N. 2005. Out-of-core tensor approximation of multi-dimensional matrices of visual data. *ACM Transactions on Graphics (ACM SIGGRAPH 2005)* 24, 3, 527–535.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, 255–264.
- XU, Y.-Q., CHEN, Y., LIN, S., ZHONG, H., WU, E., GUO, B., AND SHUM, H.-Y. 2001. Photorealistic rendering of knitwear using the Lumislice. In *Proceedings of ACM SIGGRAPH 2001*, 391–398.