Building Volumetric Appearance Models of Fabric Using Micro CT Imaging

By Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala

Abstract

Cloth is essential to our everyday lives; consequently, visualizing and rendering cloth has been an important area of research in graphics for decades. One important aspect contributing to the rich appearance of cloth is its complex 3D structure. Volumetric algorithms that model this 3D structure can correctly simulate the interaction of light with cloth to produce highly realistic images of cloth. But creating volumetric models of cloth is difficult: writing specialized procedures for each type of material is onerous, and requires significant programmer effort and intuition. Further, the resulting models look unrealistically "perfect" because they lack visually important features like naturally occurring 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, yielding a scalar density volume. This 3D data has micron resolution details about the structure of cloth but lacks all optical information. So we combine this density data with a reference photograph of the cloth sample to infer its optical properties. We show that 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 such as satin and velvet emerge automatically—all based simply on having accurate mesoscale geometry.

1. INTRODUCTION

Cloth is a fundamental material in our day-to-day lives, and creating photo-realistic renderings of cloth has been an active research topic in computer graphics for decades, with applications in areas including virtual prototyping, entertainment (movies and games), and retail.

One important aspect contributing to the appearance of cloth is its complex 3D structure, yielding complicated textures and reflectance. Further, the structure is irregular, causing difficult-to-model, but visually important randomness. Volume rendering techniques, which model such structure correctly and simulate the interaction of light with cloth explicitly, have been explored since the 1990s. These approaches address the limitations of usual surface-based models, which are visually unsatisfactory because they treat cloth as infinitely thin sheets.[8, 12, 18] Cloth exhibits a wide range of appearance, but shares a common basic structure of long, shiny fibers. The thick, fuzzy nature of cloth makes

volume models a good fit. Further, recent developments[7] 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. Previous work has primarily relied on procedural methods (special-purpose programs) for modeling these volumes, but this approach has limited generality: significant creative effort is needed to design these programs for each new material. Further, the resulting models look unrealistically "perfect" since they lack subtle irregularities that appear in real cloth.

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, and then fill in optical properties using a reference photograph. We do the latter by solving an inverse problem that statistically matches photographs and physically based renderings.

Many volume imaging technologies have been developed, including computed tomography (CT), magnetic resonance, and ultrasound, 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 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 such as 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; this volumetric appearance model produces 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 such as Bidirectional Texture Functions (BTF), which essentially consist of an exhaustive set of photographs of the surface under all possible illumination and viewing directions.[4, 5] 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 enough resolution for capturing high glossiness, and generally fails to capture or predict grazing angles, making silhouettes and edges unrealistic.

Two prominent early volume appearance models are Kajiya and Kay's[8] fur rendering and Perlin and Hoffert's[12] "hypertexture." 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[3, 18] which focused on modeling and rendering knitwear. Magda and Kriegman[11] describe a method for acquiring *volumetric textures* that combine a volumetric normal field, local reflectance functions, and occupancy information. All these approaches need significant modeling effort. Recently, Jakob et al.[7] 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 complicated optical behavior. Westin et al.[17] modeled cloth's reflectance profile by ray-tracing mesostructure models, which is related to the way cloth highlights emerge in our system. Ashikhmin et al.[2] rendered velvet and satin using hand-designed microfacet distributions. Adabala et al.[1] proposed a rendering method for woven cloth based on microfacet models, and Irawan and Marschner[6] presented an elaborate model, based on the analysis of fiber tangent directions in a range of woven fabrics, and validated it against reflectance 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.

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 close-ups, 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.[9, 13] More recently, X-ray tomography, using synchrotron facilities[16] or the rapidly improving micro-CT scanners,[10, 15] has been used to examine the structure of textiles in several applications. These studies focus on extracting geometric information related to a material's mechanical properties, but have also produced some analysis tools[15] 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 (Figure 1).

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

(a)

(b)

(c)
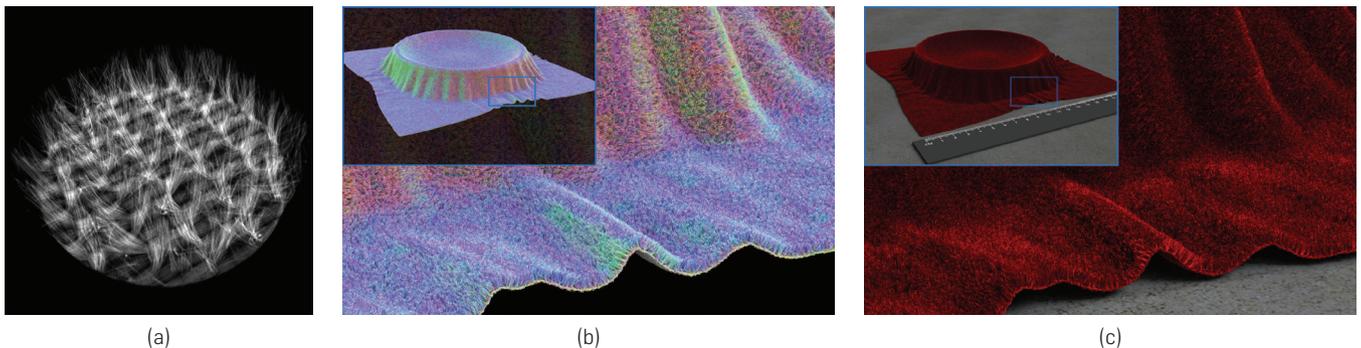
Because cloth is made of fibers, we need a volume scattering model that can handle the anisotropy of fibers; we chose a modified version of the model proposed by Jakob et al.[7] (detailed in Section 4) for this purpose. This model requires an optical density, an albedo, and two phase function parameters: an orientation vector and a specular lobe width. Intuitively, the optical density describes how often light scatters within the cloth; the albedo and the phase function respectively capture the fraction of light being absorbed and how light changes its direction at each scattering location.

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[14] to warp it.

The resulting renderings (Section 7) 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 such as 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.[7] which states that within participating media,

$$(\omega \cdot \nabla)L(\omega) + \sigma_t(\omega)L(\omega) = \sigma_s(\omega)\int_{S^2} f_p(\omega' \to \omega)L(\omega')\,d\omega'$$

where σ_s and $\sigma_t : S^2 \to \mathbb{R}$ are the anisotropic scattering and extinction coefficients, and f_p is the phase function. 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 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* $D(\mathbf{m})$ that describes the orientation $\mathbf{m}$ of (unresolved) idealized mirror flakes at every point in space. Similar to microfacet models, the phase function then involves evaluating $D(\mathbf{m})$ at the half-way direction between the incident and outgoing direction. For completeness, we reproduce the model's definition as follows:

$$\sigma_t(\omega) = a\rho\int_{S^2}|\omega \cdot \mathbf{m}|D(\mathbf{m})\,d\mathbf{m},$$
$$\sigma_s(\omega) = \alpha\,\sigma_t(\omega),$$
$$f_p(\omega' \to \omega) = \frac{a\rho\alpha}{4\sigma_s(\omega)}\Big(D\big(h(\omega,-\omega')\big) + D\big(-h(\omega,-\omega')\big)\Big)$$

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.

(a) Micro CT images

(b) Reconstructed density field and orientation field

(c) Appearance matching

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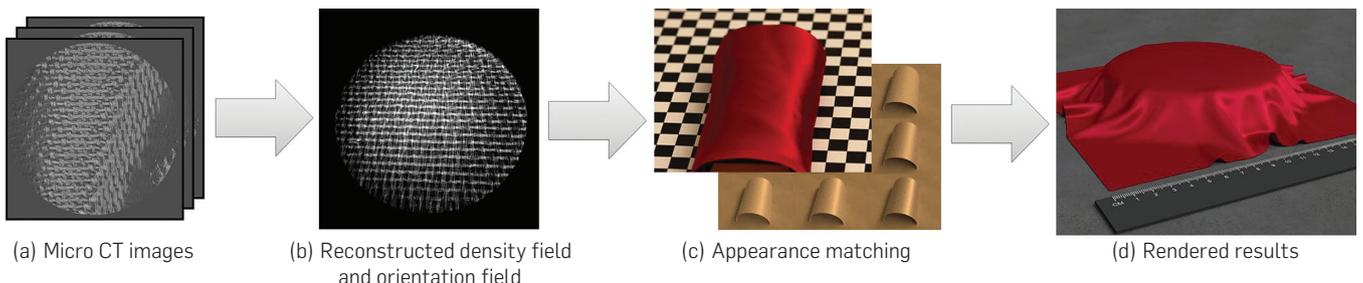
Here, ρ denotes the particle density, a is the area of a single flake, α is the associated albedo, and $h(\omega, \omega') := (\omega + \omega')/,\omega + \omega',$. Note that the above expressions are simplified by assuming the flakes have albedo independent of the scattering angle. This reduces our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

4.1. Flake distribution

We propose a flake distribution that is convenient to integrate while capturing the same key feature as the one proposed by Jakob et al.[7] We use the following density function, which specifies a truncated Gaussian centered around the great circle perpendicular to the local fiber orientation $\bar{\omega}$:

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

where the standard deviation γ determines the roughness of the fiber. More precisely, the parameters required to create renderings are:

- $\bar{\omega}$, 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, that is, $a\rho(\mathbf{x}) := d \cdot \mathrm{CT}(\mathbf{x})$, where d is a constant of proportionality.

Section 5 discusses the steps needed to obtain $\mathrm{CT}(\mathbf{x})$ and $\bar{\omega}(\mathbf{x})$; in section 6, we describe how to find α, γ, and d.

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 nonempty voxel. Our approach uses oriented filters to detect fibers, based on similar filters used by Shinohara et al.,[15] to locate fibers in CT data.

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

with the axis $\mathbf{d}$, consisting of a difference of Gaussians in distance from the axis: $q(\mathbf{d}; \mathbf{p}) := -2 \exp(-ur^2) + \exp(-wr^2)$, where $r = ,\mathbf{p} - (\mathbf{p} \cdot \mathbf{d})\mathbf{d},$ is the distance from the filter's axis and the parameters u and w (normally $u < w$) are empirically adjusted based on the size of the fibers present in the sample (see Figure 3).

The raw CT volume is thresholded at a value ϵ_d, resulting in a binary volume b: for any $\mathbf{x}$, set $b(\mathbf{x})$ to 0 if $\mathrm{CT}_{\mathrm{raw}}(\mathbf{x}) \geq \epsilon_d$ and to 1 otherwise. Then b is convolved with the filter q for each of a fixed set of orientations: $J(\mathbf{x}, \mathbf{d}) := \sum_{\mathbf{p} \in V} q(\mathbf{d}; \mathbf{p}) b(\mathbf{x} + \mathbf{p})$, where V is a cubic volume of edge length h.

As shown in Figure 3, 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 $\bar{\omega}(\mathbf{x}) = \mathbf{d}'$. In our implementation, we precompute q on a set of directions $\{\mathbf{d}_i\}$ picked from a $32 \times 32 \times 6$ cubemap. Then for each nonempty voxel $\mathbf{x}$, we set $\bar{\omega}(\mathbf{x}) = \mathbf{d}_j$, 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 b. 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

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

5.3. Data replication

The volume data needs to be replicated for rendering since our samples are very small. We explored example-based synthesis in Zhao et al.[19] which provides sophisticated tools to do this, but it is orthogonal to this paper. Here we consider 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

Figure 3. Computing function J in 2D: (a) shape of the filter q; (b) when q is aligned to the fiber; (c) when q is unaligned.

Filter $q(\mathbf{d}; \cdot)$ J has a high value J has a low value
(a)   (b)   (c)

$q < 0$ $q \geq 0$ $b = 0$ (fiber) $b = 1$ (background)
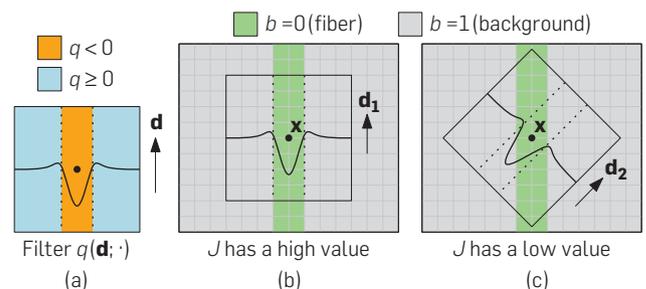
surface is simply covered by a rectangular array of tiles copied from the volume, without continuity at the tile boundaries.

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, such as 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 subrectangle 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.[14]

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 4 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 nontrivial; 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 is to repeatedly render the

Figure 4. (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.
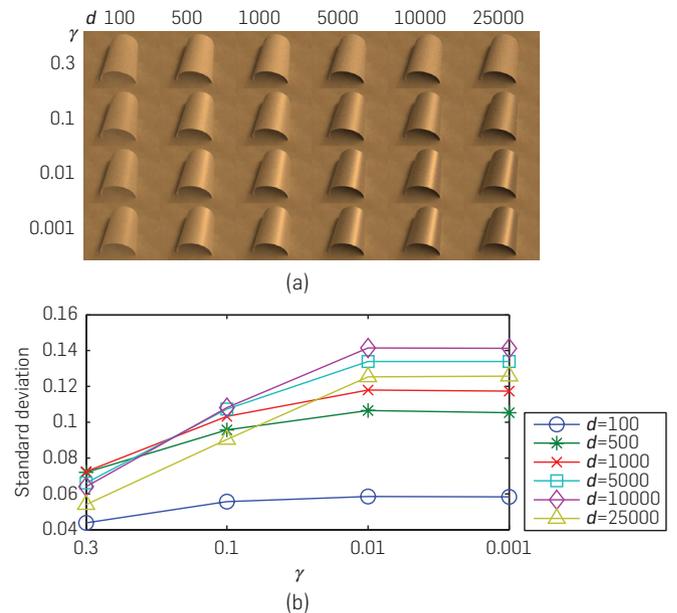
volume using our physically based renderer and to adjust the optical parameters to match certain texture statistics of the rendered images to statistics of the photograph.

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 other measures, 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 4, 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 predetermine 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, a single scalar value) using an iterative algorithm. Note that the mean and standard deviation of pixel values change monotonically with changes in α and γ, respectively. Thus, a binary search can be used to significantly improve performance as follows: first, an initial guess of γ is assumed, and 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 two or three iterations.

Finally, we take another photo under a different setup and render a corresponding image as a qualitative validation (see Section 7). Figure 5 shows the appearance matching results for two different materials.

7. 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 $5\,\mu m$ voxel size, which observed circular areas of approximately $5\,mm$ diameter. Our rendering implementation is based on the open source rendering system Mitsuba, which was extended to handle the new micro-flake distribution (Section 4). All the rendered results are generated using Monte Carlo volume path tracing.

Figure 6 shows the resulting models shell-mapped onto draped fabric geometry and rendered under environment lighting. The corresponding high-resolution renderings are available on the project Web site at http://www.cs.cornell. edu/projects/ctcloth-sg11. Fabrics are usually rendered with surface-based models, making them thin 2D sheets so that cut edges look artificial; our volumetric model, on the other hand, explicitly captures 3D structures and provides a proper impression of the thickness and weight of the fabrics. Furthermore, surface-based models show smooth edges at silhouettes while our model is able to produce fuzzy silhouettes with rich details, bringing a new level of realism to fabric rendering.

The silk satin (charmeuse) has a structure of mainly parallel fibers on the surface, resulting in a strong anisotropic highlight. In Figure 5(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 6(a). No reflectance model such as 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 5(2), the appearance matching pair is lit

Table 1. Fiber scattering model parameter values for our material samples: *d*, the density multiplier; γ, the standard deviation of the flake distribution; α, the single scattering albedo

Material	Data size	*d*	γ	α
Gabardine	$992 \times 1012 \times 181$	5000	0.1	(0.892, 0.063, 0.048)
Silk	$992 \times 1013 \times 46$	5000	0.01	(0.699, 0.030, 0.080)
Velvet	$992 \times 1012 \times 311$	500	0.1	(0.555, 0.040, 0.074)
Felt	$992 \times 1012 \times 485$	500	0.125	(0.518, 0.915, 0.365)

Figure 5. Appearance matching results: (top) silk, (bottom) gabardine. 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.
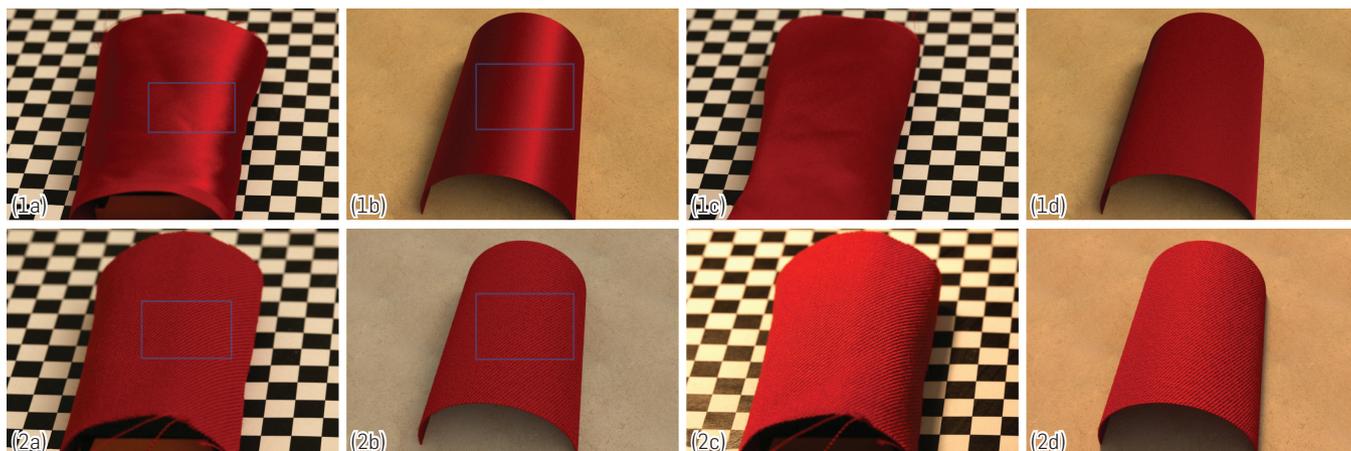
Figure 6. Fabrics in draped configurations with our volumetric appearance model: (a) silk satin, (b) gabardine, (c) velvet, and (d) felt.
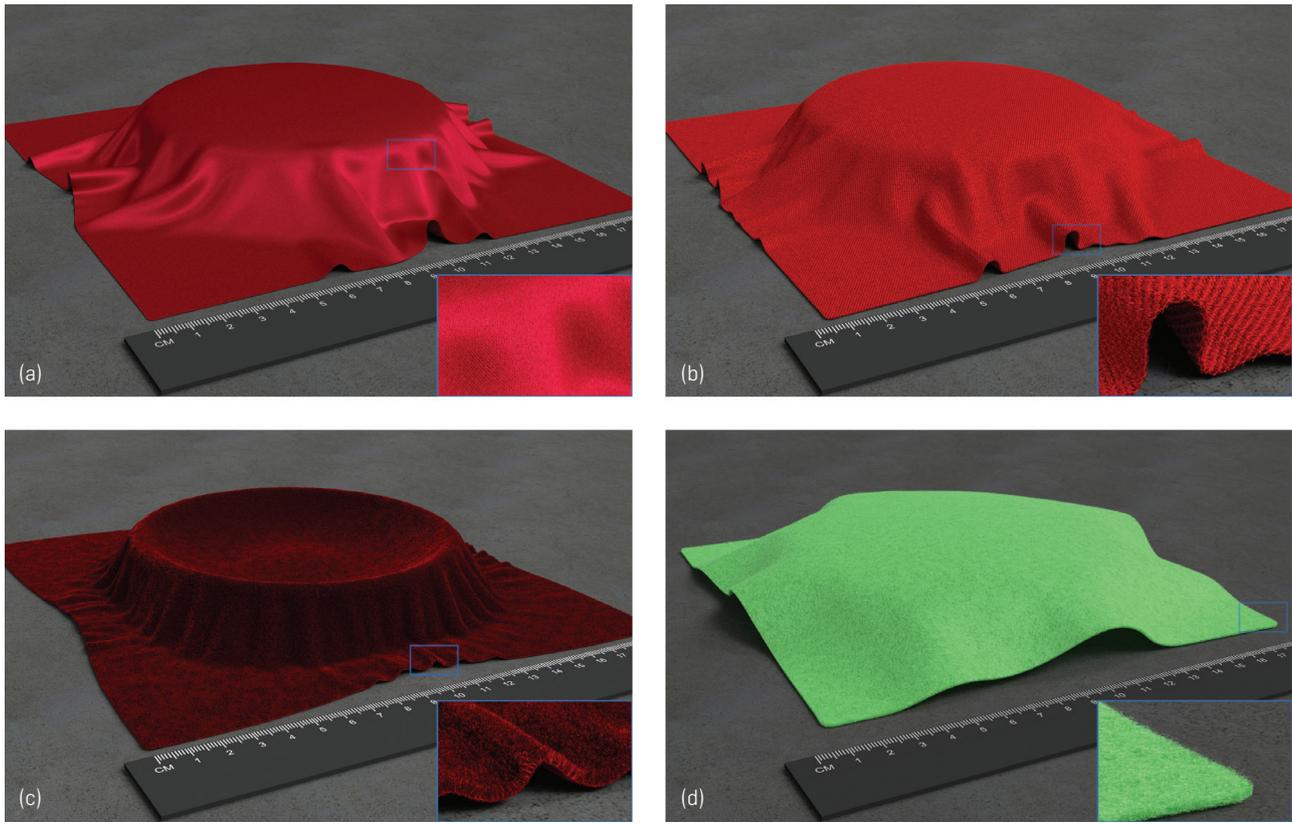
with a low-frequency environment map. The validation pair accurately predicts the texture under a different lighting condition, which involves a strong luminaire at the top. In the draped configuration in Figure 6(b), the volume model captures subtle foreshortening effects and the silhouette appearance, 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. The appearance of velvet depends on how the fibers are brushed, and our random tile rotation method produces randomly brushed velvet. In Figure 6(c), 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 6(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. In Figure 7, 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.

8. CONCLUSION

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 reflectance 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

Figure 7. Renderings obtained by editing the volumetric representation: 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).
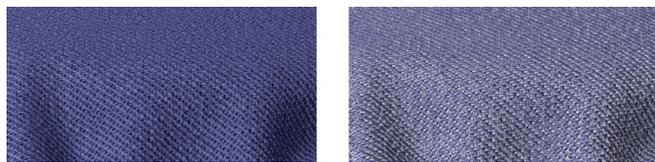
texture. By matching texture statistics we merge these two sources of information, resulting in a volume model that can produce both close-up views with rich detail of fuzz and fiber structure *and* the characteristic highlights of these materials that emerge naturally from rendering the measured structure. No reflectance 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. In addition, it is possible to CT scan a few samples with elementary weave structures and assemble the resulting volumes to form fabrics with many different designs,[19] further reducing per-design cost. 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 thickness 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. 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 (e.g., 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. **C**

References
1. Adabala, N., Magnenat-Thalmann, N., Fei, G. Visualization of woven cloth. In *14th Eurographics Workshop on Rendering* (2003), 180–185.
2. Ashikhmin, M., Premoze, S., Shirley, P.S. A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000* (2000), 65–74.
3. Chen, Y., an Hua Zhong, S.L., Xu, Y.Q., Guo, B., Shum, H.Y. Realistic rendering and animation of knitwear. *IEEE Trans. Visual. Comput. Graph. 9*, 1 (2003), 43–55.
4. Dana, K.J., van Ginneken, B., Nayar, S.K., Koenderink, J.J. Reflectance and texture of real-world surfaces. *ACM Trans. Graph. 18*, 1 (1999), 1–34.
5. Furukawa, R., Kawasaki, H., Ikeuchi, K., Sakauchi, M. Appearance based object modeling using texture database: Acquisition, compression and rendering. In *Eurographics Workshop on Rendering* (2002), 257–266.
6. Irawan, P., Marschner, S. Specular reflection from woven cloth. *ACM Trans. Graph. 31*, 1 (2012), 11:1–11:20.
7. Jakob, W., Arbree, A., Moon, J.T., Bala, K., Marschner, S. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Trans. Graph. 29*, 4 (2010), 53:1–53:13.
8. Kajiya, J.T., Kay, T.L. Rendering fur with three dimensional textures. *SIGGRAPH Comput. Graph. 23*, 3 (1989), 271–280.
9. Kawabata, S., Niwa, M., Kawai, H. The finite deformation theory of plain weave fabrics. Part I: The biaxial deformation theory. *J. Textile Instit. 64*, 1 (1973), 21–46.
10. Lomov, S., Parnas, R., Ghosh, S.B., Verpoest, I., Nakai, A. Experimental and theoretical characterization of the geometry of two-dimensional braided fabrics. *Textile Res. J. 72*, 8 (2002), 706–712.
11. Magda, S., Kriegman, D. Reconstruction of volumetric surface textures for real-time rendering. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2006), 19–29.
12. Perlin, K., Hoffert, E.M. Hypertexture. *SIGGRAPH Comput. Graph. 23*, 3 (1989), 253–262.
13. Pierce, F.T. The geometry of cloth structure. *J. Textile Instit. 28*, 3 (1937), 45–96.
14. Porumbescu, S., Budge, B., Feng, L., Joy, K. Shell maps. *ACM Trans. Graph. 24*, 3 (2005), 626–633.
15. Shinohara, T., Takayama, J., Ohyama, S., Kobayashi, A. 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 Res. J. 80*, 7 (2010), 623–630.
16. Thibault, X., Bloch, J. Structural analysis by X-ray microtomography of a strained nonwoven papermaker felt. *Textile Res. J. 80*, 7 (2002), 480–485.
17. Westin, S.H., Arvo, J.R., Torrance, K.E. Predicting reflectance functions from complex surfaces. *SIGGRAPH Comput. Graph. 26*, 2 (1992), 255–264.
18. Xu, Y.Q., Chen, Y., Lin, S., Zhong, H., Wu, E., Guo, B., Shum, H.Y. Photorealistic rendering of knitwear using the Lumislice. In *Proceedings of ACM SIGGRAPH 2001* (2001), 391–398.
19. Zhao, S., Jakob, W., Marschner, S., Bala, K. Structure-aware synthesis for predictive woven fabric appearance. *ACM Trans. Graph. 31*, 4 (2012), 75:1–75:10.

Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala ([szhao, wenzel, srm, kb]@cs.cornell.edu), Cornell University, Ithaca, NY.