An Interactive Appearance Model for Microscopic Fiber Surfaces

Z. Velinov M. B. Hullin
University of Bonn and X-Rite Graduate School on Digital Material Appearance*

Abstract

Modeling and rendering the appearance of fabrics remains one of the most challenging research topics in computer graphics. Today's most advanced models are based on volumetric fiber distributions, obtained from advanced measurement techniques like micro-CT, and only renderable through expensive volume rendering approaches. In this paper, we propose an analytical BRDF model for pile fabrics, i.e., textiles like velvet, plush or Alcantara, that are characterized by open-ended strands of yarn. A fascinating property of many such materials is that they "memorize" tactile interaction and that their appearance depends on the last direction in which the surface was brushed. Our reflectance model, inspired by recent microflake models, links the microscopic structure of a material to the interaction that caused that structure, and to the resulting macroscopic reflectance behavior. We present an end-to-end pipeline for fitting the parameters of our model to measured reflectance data, for manipulating fiber distributions using tactile input, and rendering the resulting spatially varying surface appearance in real time.

1 Introduction

Textile materials are a central ingredient in industries as diverse as clothing, furniture or automotive. Capturing and faithfully recreating the characteristic visual appearance of fabrics is therefore an important goal and an active field of research in computer graphics. So far, most research has focused on woven or knitted fabrics whose structure is tightly held together and hence static. This property enables relatively straightforward appearance acquisition by capturing images of a material sample under different lighting and viewing conditions.

In contrast, pile fabrics (velvet, alcantara and plush) contain loose fibers which causes their appearance to depend on the history of interaction. Such materials are visually intriguing, because their tilted fibers can exhibit complex specular reflectance that is not aligned with the macroscopic cloth surface. More familiar and fascinating, however, is the experience of being able to "paint" on a carpet, pillow or car seat using one's fingers. The goal of this work is to recreate this experience, which implies the need for three major ingredients: firstly, a reflectance model capable of recreating the interaction-dependent appearance of pile fabrics; secondly, a method to derive the model parameters from physical samples; and finally a rendering algorithm, along with a user interface, for displaying and manipulating virtual objects in real time.

The need for a real-time solution limits the feasible techniques that can be applied to solve the problem. Spatially varying bidirectional reflectance distribution functions (SVBRDFs) based on microfacet distributions are considered a physically correct way of modeling the appearance of opaque surfaces. They are built under certain assumptions about the surface geometry; most importantly, that the surface is comprised of microscopic mirrors distributed relative to the surface normal.

Like all textile materials, however, pile fabrics are not opaque surfaces. Rather, light that interacts with them is scattered most intricately by a volume filled by microscopic fibers. Consequently, the state of the art in fabric rendering is held by sophisticated rendering schemes that treat the material as a volume of scattering microflakes [JAM*10] or even explicit fibers. However, these techniques are computationally expensive, and modeling the interaction between external forces and every fiber would be unfeasible. Rather, we borrow ideas from these models and use them to derive a novel SVBRDF representation that approximates not only the reflectance caused by microflake distributions, but also the connection between these distributions and the way that the material has been interacted with.

To validate our approach, we measure the appearance of an alcantara sample and fit the proposed model to the obtained data. We use hemispherical bidirectional texture function (BTF) measurements, which also include non-local effects like interreflection, parallax or shadowing. To fit the model to the measured BTF data we perform a principal component analysis (PCA) on the surface distribution of normals (NDF) to derive an initial guess which is then refined using non-linear least squares optimization.

The main contributions of our work are thus threefold:

- We introduce a SVBRDF model for the local appearance of pile fabrics and the way that this appearance is affected by tactile interaction. The shading model is based on a distribution of microscopic fibers that is in general not aligned with the macroscopic material surface.
- We develop a complete pipeline for measuring, fitting and simulating materials by example.
- We demonstrate our model in an application that allows the user to interact with an object made of virtual pile fabric in real time.

The rest of the paper is structured as follows. In Section 2 we discuss the class of pile fabrics and how to approach the problem of modeling them. We discuss related work in Section 3. Then we introduce our pipeline in Section 4 by making a breakdown of each of its elements. We apply our approach to measured data and compare the results in Section 5. Section 6 discusses the strengths and weaknesses of the proposed approach. We conclude the paper with an outline of the achieved results in Section 7.

2 Background

Velvet is woven on a special loom that weaves two thicknesses of the material at the same time. They are cut afterwards to create the characteristic pile effect on each separate piece. Plush is produced in similar manner as velvet. In some weaving patterns the pile fibers

Figure 1: The surface of Alcantara is comprised of microscopic fibers whose orientation is dependent on the history of previous tactile interactions. The distribution of fibers gives its characteristic anisotropic appearance.

might be pulled from the base fabric without being cut. Alcantara is made of fine polyester microfibers. The polyester fibers are cut into small segments and then processed by a carding machine into a soft felt. The cloth-like appearance is achieved by a process of needle punching and treatment with special agents. The most characteristic property of these materials is that they quickly restore their geometry to a stationary state close to the initial after tactile interaction. However, most of the fibers are still slightly bent towards the direction of interaction (See Figure 1). The ability to reach quickly stationary state is highly dependent on the fiber length. Longer fibers as rule of thumb take more time to restore their shape. We limit our study to fabrics with short fibers which allows us to capture the appearance of different brushing patterns without requiring the acquisition of parameters used for dynamics simulation.

To reproduce the characteristic look of these materials we have devised a complete pipeline (See Figure 2). Every stage is covered - from appearance acquisition to material modeling and interactive simulation. We capture BTF of material sample brushed in eight different directions in camera dome. We process it by applying PCA and consequently nonlinear least squares optimization to fit the measured reflectance in microflake distribution adapted for surface rendering. Additional post processing is done to eliminate misalignment between neighboring pixels of the SVBRDF. Then we remap the differently brushed regions to certain states of the fabric. Afterwards we simulate it in a custom real-time ray tracing software running on the GPU. The simulation relies on the geometric information encoded within the microflake distribution to perform accurate interpolation.

Our approach is suitable for material showcases which enable the user to display different objects made of this kind of materials. It enhances the perceptual qualities of the modeled material by providing direct feedback to the user who can interact with the fabric.

3 Related Work

We identify several strands of prior work that are closely related to our problem and approach.

Time-dependent appearance. In our application, the interaction with a material leads to a change in appearance over time, a problem that has already been explored in the context of paint drying and dust accumulation [SSR*07]. In this work, a robotic arm was used to acquire gonioreflectometric data in regular time intervals, which was then used to fit the parameters of analytical BRDF models. The time dependency was then investigated in more detail by analyzing temporal series of measured BRDFs. In contrast, in our work we are not only interested in a time series but the more complex connection between tactile interaction with a material and its appearance. This leads to a higher-dimensional fitting problem. In that sense, our work is related to weathering simulation [Mil94, CXW*05, WTL*06]. The main difference here is that instead of parameters like the "exposedness" of points on the surface that affect weathering and wear, our materials respond to user interaction. To our knowledge, this effect has not been treated in computer graphics literature to date.

Appearance of fabric. Textile materials play an important role in our every-day lives and have received considerable attention in the graphics community. The models proposed in literature span a wide range. On one end are analytical BRDFs based on microfacet models [APS00] or fiber geometry [IM12, SBDDJ13]. On the other end, purely image-based models, like the bidirectional texture function (BTF) [DvGNK99], capture the view- and light-dependent reflectance of surface samples. The characteristic properties of fabric make it particularly suitable for volumetric scattering models that can either be isotropic [SKZ11] or anisotropic like microflake representations [JAM*10, ZJMB11]. Today's microflake models can be represented and rendered efficiently [HDCD15]. They have been shown to produce similar rendering quality as explicit fiber models [KSZ*15], although such models are still several orders of magnitude too expensive for real-time rendering. In our work, we use a hybrid approach. Using a BTF as input data, we fit the parameters of a custom spatially varying BRDF which, in turn, is derived from a single-layer microflake model. We derive our model from linear transport theory similarly to Hanrahan et al. [HK93]. However, this early work relied on isotropic scattering and the empirical Henyey-Greenstein phase function while we use the microflake theory to derive an anisotropic BRDF. Unlike most existing analytical BRDF models [APS00, IM12, SBDDJ13], ours allows for asymmetric or tilted microflake distributions. The only model derived under the assumptions of the microfacet theory for the purpose of simulating metalic surfaces that shares these characteristics is the ellipsoid NDF [DWMG15]. Our model decouples material parameters (gloss and color) from interaction parameters (fiber directionality), enabling the plausible manipulation of fiber distributions without leaving the appearance space of the prescribed material. Similar to other BRDF models our technique trades-off accuracy for performance, mainly, by ignoring parallax, inter-reflection and assuming simplified representations of the microgeometry. The implications of these properties are that our model would deviate from the ground-truth BTF rendering. However, our model has smaller memory footprint, requires less computation and therefore is more widely applicable.

4 Proposed Approach

Our solution is comprised of three major parts: reflectance model, parameter fitting to measured data and interactive simulation. We discuss each one of them in the following subsections.

4.1 Reflectance model

Physically correct modeling of pile fabrics under tactile interaction requires parametrization of the appearance based on the microscopic structure. Materials like Alcantara are comprised of many loose fibers that are bound together into a single fabric. This characteristic feature makes them deviate from the assumptions of microfacet theory of small mirror-reflecting surfaces. The microflake theory is a newer approach of looking at the process of modeling the microscopic geometry. The key difference is that surfaces are replaced by volumes of oriented microscopic particles. We take a

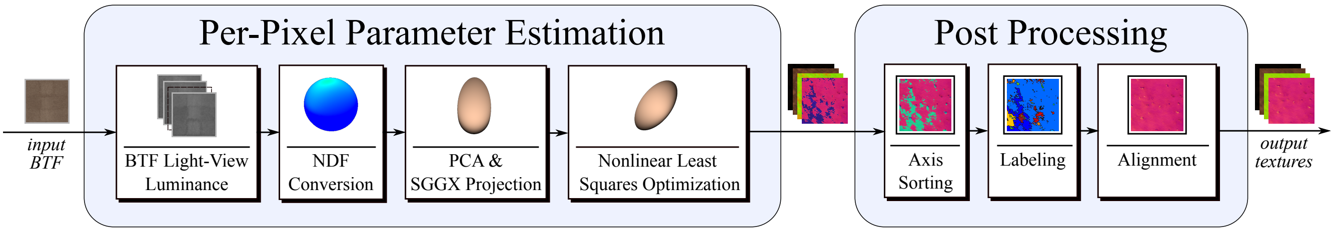

Figure 2: Our SVBRDF fitting pipeline. We use hemispherical bidirectional texture functions (BTF) measured by a camera dome as input to our pipeline. Each operation is performed per pixel in the main parameter estimation phase. Post processing is performed on the whole texture based on information about the neighboring similar regions. The output of the pipeline is a set of texture files containing SVBRDF parameters.

similar approach to Heitz et al. [HDCD15] who introduced the classical Trowbridge-Reitz (GGX) model [TR75] to microflake theory by extending its distribution of normals to the whole sphere. The preferred framework of microflake models up to this point was to use volume rendering techniques. This method of rendering is unsuitable for real-time rendering because it introduces a random walk step which severely reduces the convergence rate. We propose a different approach which can enable regular surface rendering techniques to be applied.

We start our derivation from the radiative transfer equation limited to single scattering. It is a fundamental assumption made when deriving closed form analytical BRDF models in the microfacet theory. Therefore, our approach would be comparable in quality to these models. We further assume optically thick semi-infinite medium per surface element. This assumption removes the need of performing a random walk on the surface. The reflected radiance at a given point on the surface $L(\mathbf{x},\omega_o)$ under these assumptions can be expressed as:

$$L(\mathbf{x},\omega_o) = \int_0^\infty e^{-\sigma_t(\omega_o)t}\sigma_s(\omega_o) f_p(\omega_i \to \omega_o) L_i(\omega_o,\omega_i,t)\,\mathrm{d}t. \quad (1)$$

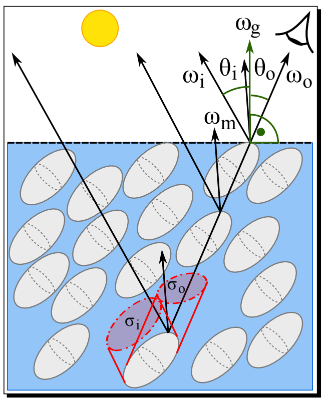

Figure 3: We assume homogeneous semi-infinite volume per surface pixel and far-field illumination. It enables us to use parallel incident lights along the scattering direction.

Light is scattered along the outgoing direction from the light source proportional to the scattering coefficient $\sigma_s(\omega_o) = k_s\rho\sigma(\omega_o)$. The $\sigma(\omega_o)$ and ρ are the projected area of the microscopic flake and the volume density, respectively. The albedo k_s determines the specular contribution in our model. Afterwards, light is exponentially attenuated along distance t at rate proportional to $\sigma_t(\omega_o) = \rho\sigma(\omega_o)$. $f(\omega_i \to \omega_o)$ represents the phase function which in our case is a specular reflectance function using the GGX distribution of normals extended on the sphere (SGGX). $L_i(\omega_o,\omega_i,t)$ represents the in-scattered radiance. We further assume far field illumination which lets us approximate the light transport as parallel rays scattered towards the outgoing direction. Complete picture of the assumptions that we apply to compute the radiance is shown in Figure 3. The resulting equation for the in-scattered radiance is:

$$L_i(\omega_o,\omega_i,t) = L_i(\omega_i)\exp\left(-\sigma_t(\omega_i)\frac{\cos\theta_o}{\cos\theta_i}t\right). \quad (2)$$

Both cosine factors are computed relative to the surface normal. They express the path difference between the incident and outgoing light. Inserting equation (2) into (1) lets us integrate the RTE under the above-stated assumptions:

$$L(\mathbf{x},\omega_o) = \underbrace{k_s\frac{f_p(\omega_i \to \omega_o)\sigma(\omega_o)}{\sigma(\omega_o)+\sigma(\omega_i)\frac{\cos\theta_o}{\cos\theta_i}}}_{f_s} L_i(\omega_i). \quad (3)$$

Note that the volume density ρ cancels out because it is part of the numerator and denumerator. The main implication of this property is that in semi-infinite medium the appearance is unaffected by scaling of the average free path length.

The phase function that we use is the mirror-like function proposed by Heitz et al. [HDCD15]:

$$f_p(\omega_i \to \omega_o) = \frac{D(\omega_h)}{4\sigma(\omega_o)}. \quad (4)$$

The SGGX distribution of normals $D(\omega_h)$ is dependent on the microscopic normal vector ω_h defined as the half-way vector between the incident and outgoing light directions. We refer to the original paper for more information. Note that we use the reciprocal parameters in our notation which is consistent with Jakob et al. [JAM*10].

We have experimented with adding a Fresnel term based on measured values of the refractive index; however, it did not result in improvement of the overall quality compared to measured data.

The model is also extendable to multiple scattering by evaluating the complete light transport. Similar approach was applied by Heitz et al. [HHdD16] in the context of the microfacet theory and Hanrahan et al. [HK93] for simulating layered materials. However, computing multiple bounces along material's surface can be quite costly. Therefore we add a regular isotropic diffuse term to compensate for the lost energy and the residual due to the complex microscopic structure. Our final model which can be substituted in the rendering equation is as follows:

$$f(\omega_i \to \omega_o) = k_s \frac{D(\omega_h)}{4\cos\theta_i \cos\theta_o \left(\frac{\sigma(\omega_o)}{\cos\theta_o} + \frac{\sigma(\omega_i)}{\cos\theta_i}\right)} + \frac{k_d}{\pi}. \quad (5)$$

The additional $\cos\theta_i$ factor accounts for the change of normals used by RTE simulation. The k_d represents the spectrum of the reflected light by the diffuse component. It is visible that the BRDF fulfills the requirement for reciprocity, i.e. $f(\omega_i \to \omega_o) = f(\omega_o \to \omega_i)$.

In the special case of isotropic BRDF where $\sigma(\omega_o) = 1$, $\sigma(\omega_i) = 1$ and $D(\omega_h) = \frac{1}{\pi}$ the model reduces to the Seeliger's Law [See88]. This property can be shown by rewriting equation (3) in the following form:

$$L(\mathbf{x}, \omega_o) = \frac{D(\omega_h)\cos\theta_i}{4(\sigma(\omega_o)\cos\theta_i + \sigma(\omega_i)\cos\theta_o)} L_i(\omega_i). \quad (6)$$

Similarly, if we apply the same constants in equation (5) we would derive the isotropic Chandrasekhar's BRDF [Cha60].

Importance sampling is essential to ensure fast convergence in ray tracing applications. The ideal importance sampling would be exactly according to the BRDF $f(\omega_i \to \omega_o)$. However, that's rarely possible. The most common approximation is to sample according to the distribution of normals. Our model is derived from tilted GGX distribution. It is one of the similarities of our model with the ellipsoid NDF derived in the supplemental material of Dong et al. [DWMG15]. Ignoring the other terms lets us sample our model in similar fashion. We refer to the original work for more details.

Modeling time-varying fabrics with this model requires the interpolation of material parameters. The geometry related parameters are contained within the distribution of normals (NDF). We use the original SGGX distribution which contains a covariance matrix decomposed into rotation and standard deviation part. Interpolation can be performed linearly on each of these parameters.

4.2 Parameter fitting

To validate the applicability of our model for approximating the appearance of real world materials we devised a pipeline which fits parameters to measurement Bidirectional Texture Functions (BTF). Our model has two main components: diffuse and specular. Fitting the specular term can be broken down into two parts: initial guess estimation and nonlinear least squares fitting. The initial guess step is required to prevent reaching suboptimal local minimum when the next step gets applied. First the BTF is sampled at a specific point on the surface and complete slice of all available pairs of light and view directions is taken with their corresponding reflected radiance to irradiance ratio. Direct inversion of the NDF with the geometry contribution would be unfeasible for the initial guess that we would like to estimate. Therefore we extract the NDF for a perfect mirror reflection by using the following relation:

$$D_{\mathrm{BTF}}(\omega_m) = \frac{1}{2\pi} \int_{\Omega^+} f_{\mathrm{BTF}}(\omega_i \to 2(\omega_m \cdot \omega_i)\omega_m - \omega_i)\, \mathrm{d}\omega_i. \quad (7)$$

This equation can be solved through numerical integration. We use stratified Monte Carlo integrator for this purpose. The resulting image is stored in parabolic map with the following coordinate mapping:

$$\begin{bmatrix} x' & y' \end{bmatrix} = \begin{bmatrix} 0.5 + 0.5\frac{x}{z+1} & 0.5 + 0.5\frac{y}{z+1} \end{bmatrix} \quad (8)$$

Proper normalization is not essential because the SGGX NDF is scale invariant.

We extract the major axes by applying PCA on the resulting NDF to determine the rotation of the SGGX covariance matrix. Given $D_{\mathrm{BTF}}(\omega_m) : \Omega^+ \mapsto R$ we perform a search for axes which satisfy the equation:

$$\omega_z = \arg\max_{\omega_m \in \Omega^+} (D_{\mathrm{BTF}}(\omega_m)) \quad (9)$$

$$\omega_x = \arg\max_{\omega_m \in \Omega^+ : \omega_z \cdot \omega_m = 0} (D_{\mathrm{BTF}}(\omega_m)) \quad (10)$$

$$\omega_y = \omega_x \times \omega_z \quad (11)$$

The only component left is the standard deviation which is proportional to the projected area on each axis. It can be extracted through a simple integration:

$$M = \begin{bmatrix} \omega_x & \omega_y & \omega_z \end{bmatrix}$$

$$S_v = \int_{\Omega^+} (M^T \omega_m^T) D_{\mathrm{BTF}}(\omega_m)\, \mathrm{d}\omega_m. \quad (12)$$

Having determined the major axes and the projected area we have a good enough initial guess on the specular contribution. To derive optimal parameters for the complete BRDF we take the estimated specular term and apply nonlinear least squares optimization. We compute the residual as the difference of the samples contained within the BTF and our BRDF function leading us to the following equations representing the residual:

$$g_k(W^k_{\omega_i \to \omega_o}) = f_{\mathrm{BTF}}(W^k_{\omega_i \to \omega_o}) - f(\omega_i^k \to \omega_o^k), \quad (13)$$

The sum of squares is expressed as:

$$G(W^k_{\omega_i \to \omega_o}) = \frac{1}{2} \sum_{k=1}^{N_L N_V} (g_k(W^k_{\omega_i \to \omega_o}))^2. \quad (14)$$

Given $G : R^n \mapsto R$. We must find global minimizer:

$${W^k_{\omega_i \to \omega_o}}^+ = \arg\min\{G(W^k_{\omega_i \to \omega_o})\} \quad (15)$$

This is a popular problem in mathematical optimization and many solutions exist to solve this problem. We used Levenberg-Marquardt method for this purpose by applying it to the specular and diffuse components of the BRDF.

4.3 Post processing

The resulting SVBRDF from the fitting procedure has many areas where the tangent directions of neighboring pixels are pointing towards opposite directions. It is essentially caused by symmetry of the BRDF as opposite bases would generate the same appearance. Interpolating this kind of bases would result in visible artifacts. The first step to alleviate this issue is to sort the basis according to standard deviation. By applying this method we are exploiting the mostly homogeneous structure of the material. However, due to symmetry in the BRDF it is still possible to have misalignment. Therefore, we further detect the outlines of areas where this behavior occurs and label the individual regions. Afterwards, we reorient each individual region towards the dominant direction (the region covering most area). It can be further realigned, so that the dominant direction points towards the preferred direction, ensuring smooth interpolation during interaction. All of these operations don't change the overall appearance as we are exploiting symmetries in the BRDF function.

5 Implementation and Results

We have implemented the parameter fitting pipeline in C++. We have experimented with CUDA implementation for some parts of the pipeline. However, it was mostly memory bandwidth limited, so the overall performance improvements were fairly modest. We use separate tools for the per pixel parameter estimation and the post processing step for faster development iteration when testing different experimental techniques.

We work with BTF that was captured with 30° rotations of the camera dome turntable and 11 cameras. It was rectified as a set of images with resolution 1024×1024 and resampled with 151×151 light-view vector pairs. It was compressed in Full Matrix Factorization (FMF) format [KMBK03]. Our SVBRDF fitting procedure takes about 150 minutes per 200×200 patch on an Intel i7-5820k system running at 3.30 GHz. We work with chunks of 64×64 randomly selected light-view vector pairs to ensure low error. It is the leading factor affecting the current execution times. The main bottleneck is the nonlinear least squares optimization followed by the Monte Carlo integration used for estimating the NDF. We have experimented with some photometric normals techniques for estimating the initial guess of the SGGX basis axes; however, they did not yield better results than the NDF estimation technique that we proposed in this paper. We use a configuration of 4 dispatcher threads which manage information for 4 different pixels at the same time and a shared thread pool that processes the bulk of the data. Dispatcher threads assist in the data processing while they wait for completion of related tasks. This architecture ensures full utilization of the CPU.

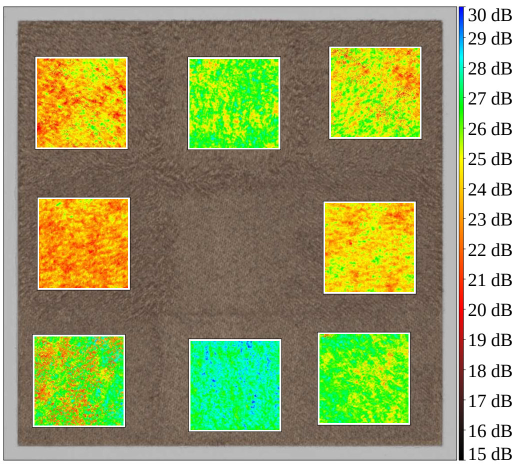

Figure 4: PSNR in the angular domain for each pixel of the SVBRDF produced for Alcantara material sample. The original BTF lit from top is shown as a background. Some of the areas show distinct striped patterns which can be explained by inter-reflectance and shadowing effects.

After the initial per pixel parameter step is complete we get above 20 dB peak signal-to-noise (PSNR) ratio in the angular domain for most of the samples as can be seen in Figure 4. It guarantees us below 10% error for most of the surface. Increasing this ratio would require further fitting of multiple lobes and taking into account light transport related effects. It is possible to apply multiple lobes in the parameter estimation procedure; however, it would be more unstable. It should be noted that our results are based on a plane estimation and subsequent rectification of photographs which results in different quality between the samples in the angular domain. The main reasons for this quality disparity is the difference in resolution

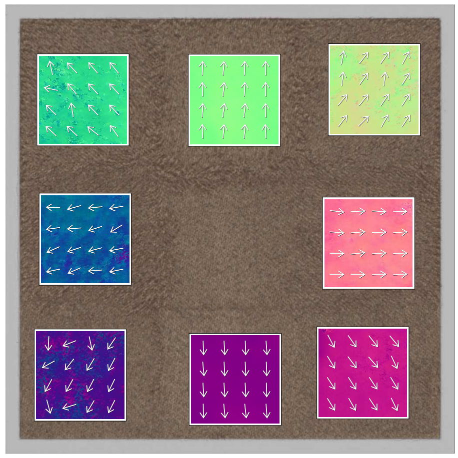

Figure 5: Orientation of tangent vectors of the SGGX basis after postprocessing. It is visible that the appearance is highly dependent on direction of material manipulation.

between top and grazing view angles and parallax inside the measured data. Our technique does not take parallax into account which is required to further improve the visual quality. Standard solution to this problem is to measure geometry by using structured lighting techniques. Applying structured lighting techniques would be challenging as most of the surface detail is close to the error caused by the instability of projectors. There are also inter-reflection and subsurface scattering effects which impede exact geometry acquisition. These problems can be bypassed by working only with the top-down view of the BTF material. Example images are provided in Figure 8. Ground truth rendering of the original BTF with the same scene is provided in Figure 9.

Post processing the data is crucial to guarantee smooth interpolation across the whole surface. We have chosen eight patches which are brushed in the same direction. Essentially they represent eight different states of the material. The objective of the post processing stage is to also align them so that we can guarantee smooth transitions when tactile interaction is applied. Overall the resulting vector field is smooth after our post processing step which is confirmed by Figure 5. Further confirmation of the alignment of the SGGX basis with the brushing directions is shown in Figure 6. The post processing step as we apply it works only on the SGGX basis vectors and the standard deviation and does not introduce any errors as it performs swapping and axis reversal operations.

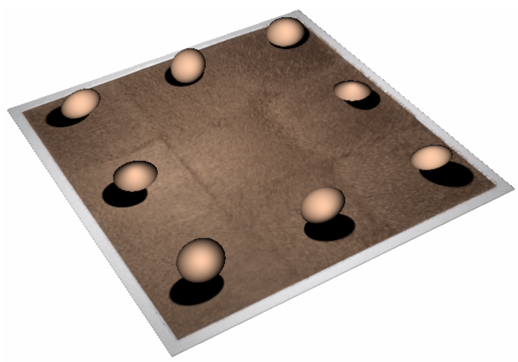

Figure 6: Ellipsoid distributions fit to the original BTF representation. The distributions are leaning towards the brushing direction connecting appearance and microscopic structure.

We have developed an application that demonstrates the capability of our SVBRDF technique to simulate tactile interaction. We provide images rendered by our framework in Figure 7. We have used the real-time GPU ray tracing backend in our custom rendering framework. However, it is directly applicable in regular GPU rasterization based deferred shading by storing information about tangent space and one additional specular power component. Therefore we consider our technique as widely applicable for rendering in many scenarios. It is visible that the change of appearance when interacting with the material is handled by our approach. Our painting solution is based on software rasterization on the GPU of two-dimensional capsules with linear fall-off along the edges. Interpolation is performed linearly with spherical linear interpolation for the SGGX basis. Editing can be performed by taking into account that smoothness corresponds to lower standard deviation values and anisotropy can be boosted by reducing the value of the smallest ellipsoid axis.

6 Discussion

We are working with SVBRDF representation of measured material. This class of techniques shares similar set of simplification related to handling parallax, inter-reflection and microgeometry.

Parallax is usually handled by geometry acquisition. In our case we work with a material with fine fiber structure which is challenging for the currently existing structured light techniques. They are negatively affected by the instability of the projected structured patterns on the surface. Inaccuracies in the rotations of the turntable sample holder can affect the final quality of the reconstruction. Also, they are negatively affected by inter-reflection and subsurface scattering effects. All of these factors make the exact acquisition of reconstructed geometry extremely challenging. Our research does not focus on resolving these issues; therefore, this is left as future work.

Inter-reflection can be handled by our model. Full light transport simulation is prohibitively expensive for real-time rendering of complex scenes. Another concern is that the parameter estimation would take much longer time. Therefore, within this paper we demonstrated only single scattering data fits. It is left as future study to determine the quality change when multiple scattering is applied for high quality offline rendering.

Complex microgeometry can be handled through multiple lobe simulation. The main concerns when applying multiple lobes is that the parameter estimation becomes unstable as we are working with ill-posed problem. Another concern is that applying interpolation becomes more ambiguous when simulating tactile interactions. These are legitimate concerns and we are looking forward to address in future work.

7 Conclusions

We introduced a SVBRDF model based on homogeneous infinite volume approximation of pile fabrics which can simulate the change of appearance based on tactile interactions. Our model is simple, efficient and easy to implement both in ray tracing and traditional rasterization graphics pipelines.

We developed a complete pipeline for measuring, fitting and simulating materials with microscopic fiber surfaces. We validated our pipeline against measurements of Alcantara material and proven that it provides good quality for materials simulated with our low-dimensional SVBRDF model. Parallax and inter-reflection still remain a challenging problem. However, they would require introducing sophisticated methods for geometry acquisition, which are out of the scope of our research.

We demonstrated the applicability of our SVBRDF model for modeling the tactile interaction between human hand and object made of this class of materials in real-time. We further showed that our model is easy to edit which makes it applicable for generating synthetic materials for the purpose of designing objects made of pile fabrics.

Acknowledgements

We would like to thank Sebastian Werner for the fruitful discussions, and Sebastian Merzbach and Dennis den Brok for assistance in acquisition and processing of BTF data. We would also like to thank the anonymous reviewers for their helpful comments and suggestions. This work was funded by the X-Rite Chair and Graduate School on Digital Material Appearance.

References

[APS00] ASHIKMIN M., PREMOŽE S., SHIRLEY P.: A microfacet-based BRDF generator. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2000), SIGGRAPH '00, ACM Press/Addison-Wesley Publishing Co., pp. 65–74.

[Cha60] CHANDRASEKHAR S.: *Radiative transfer*. 1960.

[CXW*05] CHEN Y., XIA L., WONG T.-T., TONG X., BAO H., GUO B., SHUM H.-Y.: Visual simulation of weathering by γ-ton tracing. *ACM Trans. Graph. 24*, 3 (July 2005), 1127–1133.

[DvGNK99] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. *ACM Trans. Graph. 18*, 1 (Jan. 1999), 1–34.

[DWMG15] DONG Z., WALTER B., MARSCHNER S., GREENBERG D. P.: Predicting appearance from measured microgeometry of metal surfaces. *ACM Trans. Graph. 35*, 1 (Dec. 2015), 9:1–9:13.

[HDCD15] HEITZ E., DUPUY J., CRASSIN C., DACHSBACHER C.: The SGGX microflake distribution. *ACM Trans. Graph. 34*, 4 (July 2015), 48:1–48:11.

[HHdD16] HEITZ E., HANIKA J., D'EON E., DACHSBACHER C.: Multiple-scattering microfacet BSDFs with the smith model. *ACM Trans. Graph. 35*, 4 (July 2016), 58:1–58:14.

[HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1993), SIGGRAPH '93, ACM, pp. 165–174.

[IM12] IRAWAN P., MARSCHNER S.: Specular reflection from woven cloth. *ACM Trans. Graph. 31*, 1 (Feb. 2012), 11:1–11:20.

[JAM*10] JAKOB W., ARBREE A., MOON J. T., BALA K., MARSCHNER S.: A radiative transfer framework for rendering materials with anisotropic structure. In *ACM SIGGRAPH 2010 Papers* (New York, NY,

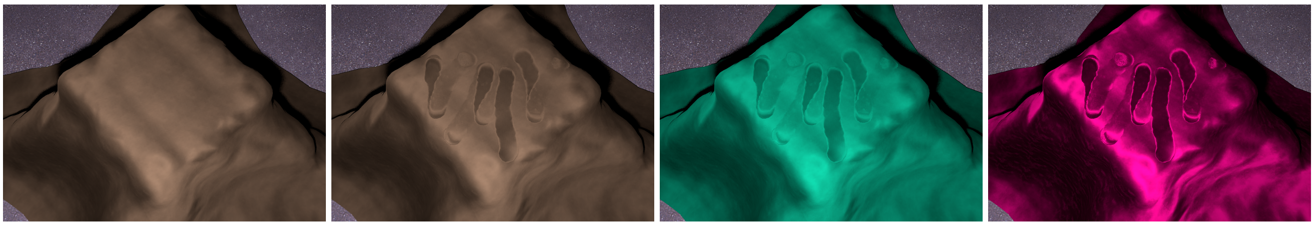

Figure 7: Our model can be manipulated in real-time and it contains parameters that are intuitive for editing. The first image on the left is a surface whose SVBRDF was generated by our parameter fitting pipeline. The second image displays tactile interaction with the material. The third image presents simple recoloring of the SVBRDF to green color. The last image on the right presents recoloring and boosting the anisotropy which is controlled by the standard deviation.

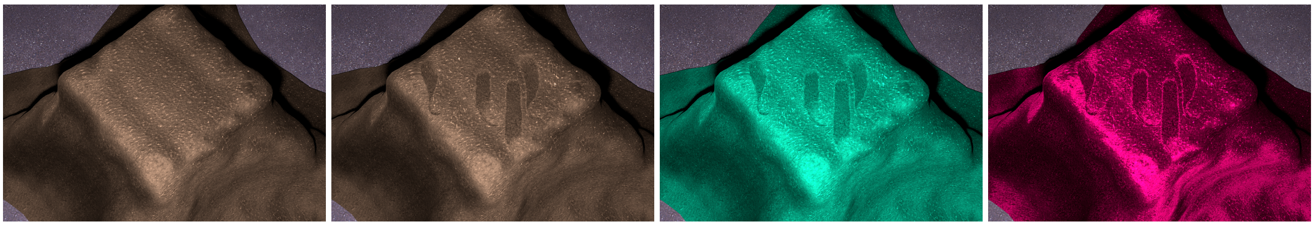

Figure 8: Higher level of detail can be achieved if the top-down view of the BTF is taken which is not affected by parallax and rectification errors. The first image on the left is a surface whose SVBRDF was generated by our parameter fitting pipeline. The second image displays tactile interaction with the material. The third presents simple recoloring of the SVBRDF to green color. The last image presents recoloring and boosting the anisotropy which is controlled by the standard deviation.

USA, 2010), SIGGRAPH '10, ACM, pp. 53:1–53:13.

[KMBK03] KOUDELKA M. L., MAGDA S., BELHUMEUR P. N., KRIEGMAN D. J.: Acquisition, compression, and synthesis of bidirectional texture functions. In *In ICCV 03 Workshop on Texture Analysis and Synthesis* (2003).

[KSZ*15] KHUNGURN P., SCHROEDER D., ZHAO S., MARSCHNER S., BALA K.: Matching microappearance models to real fabrics. *ACM Trans. Graph.* (2015).

[Mil94] MILLER G.: Efficient algorithms for local and global accessibility shading. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1994), SIGGRAPH '94, ACM, pp. 319–326.

[SBDDJ13] SADEGHI I., BISKER O., DE DEKEN J., JENSEN H. W.: A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph. 32*, 2 (Apr. 2013), 14:1–14:12.

[See88] SEELIGER R.: *Munch. Akad. II. Kl. Sitzungsber 18.* 1888.

[SKZ11] SCHRÖDER K., KLEIN R., ZINKE A.: A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum (Proceedings of EGSR 2011) 30*, 4 (July 2011), 1277–1286.

[SSR*07] SUN B., SUNKAVALLI K., RAMAMOORTHI R., BELHUMEUR P., NAYAR S.: Time-Varying BRDFs. *IEEE Transactions on Visualization and Computer Graphics* (Mar 2007).

[TR75] TROWBRIDGE T. S., REITZ K. P.: Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am. 65*, 5 (May 1975), 531–536.

[WTL*06] WANG J., TONG X., LIN S., PAN M., WANG C., BAO H., GUO B., SHUM H.-Y.: Appearance manifolds for modeling time-variant appearance of materials. *ACM Trans. Graph. 25*, 3 (July 2006), 754–761.

[ZJMB11] ZHAO S., JAKOB W., MARSCHNER S., BALA K.: Building volumetric appearance models of fabric using micro CT imaging. In *ACM SIGGRAPH 2011 Papers* (New York, NY, USA, 2011), SIGGRAPH '11, ACM, pp. 44:1–44:10.

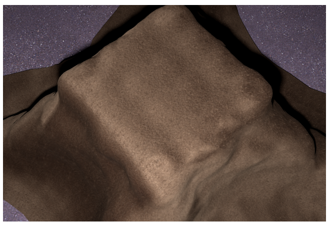

Figure 9: Ground truth BTF rendering. Note that due to the rectification error the effective resolution of the BTF decreases at lower viewing angles, making the appearance smoother. The problem is further emphasized by intermixing caused by the parallax error.