

# Towards Material Digitization with a Dual-scale Optical System

ELENA GARCES, Universidad Rey Juan Carlos & SEDDI, Spain

VICTOR ARELLANO, SEDDI, Spain

CARLOS RODRIGUEZ-PARDO, Universidad Rey Juan Carlos & SEDDI, Spain

DAVID PASCUAL-HERNANDEZ, SEDDI, Spain

SERGIO SUJA, SEDDI, Spain

JORGE LOPEZ-MORENO, Universidad Rey Juan Carlos & SEDDI, Spain



Fig. 1. Left, virtual scene rendered using eight different materials digitized with our optical device. Right, real photos of four of the materials taken under diverse illumination conditions: area lights, diffuse lighting, and directional lighting of high-resolution patch. The full set of real images is shown in the paper and in the supplementary material.

Existing devices for measuring material appearance in spatially-varying samples are limited to a single scale, either micro or mesoscopic. This is a practical limitation when the material has a complex multi-scale structure. In this paper, we present a system and methods to digitize materials at two scales, designed to include high-resolution data in spatially-varying representations at larger scales. We design and build a hemispherical light dome able to digitize flat material samples up to 11x11cm. We estimate geometric properties, anisotropic reflectance and transmittance at the microscopic level using polarized directional lighting with a single orthogonal camera. Then, we propagate this structured information to the mesoscale, using a neural network trained with the data acquired by the device and image-to-image translation methods. To maximize the compatibility of our digitization, we leverage standard BSDF models commonly adopted in the industry. Through extensive experiments, we demonstrate the precision of our device and the quality of our digitization process using a set of challenging real-world material samples and validation scenes. Further, we demonstrate the optical

Authors' addresses: Elena Garces, elena.garces@urjc.es, Universidad Rey Juan Carlos & SEDDI, Spain; Victor Arellano, victorarella@gmail.com, SEDDI, Spain; Carlos Rodriguez-Pardo, carlos.rodriguezparadojimenez@gmail.com, Universidad Rey Juan Carlos & SEDDI, Spain; David Pascual-Hernandez, david.pascual@seddi.com, SEDDI, Spain; Sergio Suja, sergio.suja@seddi.com, SEDDI, Spain; Jorge Lopez-Moreno, jorge@jorg3.com, Universidad Rey Juan Carlos & SEDDI, Spain.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3592147>.

resolution and potential of our device for acquiring more complex material representations by capturing microscopic attributes which affect the global appearance: we characterize the properties of textile materials such as the yarn twist or the shape of individual fly-out fibers. We also release the SEDDIDOME dataset of materials, including raw data captured by the machine and optimized parameters.

CCS Concepts: • **Computing methodologies** → **Rendering; Reflectance modeling**; • **Hardware**;

Additional Key Words and Phrases: digitization, capture system, appearance capture

## ACM Reference Format:

Elena Garces, Victor Arellano, Carlos Rodriguez-Pardo, David Pascual-Hernandez, Sergio Suja, and Jorge Lopez-Moreno. 2023. Towards Material Digitization with a Dual-scale Optical System. *ACM Trans. Graph.* 42, 4, Article 1 (August 2023), 13 pages. <https://doi.org/10.1145/3592147>

## 1 INTRODUCTION

A core problem in computer graphics research is acquiring and modeling the appearance of real-world materials [Dorsey et al. 2010; Guarnera et al. 2016]. Ideal digital representations should be both realistic and easy to edit, enhancing 3D experiences while minimizing the complexity of scene editing operations.

Bidirectional texture functions (BTFs) [Dana 2001] are an accurate way to represent the optical properties of materials, modeling

radiance as a 6D function parameterized by light and camera directions. However, acquiring this representation is difficult; requiring expensive scanning hardware and demanding a lot of memory. Furthermore, BTFs lack the flexibility of analytic material models, such as microfacet-based models [Cook and Torrance 1981; Torrance and Sparrow 1967] which are generally more compact, easy to compute, and can be designed to be more intuitive to manipulate by artists [Burley 2015]. Due to their widespread use in real-time rendering, these models have become the standard in the industry. We also adopt the spatially-varying bidirectional scattering distribution function (SVBSDF), including anisotropy and transmittance effects.

In recent years, estimating spatially-varying parameters for microfacet based models from images or videos, even acquired with casual devices such as smartphones, has become a more constrained and well-defined problem thanks to the ability of neural networks to learn data-driven priors [Deschaintre et al. 2018, 2019, 2020; Dong 2019; Guo et al. 2021; Henzler et al. 2021; Li et al. 2018; Martin et al. 2022; Vecchio et al. 2021; Zhou and Kalantari 2021, 2022]. However, all these methods are limited to isotropic and opaque materials, not covering more complex phenomena, such as anisotropy and translucency. The works of Merzbach *et al.* [2019], Ma *et al.* [2021], and Vidaurre *et al.* [2019] are some of the few exceptions which can recover anisotropic parameters [Ashikhmin and Shirley 2000; Geisler-Moroder and Dür 2010] in calibrated and controlled setups (*e.g.*, known geometry), although the latter does not recover textures, and none of them take into account transmittance effects, which, in general, are costly and difficult to acquire for heterogeneous materials [Frisvad et al. 2020]. Further, none of the existing methods is able to capture simultaneously the material at the micro and the mesoscale, a desirable property for complex multi-scale structures (*e.g.*, textiles, leathers, fur, etc.)

We present an optical acquisition system that captures materials at two scales in an end-to-end pipeline. At the microscale, we reach an optical resolution of  $1.8\mu\text{m}/\text{px}$  (or 14200PPis), sufficient to observe small-scale structures. At the mesoscale, we get a resolution of 1036 PPIs which cover mid-level patterns. Further, we capture transmittance effects through backlighting. Designing this device posed several challenges to minimize occlusions and maximize optical resolution, quality, and effective field of view. Additionally, we had to implement specific calibration procedures to set the radiance of all cameras and lights to the same calibration space.

Our method to estimate material parameters has two steps: In the first step, we leverage light polarization and directional lighting to optimize the SVBSDF through differentiable rendering. In the second step, we propagate these maps to the mesoscale using a neural network trained per material using captured data. Although our technique has resemblances with previous work on material acquisition [Merzbach et al. 2019; Nam et al. 2016], to our knowledge, we are the first method to propose this kind of dual-scale end-to-end estimation pipeline. To demonstrate the accuracy of our methods, we gather a dataset of textiles and leather materials with complex and varied micro-structures that we digitize with our machine. We additionally capture several images of the materials in various illumination settings that help us validate that our digitizations match the real world.

In summary, we propose the following contributions:

- A custom-built optical capture system for capturing reflectance properties at both micro and meso scale, capable of acquiring, in an end-to-end pipeline, anisotropic reflectance, transmittance, and micro-level details.
- We evaluate different design choices and optimization approaches to estimate a standard physically-based microfacet SVBSDF, including anisotropy and transmittance parameters.
- We evaluate our digital materials by comparing them to real-world scenes with different lighting conditions.
- The SEDDIDOME dataset that contains 36 materials with raw captured data and optimized parameters.

## 2 RELATED WORK

Capturing materials requires a strategy that takes into account appearance factors such as isotropic/anisotropic behavior, opacity, and texture, as well as hardware factors such as cost, accuracy, and measurement resolution.

*Rigid Capture Devices.* BTF Lightdome or gonioreflectometers are the most controllable means to measure spatially-varying material appearance [Dana 2001; Marschner et al. 2005; Schwartz et al. 2014]. Alternatively, movable gantries can be used for point-based BRDF measurements [Dupuy and Jakob 2018; Filip and Vávra 2014; Matusik et al. 2003a; Nicodemus et al. 1992; White et al. 1998]. Since these approaches require a large amount of storage and capture time, much effort has focused on optimizing the number of samples to minimize capture time and space while preserving the optical appearance of the material [Dong et al. 2010; Matusik et al. 2003b; Nielsen et al. 2015; Yu et al. 2016].

Additional factors to consider are the size and portability of the device [Havran et al. 2017], wherein pixel resolution is often sacrificed for sample size. Several existing systems can capture at the microscale: Nam *et al.* [2016] reaches a resolution of  $0.62\mu\text{m}$  per pixel, overcoming the challenges of capturing at such a small scale by using vibration absorption mechanisms and high-precision hardware. Alcain *et al.* [2019] use a single monocular camera, a filter wheel, and rail-based motorized focus with collimated light sources and a vertical sample holder. Graham *et al.* [2013] captures a patch of skin of  $10\mu\text{m}$  with a twelve-light dome hemisphere. In the case of metallic surfaces, microgeometry has been measured using specific hardware such as an optical profilometer [Dong et al. 2015].

Our system achieves high accuracy and captures data at a fine scale. Unlike Alcain *et al.*, we rely on a liquid lens-based optics that minimizes movable components and a denser light sampling distribution. Importantly, to accommodate flexible materials, we use a horizontal design for the holder and dome. In contrast to prior work, our method supplements the main microscopic system with an additional camera set to measure mesostructure, combining the benefits of high optical resolution and detail with greater sample sizes and optical effects observable by the human eye.

*Lightweight Capture Systems.* Casual systems use readily available devices (like smartphones) [Dong 2019; Guarnera et al. 2016]. Ngan and Durand [2006] capture BTFs by transferring texture statistics to a low quality texture sample. Ren *et al.* [2011] uses custom-made BRDF calibration chart containing several material appearances.

Wang *et al.* [2011] models glossy stationary surfaces using step lighting. Also leveraging stationarity, Aittala *et al.* [2015] required two pictures of the material (a flash and non-flash version). Some methods use neural methods trained on datasets to relax the capture constraints, requiring a single image [Deschaintre *et al.* 2018, 2019, 2020; Dong 2019; Guo *et al.* 2021; Henzler *et al.* 2021; Li *et al.* 2018; Martin *et al.* 2022; Shi *et al.* 2020; Vecchio *et al.* 2021; Zhou and Kalantari 2021, 2022], or multiple images [Deschaintre *et al.* 2019; Guo *et al.* 2020]. Recent research combined micro-flake models and procedural yarns to digitize simple woven structures [Guarnera *et al.* 2017; Jin *et al.* 2022].

*Fiber-Level Capture.* The look of fiber-based materials is closely related to their microstructures. Advanced rendering models can simulate the light scattering at the fiber level, but even simple hair and fur models [Marschner *et al.* 2003] require complex laboratory measurements to capture radiance. For fabric materials, it is important to characterize the cross-sections and the dye molar concentrations to correctly simulate the light transport inside the fibers [Aliaga *et al.* 2017]. The aggregated effect of the fibers forming yarns can be approximated by generative models [Khungurn *et al.* 2015; Sadeghi *et al.* 2013] or measuring the volume with a micro CT scan [Zhao *et al.* 2011, 2012, 2016]. Our optical system, thanks to reaching microscopic resolution, can be used to obtain fine grained details at a fraction of the cost and time required by the complex and expensive micro-tomography systems described above. For further details in fabric rendering models acquisition systems, we refer to the survey by Castillo *et al.* [2019].

*Translucency.* Research on the importance of translucency for industrial applications (digital twins, 3D printing, ...) has increased in recent years, and some systems can measure it with only three spectral samples [Iser *et al.* 2022]. For fabrics, it has huge impact causing small geometric details to blur and increase color saturation and luminance for certain illumination directions.

Transmittance (or phase anisotropy) has been studied in the field of perception [Fleming and Bühlhoff 2005; Xiao *et al.* 2014, 2020], but quantifying its contribution in a given image has not solved to date. Current solutions require complex devices to capture the 8D BSSRDF of homogeneous materials given known geometry. Simpler models, such as the dipole [Jensen *et al.* 2001], can capture subsurface scattering with a single camera. However, objects with arbitrary geometry and heterogeneous transmittance remain a challenge. Accurate geometrical measurements are difficult when the local curvature is comparable to the depth of the subsurface light transport, as observed in commercial 3D scanners [Guerra *et al.* 2019]. For more information on acquisition models, refer to Frisvad *et al.* [2020]'s recent survey. For thin objects, direct transmittance can be modelled using a BSDF, with light exiting at the same sampled point, ignoring multiple-scattering effects, or partially aggregating them with other components (*e.g.*, diffuse reflectance) [Burley 2015]. We follow this approach, which is described in Section 6.2.

### 3 OPTICAL ACQUISITION SYSTEM WITH TRANSMITTANCE

#### 3.1 System Design

We now discuss our acquisition setup, which presents several challenges motivated by the following goals:

- (a) Capture *microscopic* details at the fiber level (3–20  $\mu\text{m}$ ).
- (b) Capture *mesoscale* textures of a medium-size patch ( $\sim 11 \times 11$  cm).
- (c) Leverage optical components whenever possible to ease the optimization processes.
- (d) Include transmittance lighting.
- (e) Enable an end-to-end automatic capture pipeline.
- (f) Limited number of movable pieces to reduce hardware failure.

Prior work on measuring material appearance either focuses on a single scale [TAC 2016; Alcain *et al.* 2019; Nam *et al.* 2016], a single point [Dupuy and Jakob 2018], or inexpensive and casual hardware –neglecting transmittance or anisotropic effects– so none of the existing solutions meet our goals.

The microscopic optical system is the most critical piece of our design. It should be small enough to avoid potential occlusions and shadow projections while allowing for great magnification. Our design includes a miniaturized high-resolution industrial camera and lens pair, and a liquid lens with electrically tunable focus<sup>1</sup>. As opposed to previous work [Alcain *et al.* 2019], the usage of a tunable lens circumvents the usage of motorized pieces, which further reduces the overall size and maintainability costs of our system. Besides, this tunable lens serves as an extension tube, which helps us reduce the minimum working distance of our optical system.

Contrary to existing four-dimensional gantry-based capture setups [Dupuy and Jakob 2018; Holroyd *et al.* 2010], we rely on a solid 3D-printed structure (dome), where cameras and lights are set at a fixed position and controlled electronically. However, since the field of view of the microscopic camera is less than a  $\text{cm}^2$ , covering a small sample of the material, we have implemented a movable holder structure with sub-millimetric precision. This allows us to position the material at any point in the XY axis, and capture any part of the material sample without the need of manual intervention. The holder's frame is designed to minimize the light occlusion at grazing angles. During capture, the material is placed on top of a 3mm transparent glass with an anti-reflection coating to avoid light inter-reflections while enabling backlighting.

Figure 2 shows a schematic of the device and several photos. It has four cameras: a microscopic one with a polarized tunable lens, a mid-range high-resolution camera, and two polar cameras at oblique angles. It has three lighting setups: directional polarized lighting, diffuse lighting, and backlighting. We provide detailed descriptions of these components in the following paragraphs.

*Light Dome.* The hemispheric light dome of radius 0.3 m is printed in solid PU material. It has apertures for placing individual collimated lights (or camera) modules (Figure 2 (d)). Each collimated LED light has a diameter of 19.89 mm and is composed of a high-power white LED, a collimating lens, and a linear polarizer (Figure 2 (b)) with s-polarization. To avoid speckles, the lens of the collimated LEDs is set to work in an out-focus-plane. This causes speckles to be blurred at the sample, with negligible impact on the quality of

<sup>1</sup><https://www.optotune.com/focus-tunable-lenses>

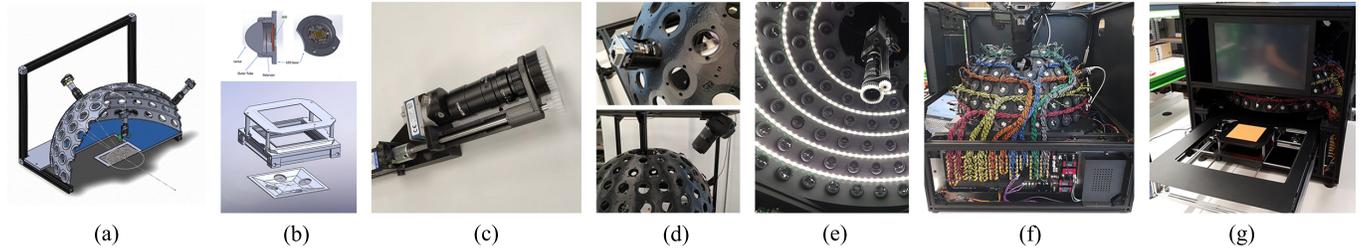


Fig. 2. Schemes and photos of our optical device. (a) Schema of a cross-section of the hemisphere including micro camera and one polar camera. (b) (Top) Schema of the collimated LED design to account for the polarizer and collimating lens; (Bottom) Schema of the main holder and backlight support. (c) Microscopic optical setup. (d) (Top) A Polar Camera; (Bottom) Mid-distance Camera. (e) Interior of the dome with the diffuse LED strip activated. (f) Exterior of the dome and wiring. (g) Holder and exterior cover.

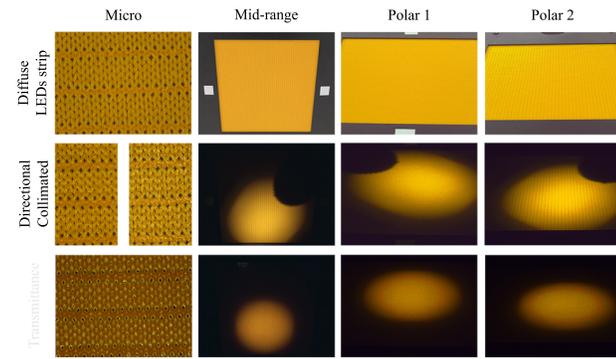


Fig. 3. Images captured by each camera (x-axis) and lighting setup (y-axis).

the captured images. We use a total of 127 white LEDs of 4000k and 400lm. The same type of lighting modules is placed at the back of the holder (Figure 2 (b-c)). We use a total of four backlights, which we aggregate to simulate diffuse transmissive illumination. All the collimated LED light modules point towards the central axis of the dome, where the microscopic camera faces. Consequently, we assume that parallel rays hit the sample in the collimated area during the inverse rendering optimization process (described in Section 6). In addition, we provide semi-diffuse radial illumination through several LED strips of 6000k and 1200lm (Figure 2 (e)). See Supplementary Material for further detail about the measured spectrum and color of the different lighting configurations. For the rest of the paper, the light coming from the diffuse LEDs will be referred to as *diffuse* lighting, while the light coming from the directional collimated LEDs will be referred to as *directional* LEDs or *directional* lighting.

**Cameras.** We capture the material at two scales. At the microscale we have placed an industrial camera Genie Nano C4900 of 18MP (pixel size: 1.25 $\mu$ m) with a 50mm lens and the auto-focusable liquid lens reaching a FOV of 8.9 x 6.7 mm and 14020 PPIs. At the mesoscale level, we have three cameras: two industrial cameras of the same model with a 16mm lens located at 30 degrees of elevation w.r.t the sample plane (azimuth positions can be visualized in Figure 6), reaching 490 PPIs; and a mid-range DSLR Canon EOS 5DSR paired with a 70mm macro lens, located at 60 degrees, reaching 1036 PPIs.

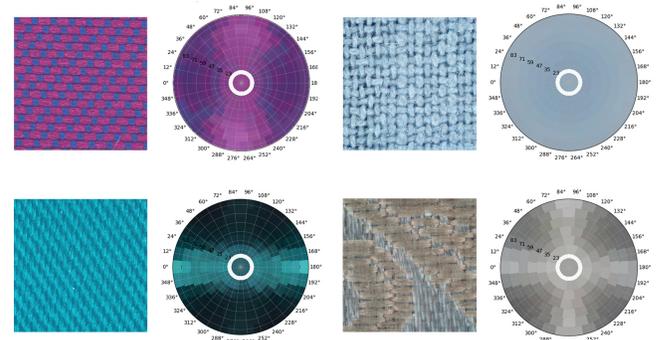


Fig. 4. Example of micro captures with diffuse lighting and circular plot corresponding to the average radiance for each directional LED projected in a circle.

Figure 3 shows the images captured by each camera and lighting setup, and Figure 6 shows the location in a 2D projection of the directional lights and cameras. Figure 4 includes microscale captures with diffuse lighting and a circular plot, where each cell contains the average radiance for each directional LED. Note that the directions of anisotropy for woven fabrics are clearly visible.

**Polarization.** The micro camera has an extra module with a polarization filter micro-controlled to adjust the polarization angle dynamically depending on the active luminary (Figure 2 (c)). The use of polarized imaging is important to separate diffuse and specular components [Ma et al. 2007], facilitating the estimation of the SVBSDF (Section 6) and other operations done at the micro level (Section 8). We polarize the light in the perpendicular direction with respect to the plane of incidence (s-polarization). In the camera, we have a motorized filter that allows to pass only the light that is polarized in the same direction as the source (P90) or in the opposite direction (P0). The light that only bounces once in the surface –specular reflection– will keep the polarization angle of the light source, while the light that enters in the material –subsurface scattering and single scattering– will be depolarized after bouncing multiple times. Thus, in ideal conditions, by subtracting the image taken with P0 from the image taken with P90, we obtain an image that contains mostly direct reflectance (examples in Figure 5).

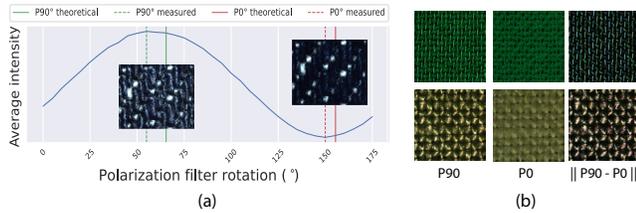


Fig. 5. (a) Influence of polarization in reflectance: Average intensity of the captured image (y-axis) varying the angle of polarization of the micro camera (x-axis). The insets are captured images at those peak angles. (b) A GREEN SATIN (top) and a LINEN (bottom) captured with different polarization modes.

#### 4 SVBSDF MATERIAL MODEL

Our optical device captures the appearance of complex translucent materials. Our model is based on Disney 2015 [Burley 2015] for maximum compatibility. However, we focus on two factors often overlooked in surface-based material acquisition models: anisotropy and transmittance. We roughly describe the appearance model in this section and the estimation pipeline in the following ones. The complete equations of the model are explained in the supplementary material.

Anisotropy is important, for instance, for rendering cloth – particularly wovens – since the directionality of the fibers and yarns causes light to be reflected or transmitted differently depending on the angle of incidence. Our anisotropic reflectance lobe has 16 spatially-varying parameters: basecolor or diffuse albedo  $\mathbf{b} \in \mathbb{R}^3$ , roughness  $\sigma_r \in \mathbb{R}$ , IOR  $\in \mathbb{R}$ , degree of anisotropy  $\alpha \in \mathbb{R}$ , specular tint (intensity of the colorization of the specular reflection)  $\rho_s \in \mathbb{R}$ , normals  $\mathbf{n} \in \mathbb{R}^3$ , and tangents (major axis of the anisotropy)  $\mathbf{t}_g \in \mathbb{R}^3$ .

Accurately capturing the transmittance effect is crucial for achieving realism when rendering non-opaque or thin-layered materials with volumetric appearance. Using just approximate visibility techniques, such as alpha maps, which do not fully consider transmittance and multiple scattering effects, can make the rendered material appear darker than intended. Transmittance is usually modeled by simulating light transport as it interacts with the internal structure of the material, simulating scattering, absorption, and refraction effects. Additionally, some real-time rendering engines have also developed techniques to approximate the transmittance effect with less computational cost, such as screen-space subsurface scattering [Jimenez and Gutierrez 2010]. We model transmittance using a single diffuse and spatially-varying lobe  $\mathbf{t} \in \mathbb{R}^3$  modulated by the angle of the underlying geometry. We found that for the transmissive samples in our set is enough (*i.e.*, fabrics), as light is quite scattered within, but more transparent materials would require taking into account the optimized surface normal and fitting a specular transmissive lobe, which is later discussed as future work. We further include an opacity map  $\tau \in \mathbb{R}$  necessary for modeling actual holes and cavities, which accounts for unscattered light directly passing through the material.

#### 5 OVERVIEW

Our parameter-estimation pipeline is illustrated in Figure 6. Our method has two distinct steps, in which the SVBSDF is estimated at the micro level first (Section 6), and then, propagated to the mesoscale (Section 7). If the material is spatially-varying and not all the variation is captured at the micro level in a single capture, we estimate as many micro SVBSDFs as necessary to cover the whole appearance variability.

As input to both steps, we use a microscale *photometric dataset*, which is composed of images taken with the micro camera and directional lights  $k$ , polarized at 0 degrees  $\{i_k^{P0}\}$ , polarized at 90 degrees  $\{i_k^{P90}\}$ , and the combination of both  $\{i_k = i_k^{P0} + i_k^{P90}\}$ . The latter images would be equivalent to non-polarized illumination. As input for the mesoscale propagation step we use the mid-range camera image taken under the diffuse lighting  $G$ : its RGB values will guide the propagation of the micro-level fitted parameters to larger samples, effectively preventing us from repeating the optimization process done at the microscale.

#### 6 MICRO OPTIMIZATION

At the micro level, we compute reflectance and transmittance in two steps, described in Section 6.1 and 6.2, respectively. First, we estimate the anisotropic reflectance parameters through an optimization process using the majority of lights  $Q_F$  of the device, as shown in Figure 6. Then, for transmittance and opacity (Section 6.2), we use the basecolor  $\mathbf{b}$ , and normals  $\mathbf{n}$ , resulting of the previous step, as well as the transmittance views  $Q_T$ .

##### 6.1 Anisotropic Reflectance Estimation

*Baseline Initialization.* The parameters are initialized with the data obtained using our hardware components, which is already a good starting point for the optimization process and works reasonably well for some simple cases. In addition, at the beginning of each optimization step, we set additional constraints (see Figure 11). The normals  $\mathbf{n}$  are initialized using classical photometric stereo techniques [Woodham 1980] using diffuse polarized images  $i_k^{P0}$ . The tangents are naïvely initialized in x direction. The roughness  $\sigma_r$  is initialized at a mid-range value (0.5). Smaller initial values (*e.g.*, 0.2) produce enough diffuse information in the image, but on average convergence is faster in this fashion. Extreme values such as 0.0 should be avoided, as they would create mostly-black images with some sparse highlights, usually far from the average sample. The initial diffuse color  $\mathbf{b}$  is the photo taken with the diffuse lighting, and the anisotropy  $\alpha$  and specular tint  $\rho_s$  are initialized with constant values of 0.1 and 0, respectively. We establish a minimal anisotropic value so that the modification of related parameters such as tangents affects the loss function. Once the loss function has reached a certain tolerance we remove this constraint, and anisotropy can get to the 0.0 limit.

The IOR is initialized using the difference between our two polarization modes following previous work [Ma et al. 2007]. However, we found a considerable limitation of the standard BSDF model to represent partially transmissive materials. Even though most dielectric samples in our database are theoretically in the 1.5-1.8 range, we observed significant specular reflection which was not

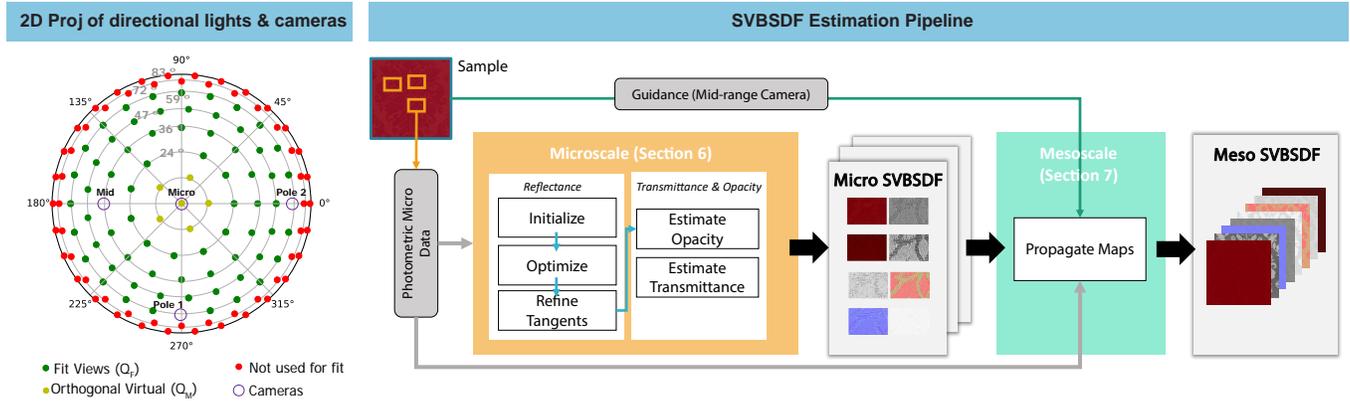


Fig. 6. Left, 2D projections of the directional lights and cameras of the dome. Right, an overview of the parameter estimation pipeline. Our method estimates the SVBSDF at two scales. At the microscale, we fit the parameters of the anisotropic reflectance lobe using directional lighting. The transmittance and opacity maps are computed next using backlighting and some reflectance maps. Once the micro SVBSDF is estimated, we propagate these values to the mesoscale using a neural network trained with data captured with the optical device.

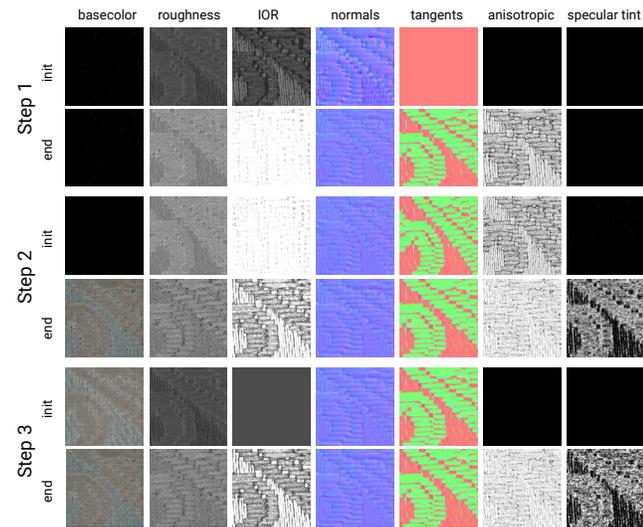


Fig. 7. Visualization of the inputs and outputs of the three optimization steps for the JARQUARD-SILVER fabric.

fully polarized and isolated by our filters. These secondary lobes aggregate transmitted, and then reflected-back photons, no longer maintaining the original wavelength (color) of the light source.

To augment the expressiveness of the model, we mapped the observed IOR values to  $([1, 4])$ , allowing otherwise unlikely, high specular values at surface normals oriented towards the camera plane. This combined with the  $\rho_s$  parameter, increases the accuracy of our maps for several cases (see Figure 11).

*Optimization.* Optimizing the SVBDF model can be quite complicated due to the non-orthogonal nature of its parameters. Even when

not optimizing for geometric data such as normals or tangents, multiple combinations can produce similar outcomes. Furthermore, working with an extended IOR range increases the non-orthogonality, making it less apparent the difference between a diffuse lobe and a rough, tinted specular lobe, which in turn increases the probability of finding local minima. This can result in local noise and artifacts on each map.

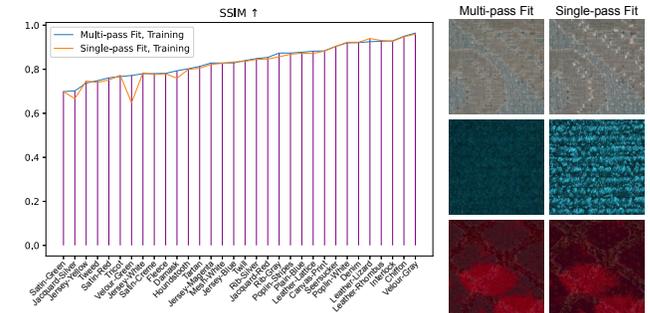


Fig. 8. On the left, we show SSIM errors on the validation views comparing direct (single-pass) and three-step (multi-pass) optimization approaches. Right: optimized base color (albedo) for three materials. Top Row: The multipass fit reduces dark and bright artifacts at the extreme of horizontal yarns. Middle row: The velvet albedo in a single pass accumulates a specular component, lowering the SSIM in multiple views. Bottom row: Leather does not have high transmissive-reflectance effects, and thus the IOR treatment is not as relevant. However, the normals are better estimated, reducing artifacts in albedo.

In Figure 8, we present the result of a direct optimization approach. Although the SSIM values of the rendered outcome are similar, the error on individual maps is significant, hindering future editing operations with the texture stack. To overcome this issue, we introduce a novel optimization strategy and losses aimed at minimizing local minima.

We divide our process into three steps, following common strategies: identifying ranges and relationships between parameters, isolating their optimization cycles when possible, and controlling their modification rate with respect to other parameters if they are jointly optimized. We avoided regularization terms, as they tend to produce implicit smoothing on the maps, which would remove many important features at the micro level, such as scratch marks or loose fibers. The steps, illustrated in Figure 7, are the following:

*Step 1* First we estimate the specular reflectance, without the influence of the diffuse lobe (basecolor), which is set to 0. The rest of the maps are freely optimized. The specular max value is highly correlated to normals and tangents, setting a good starting point for their correspondent maps.

*Step 2* Then, we incorporate the diffuse lobe (basecolor) in the optimization. This narrows the amplitude of the specular lobe, in turn refining the direction of the normals and tangents implied in the shading.

*Step 3* The tangents and normals from the previous step are quite precise so we fix them, initializing the rest of the maps following the baseline initialization, and optimize them. This last step is necessary due to the inherent ambiguity of the BSDF model, between diffuse and specular lobes, particularly for combinations of high IOR and high roughness. We introduce a prior, favoring a diffuse lobe first, as it is the most physically plausible scenario by default. In this fashion, the estimated values become more coherent with their neighbours at each map, reducing noise and artifacts, as we can see in Figure 8.

Our loss function used for all the optimization steps is composed of three terms:  $\mathcal{L}_{\text{rec}}$ ,  $\mathcal{L}_{\text{ortho}}$ , and  $\mathcal{L}_{\text{ior}}$ . The first term is the reconstruction error that measures the similarity between the captured image  $i_k$  and the rendered  $r_k$ . We use the smoothL1 ( $\hat{\ell}_1$ ) loss which is less sensitive to outliers than MSE.

$$\mathcal{L}_{\text{rec}} = \frac{1}{|Q_F|} \sum_{k \in Q_F} \hat{\ell}_1(i_k, r_k). \quad (1)$$

The second term is meant to ensure that specular highlights do not concentrate on orthogonal angles that are not used in the fit ( $\theta < 24^\circ$ ). Thus, assuming that we deal with dielectric materials, we penalize peak reflections that are higher than the maximum observed, a common practice when analyzing reflectance from grazing angles [Lavoué et al. 2021].

$$\bar{i} = \max_{1 \leq k \leq |Q_F|} i_k, \quad (2)$$

$$\mathcal{L}_{\text{ortho}} = \frac{1}{|Q_M|} \sum_{k \in Q_M} \max(0, r_k - \bar{i}), \quad (3)$$

where  $\bar{i}$  is the maximum reflection observed in the fit views  $Q_F$ , and the loss is only applied to certain virtual views  $Q_M$  uniformly distributed to fill the gaps of our capture machine in orthogonal positions (see Figure 6 (c)).

Finally, the third term establishes a soft upper bound for the IOR, favoring physically valid values as much as possible. We set it to

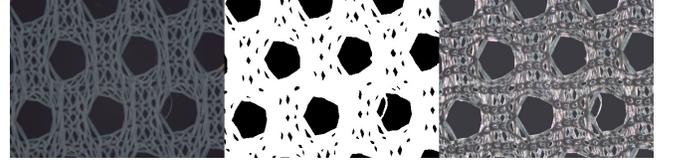


Fig. 9. From left to right: basecolor from baseline initialization, opacity map, and transmittance.

$\bar{\text{IOR}} = 1.78$  accounting for most common dielectric materials.

$$\mathcal{L}_{\text{ior}} = \left( \frac{\max(0, \text{IOR} - \bar{\text{IOR}})}{4 - \bar{\text{IOR}}} \right) \quad (4)$$

The final loss is built by aggregating the three terms as follows,

$$\mathcal{L} = k_1 \mathcal{L}_{\text{rec}} + k_2 \mathcal{L}_{\text{ortho}} + k_3 \mathcal{L}_{\text{ior}}, \quad (5)$$

with  $k_1 = 35.0$ ,  $k_2 = 1.0$ ,  $k_3 = 0.01$ .

*Tangent refinement.* As a final postprocessing, without further impact, we refine the tangent map to reduce ambiguity by doing a modulo over the range  $[-\pi/2, 3\pi/2]$ , constraining its values to the hemisphere.

## 6.2 Opacity and Transmittance Estimation

We compute the opacity map  $\tau$  using the optimized albedo, IOR, and normals from the previous step as input. Our opacity map is built from the intersection between three masks,  $\tau = \mathcal{B}_b \cap \mathcal{B}_{\text{IOR}} \cap \mathcal{B}_n$ , as only using the albedo or the image as a cue was not sufficiently accurate. A pixel will be opaque (1) in each mask if the value of each map is greater than a certain threshold. The thresholds are  $0.6\mu(b)$  and  $0.3\mu(\text{IOR})$  for  $\mathcal{B}_b$  and  $\mathcal{B}_{\text{IOR}}$ , respectively. Here  $\mu(\cdot)$  refers to the mean value of the map. Meanwhile, for estimating  $\mathcal{B}_n$ , we will consider as opaque those pixels that, given the absolute difference between  $\mathbf{n}$  and a perfectly orthogonal normal map  $\mathbf{n}_\perp$ , have a greater value than 0.015.

We compute the transmittance map  $t$  by taking the minimum value of the four transmittance images  $\{i_k^{P0}\}_{k \in Q_T}$ , where  $Q_T$  is the transmittance LEDs, taken with the diffuse polarization mode (P0). Figure 9 illustrates the opacity and transmittance maps resulting from this step.

## 7 MESOSCALE PROPAGATION

Estimating the SVBSDF at the microscale is a relatively straightforward and well-constrained problem given that our system is designed for this level. However, replicating the same optimization process at the mesoscale level poses some challenges. One issue is that directional lights, which collimate at a small area where the micro camera faces, cause non-homogeneous light spots at the mesoscale (see Figure 3). Additionally, the mid-range camera's tilted location (see Figure 2 (d)) to avoid occlusions with the micro setup causes soft gradients of specularly in some materials.

One way to address these challenges is to characterize the spatially-varying vector of illumination at each point of the sample, as in previous work [Nam et al. 2016], and repeat the fitting process done at the microscale. Other methods [Deschaintre et al. 2020; Ngan

and Durand 2006; Ruiters et al. 2013; Steinhausen et al. 2014] estimate the material properties in a localized area and extended it to larger samples. We follow the latter approach and build upon the work of Rodriguez-Pardo and Garces [2021], which leverage images of the material taken under different illumination conditions to train a neural network that propagates individual maps. We use a photometric dataset obtained with the micro setup, the spatially-varying parameters of the optimized micro SVBSDF (as described in Section 6), and a guidance image that will serve as a reference for the mesoscale propagation. The guidance image is taken with the mid-range camera and diffuse lighting. We introduce some modifications to the original architecture and training process to account for transferring multiple maps at the same time and improving overall accuracy described below:

*Training Data.* To improve the model’s generalization, we increase the training data by using a large photometric dataset and adding random Gaussian blurs. We further remove affine distortions and train the model for more iterations and larger batch size.

*Large Spatially-varying Materials.* We can account for materials that have variable appearances across multiple micro captures by training the propagation network with as many micro SVBSDFs as needed. See Figure 10 (a) for an example that required two microscale captures. To avoid color inconsistencies, we disable color-invariance data augmentation when multiple micro SVBSDFs are available.

*Architecture and Losses.* We observed that the original architecture was unable to effectively transfer the large number of parameters in our material model. Therefore, inspired by recent work [Janner et al. 2017; Rodriguez-Pardo and Garces 2022], we introduce a separate decoder for every map we aim to transfer. We introduced residual connections [Diakogiannis et al. 2020; He et al. 2016] into every layer, and substituted Batch Normalization [Ioffe and Szegedy 2015] for Group Normalization [Wu and He 2018]. When multiple microscale captures are available, we used a model with additional filters in every layer to better learn from these extended datasets. This design generates sharper texture maps that preserve the statistics and appearance of each fitted microscale map more accurately. We also added a multi-channel perceptual component [Chambon et al. 2021] to the loss function, which is a powerful regularizer for texture synthesis. Figure 10 (b-c) shows the difference with respect to the previous approach. For detailed implementation details and results, see the supplementary material.

## 7.1 Evaluation of the SVBSDF Estimation

We evaluate our model and fits using images taken with the capture device, as well as other validation scenes built for this purpose. We further compare several versions of our material model.

*Material Model.* We evaluate different versions of the material model in Figure 11. These include the baseline initialization, an isotropic version, one without a specular tint, and one with a limited range of the index of refraction (IOR). This last variation is equivalent to the Disney 2012 model [Burley and Studios 2012], which uses the Schlick approximation instead of the IOR. We present the results in terms of SSIM, averaging the error across views and include several examples. Overall, the smaller error is always obtained using the full model, followed by the model with the IOR clamped and the model

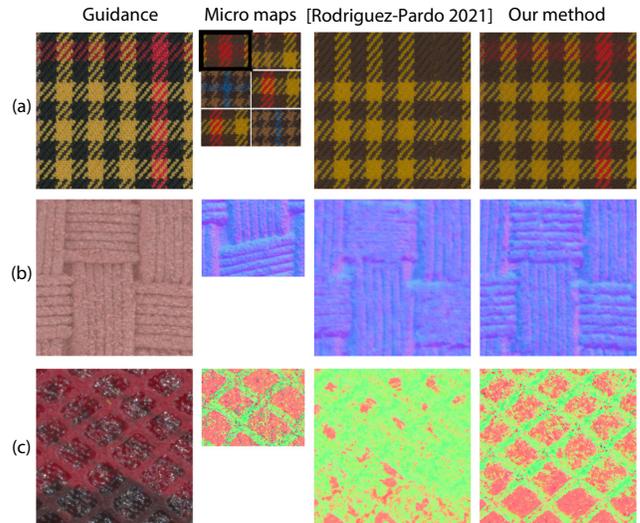


Fig. 10. Comparison with the method of Rodriguez-Pardo and Garces [2021] for mesoscale maps propagation. The first row (a) showcases an example where multiple captures at the microscale were needed to cover the spatially-varying albedo of the material. (b) and (c) required a single capture, we show the result of normals and tangents maps compared with previous work.

without specular tint. Fabrics that require specular tint are the satins (SATIN-CREME, SATIN-RED), the velvet (VELOUR-GREEN) and the red jacquard (JACQUARD-RED). These fabrics, although not metallic, reflect a colored specular because of their internal microstructure. The model with the IOR clamped works reasonably well for diffuse materials that do not present intense specular highlights such as plain wovens (TRICOT), single jerseys (JERSEY-YELLOW) or leathers (LEATHER-RHOMBUS, LEATHER-LIZARD). Anisotropy is necessary in most materials, except for the fleece (FLEECE), the leathers, and some hairy knits (JERSEY-HEATHER, JERSEY-YELLOW). This is expected, as these materials lack the structured internal organization present in woven fabrics, which causes anisotropic reflections when light strikes the main direction of the weaving pattern. Finally, we can see that the baseline initialization provides the worse results in general, only performing well for very diffuse fabrics such as the heather knit (JERSEY-HEATHER) and the fleece (FLEECE).

*Physical Validation Setups.* We build five distinct physical setups to validate different aspects of the quality of the reconstruction qualitatively, illustrated in Figure 12. The S-validation setups aim to assess the global positioning and integrated appearance of the specular lobes (S-validation/Specular) and sheen-like effects (S-validation/Sheen). It is composed of a solid S-shaped PU structure 21 cm wide, illuminated by point lights at different positions. The Hanging and Stretch scenes are samples of fabrics of 50x50 cm wide hung with magnetic pins illuminated by two area lights. The Bias scene is a crop of 10x10 cm over an approximate black and white cardboard illuminated by the diffuse lighting of our optical device. All the images have been captured with the same camera sensor, a Canon EOS 5DSR.



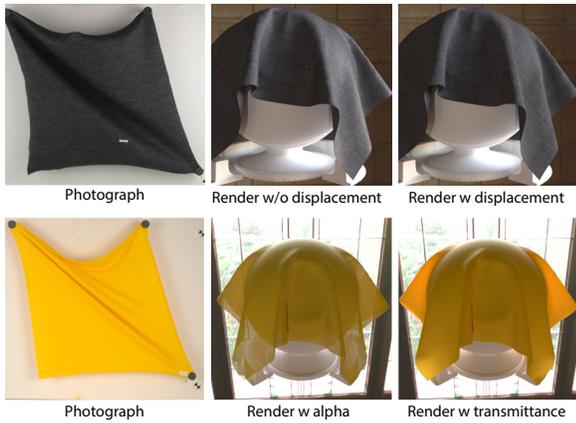


Fig. 13. Top: validation real image, and render of a diffuse fabric without and with displacement. Bottom: validation real image, and render of a fabric with transmittance by only using a naive alpha approximation. Using displacements and real transmittance maps are important for a realistic appearance in textile materials.

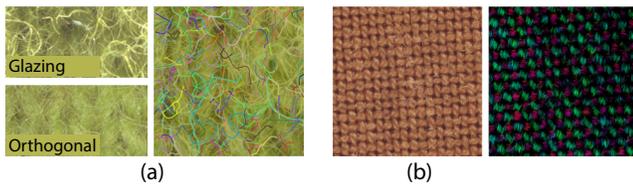


Fig. 14. (a) Image-based fly-out fibers detection. (b) Result of the fiber twist estimation. This microscale information can be applied to design specific fabric shading models such as the one shown in Figure 15.



Fig. 15. Example of microstructure influence in final appearance. Explicit yarns are modeled and twist is used as a property to control the roughness and orientation of the specular lobe in a real-time shader. Left: detailed views. Middle: high irregular twist. Right: no twist, characteristic of silk and satin fabrics.

the path of the minimum gradient. The full implementation details are in the supplementary material. Some results in Figure 14.

## 9 CONCLUSIONS AND FUTURE WORK

In this paper we have presented an optical acquisition system to digitize flat samples of anisotropic and translucent materials using a microfacet SVBSDF. Our device can capture the material at two

scales, which we leverage to propagate the SVBSDF parameters obtained from microscale to the mesoscale by using a neural network trained per material using data from the capture device.

We have validated the accuracy of our digitizations using a variety of real-world scenes with diverse illuminations, gathering a dataset of captured textiles and leathers that we will release to be used by the community.

Our current pipeline is quite robust, covering a reasonable range of dielectric materials. However, the SVBSDF model does not support directional specular transmittance and metallics. The former could be optimized by leveraging the individual contribution of each of the four directional LEDs that provide backlighting. While most physically-based rendering systems make a clear distinction between metallic and dielectric surfaces, it is common to use tinted specular to emulate the metallic effect [Li et al. 2019], as we have done. Even with constrained inputs, previous metallicness estimation methods are not very accurate [Collins et al. 2022]. Our optimization model could include spatially varying metallicness as an additional parameter, but the added complexity will require new strategies to keep it robust. Our database of materials will have to be extended to include enough samples of metallic and transparent objects. Our polar cameras have been used only for validation and calibration purposes, but we devised them to estimate view-dependent effects, such as inter-yarn occlusion in fabrics. Our current optimization is limited to an orthogonal view, missing information to estimate additional parameters such as sheen, and other scattering effects at grazing angles. There is also the possibility of analyzing the mesoscale microfacet geometry term, and compare it to occlusion produced by microgeometry displaced according to our micro fit process. Furthermore, our differentiable rendering approach has some limitations: the microstructure 3D surface is not reconstructed, thus being unable to generate cast shadows for incoming light at grazing angles. More sophisticated inverse rendering approaches could be explored [Jakob et al. 2022], either with surface or volumetric microflake representations, which could reproduce inter-reflections, transmission and occlusion at micro scale level [Jin et al. 2022]. Finally, as we have shown with our fiber twist and fly-out characterization, we think that our system could be specially relevant to advanced material models [Jin et al. 2022; Montazeri et al. 2020], which require statistical characterization of spatially-varying microscopic properties.

## ACKNOWLEDGMENTS

We wish to thank the reviewers for their helpful comments. We also thank Javier Fabre for help with the renders. Luis Romero for overall support with the 3D scenes and the validation setups. Sofia Dominguez for her help capturing data. We would like to thank Carlos Heras, Iñigo Salinas, Raul Alcain, Carlos Aliaga, and Enrique Pellejer for their help with hardware prototyping. Elena Garces was partially supported by a Juan de la Cierva - Incorporacion Fellowship (IJC2020-044192-I). This publication is part of the project TaiLOR, CPP2021-008842 funded by MCIN/AEI/10.13039/501100011033 and the NextGenerationEU / PRTR programs.

## REFERENCES

2016. X-RITE: Tac7. <https://www.xrite.com/categories/appearance/total-appearance-capture-ecosystem/tac7>. Accessed: 2022-04-28.
- Miika Aittala, Tim Weyrich, Jaakko Lehtinen, et al. 2015. Two-shot SVBRDF capture for stationary materials. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 110–1.
- Raúl Alcain, Carlos Heras, Iñigo Salinas, Jorge López, and Carlos Aliaga. 2019. Microscale Optical Capture System for Digital Fabric Recreation. In *Photoptics*. 114–119.
- Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. 2017. An appearance model for textile fibers. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 35–45.
- Michael Ashikhmin and Peter Shirley. 2000. An anisotropic phong brdf model. *Journal of Graphics Tools* 5, 2 (2000), 25–32.
- Brent Burley. 2015. Extending the Disney BRDF to a BSDF with integrated subsurface scattering. *Physically Based Shading in Theory and Practice SIGGRAPH Course* (2015).
- Brent Burley and Walt Disney Animation Studios. 2012. Physically-based shading at disney. In *ACM SIGGRAPH*, Vol. 2012. vol. 2012, 1–7.
- Carlos Castillo, Jorge López-Moreno, and Carlos Aliaga. 2019. Recent Advances in Fabric Appearance Reproduction. *Computers & Graphics* 84, C (Nov 2019), 103–121.
- Thomas Chambon, Eric Heitz, and Laurent Belcour. 2021. Passing Multi-Channel Material Textures to a 3-Channel Loss. In *ACM SIGGRAPH 2021 Talks*. 1–2.
- Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu, Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, et al. 2022. Abo: Dataset and benchmarks for real-world 3d object understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21126–21136.
- Robert L Cook and Kenneth E Torrance. 1981. A reflectance model for computer graphics. *ACM SIGGRAPH Computer Graphics* 15, 3 (1981), 307–316.
- Kristin J Dana. 2001. BRDF/BTF measurement device. In *Proceedings Eighth IEEE International Conference on Computer Vision, ICCV 2001*, Vol. 2. IEEE, 460–466.
- Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.
- Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2019. Flexible svbrdf capture with a multi-image deep network. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 1–13.
- Valentin Deschaintre, George Drettakis, and Adrien Bousseau. 2020. Guided Fine-Tuning for Large-Scale Material Transfer. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 91–105.
- Foivos I Diakogiannis, François Waldner, Peter Caccetta, and Chen Wu. 2020. ResUNet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing* 162 (2020), 94–114.
- Kirill Dmitriev and Evgeny Makarov. 2011. Generating Displacement from Normal Map for Use in 3D Games. In *ACM SIGGRAPH 2011 Talks* (Vancouver, British Columbia, Canada) (*SIGGRAPH '11*). Association for Computing Machinery, New York, NY, USA, Article 9, 1 pages.
- Yue Dong. 2019. Deep appearance modeling: A survey. *Visual Informatics* 3, 2 (2019), 59–68.
- Yue Dong, Jiaping Wang, Xin Tong, John Snyder, Yanxiang Lan, Moshe Ben-Ezra, and Baining Guo. 2010. Manifold bootstrapping for SVBRDF capture. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 1–10.
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P Greenberg. 2015. Predicting appearance from measured microgeometry of metal surfaces. *ACM Transactions on Graphics (TOG)* 35, 1 (2015), 1–13.
- Julie Dorsey, Holly Rushmeier, and François Sillion. 2010. *Digital modeling of material appearance*. Elsevier.
- Jonathan Dupuy and Wenzel Jakob. 2018. An adaptive parameterization for efficient material acquisition and rendering. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–14.
- Jiri Filip and Radomír Vávra. 2014. Template-based sampling of anisotropic BRDFs. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 91–99.
- Roland W. Fleming and Heinrich H. Bühlhoff. 2005. Low-Level Image Cues in the Perception of Translucent Materials. *ACM Transactions on Applied Perception* 2, 3 (jul 2005), 346–382.
- Jeppe Revall Frisvad, Soeren A Jensen, Jonas Skovlund Madsen, António Correia, Li Yang, Søren Kimmner Schou Gregersen, Youri Meuret, and P-E Hansen. 2020. Survey of models for acquiring the optical properties of translucent materials. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 729–755.
- David Geisler-Moroder and Arne Dür. 2010. A new ward brdf model with bounded albedo. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 1391–1398.
- Paul Graham, Borom Tunwattanapong, Jay Busch, Xueming Yu, Andrew Jones, Paul Debevec, and Abhijeet Ghosh. 2013. Measurement-based synthesis of facial microgeometry. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 335–344.
- Darya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. 2016. BRDF representation and acquisition. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 625–650.
- Giuseppe Claudio Guarnera, Peter Hall, Alain Chesnais, and Mashhuda Glencross. 2017. Woven fabric model creation from a single image. *ACM Transactions on Graphics (TOG)* 36, 5 (2017), 1–13.
- MG Guerra, SS Gregersen, Jeppe Revall Frisvad, Leonardo De Chiffre, F Lavecchia, and LM Galantucci. 2019. Measurement of polymers with 3D optical scanners: evaluation of the subsurface scattering effect through five miniature step gauges. *Measurement Science and Technology* 31, 1 (2019), 015010.
- Jie Guo, Shuichang Lai, Chengzhi Tao, Yuelong Cai, Lei Wang, Yanwen Guo, and Ling-Qi Yan. 2021. Highlight-aware two-stream network for single-image SVBRDF acquisition. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14.
- Yu Guo, Cameron Smith, Miloš Hašan, Kalyan Sunkavalli, and Shuang Zhao. 2020. MaterialGAN: reflectance capture using a generative SVBRDF model. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–13.
- Vlastimil Havran, Jan Hošek, Šárka Němcová, Jiří Čáp, and Jiří Bittner. 2017. Light-drum—Portable light stage for accurate BTF measurement on site. *Sensors* 17, 3 (2017), 423.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 770–778.
- Philipp Hertzner, Valentin Deschaintre, Niloy J Mitra, and Tobias Ritschel. 2021. Generative Modelling of BRDF Textures from Flash Images. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–13.
- Michael Holroyd, Jason Lawrence, and Todd Zickler. 2010. A coaxial optical scanner for synchronous acquisition of 3D geometry and surface reflectance. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 1–12.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*. PMLR, 448–456.
- Tomáš Iser, Tobias Rittig, Emilie Nogué, Thomas Klaus Nindel, and Alexander Wilkie. 2022. Affordable Spectral Measurements of Translucent Materials. *ACM Transactions on Graphics (TOG)* 41, 6, Article 199 (Nov 2022), 13 pages.
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. <https://mitsuba-renderer.org>.
- Michael Janner, Jiajun Wu, Tejas D Kulkarni, Ilker Yildirim, and Josh Tenenbaum. 2017. Self-supervised intrinsic image decomposition. *Advances in Neural Information Processing Systems* 30 (2017).
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A Practical Model for Subsurface Light Transport (*SIGGRAPH '01*). Association for Computing Machinery, New York, NY, USA, 511–518.
- Jorge Jimenez and Diego Gutierrez. 2010. *GPU Pro: Advanced Rendering Techniques*. AK Peters Ltd., Chapter Screen-Space Subsurface Scattering, 335–351.
- Wenhua Jin, Beibei Wang, Miloš Hašan, Yu Guo, Steve Marschner, and Ling-Qi Yan. 2022. Woven Fabric Capture from a Single Photo. In *Proceedings of SIGGRAPH Asia 2022*.
- Pramoek Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. *ACM Transactions on Graphics (TOG)* 35, 1 (2015), 1–1.
- Guillaume Lavoué, Nicolas Bonneel, Jean-Philippe Farrugia, and Cyril Soler. 2021. Perceptual quality of BRDF approximations: dataset and metrics. *Computer Graphics Forum* 40, 2 (2021), 327–338. <https://doi.org/10.1111/cgf.142636>
- Bo Li, Jie Feng, and Bingfeng Zhou. 2019. A SVBRDF Modeling Pipeline using Pixel Clustering. <https://arxiv.org/abs/1912.00321>
- Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. 2018. Materials for masses: SVBRDF acquisition with a single mobile phone image. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 72–87.
- Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques*. 183–194.
- Xiaohu Ma, Kaizhang Kang, Ruisheng Zhu, Hongzhi Wu, and Kun Zhou. 2021. Free-form scanning of non-planar appearance with neural trace photography. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–13.
- Stephen R Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light scattering from human hair fibers. *ACM Transactions on Graphics (TOG)* 22, 3 (2003), 780–791.
- Stephen R Marschner, Stephen H Westin, Adam Arbree, and Jonathan T Moon. 2005. Measuring and modeling the appearance of finished wood. In *ACM SIGGRAPH 2005 Papers*. 727–734.
- Rosalie Martin, Arthur Roullier, Romain Rouffet, Adrien Kaiser, and Tamy Boubekeur. 2022. MaterIA: Single Image High-Resolution Material Capture in the Wild. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 163–177.
- Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2003a. A data-driven reflectance model. *ACM Transactions on Graphics (TOG)* 22, 3 (2003), 759–769.

- Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. 2003b. Efficient isotropic BRDF measurement. In *Proceedings of the 14th Eurographics Workshop on Rendering*, 241–247.
- Sebastian Merzbach, Max Hermann, Martin Rump, and Reinhard Klein. 2019. Learned fitting of spatially varying brdfs. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 193–205.
- Zahra Montazeri, Søren B. Gammelmark, Shuang Zhao, and Henrik Wann Jensen. 2020. A Practical Ply-Based Appearance Model of Woven Fabrics. *ACM Trans. Graph.* 39, 6, Article 251 (nov 2020), 13 pages. <https://doi.org/10.1145/3414685.3417777>
- Giljoo Nam, Joo Ho Lee, Hongzhi Wu, Diego Gutierrez, and Min H Kim. 2016. Simultaneous acquisition of microscale reflectance and normals. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 185–1.
- Addy Ngan and Frédo Durand. 2006. Statistical acquisition of texture appearance. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques*. 31–40.
- FE Nicodemus, JC Richmond, JJ Hsia, IW Ginsberg, and T Limperis. 1992. Geometrical considerations and nomenclature for reflectance. In *Radiometry*. 94–145.
- Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–11.
- Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo. 2011. Pocket reflectometry. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–10.
- Carlos Rodriguez-Pardo and Elena Garces. 2021. Neural Photometry-guided Visual Attribute Transfer. *IEEE Transactions on Visualization and Computer Graphics* (2021), 1–1.
- Carlos Rodriguez-Pardo and Elena Garces. 2022. SeamlessGAN: Self-Supervised Synthesis of Tileable Texture Maps. *IEEE Transactions on Visualization and Computer Graphics* (2022).
- Roland Ruiters, Christopher Schwartz, and Reinhard Klein. 2013. Example-based Interpolation and Synthesis of Bidirectional Texture Functions. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 361–370.
- Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Transactions on Graphics (TOG)* 32, 2 (2013), 1–12.
- Christopher Schwartz, Ralf Sarlette, Michael Weinmann, Martin Rump, and Reinhard Klein. 2014. Design and implementation of practical bidirectional texture function measurement devices focusing on the developments at the University of Bonn. *Sensors* 14, 5 (2014), 7753–7819.
- Liang Shi, Beichen Li, Miloš Hašan, Kalyan Sunkavalli, Tamy Boubekeur, Radomir Mech, and Wojciech Matusik. 2020. Match: differentiable material graphs for procedural material capture. (2020).
- Heinz Christian Steinhausen, Dennis den Brok, Matthias B Hullin, and Reinhard Klein. 2014. Acquiring bidirectional texture functions for large-scale material samples. (2014).
- Kenneth E Torrance and Ephraim M Sparrow. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America* 57, 9 (1967), 1105–1114.
- Giuseppe Vecchio, Simone Palazzo, and Concetto Spampinato. 2021. SurfaceNet: Adversarial SVBRDF Estimation from a Single Image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12840–12848.
- Raquel Vidaurre, Dan Casas, Elena Garces, and Jorge Lopez-Moreno. 2019. Brdf estimation of complex materials with nested learning. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 1347–1356.
- Chun-Po Wang, Noah Snaveley, and Steve Marschner. 2011. Estimating dual-scale properties of glossy surfaces from step-edge lighting. In *Proceedings of the 2011 SIGGRAPH Asia Conference*. 1–12.
- D Rod White, Peter Saunders, Stuart J Bonsey, John van de Ven, and Hamish Edgar. 1998. Reflectometer for measuring the bidirectional reflectance of rough surfaces. *Applied Optics* 37, 16 (1998), 3450–3454.
- Robert J Woodham. 1980. Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1 (1980), 139–144.
- Yuxin Wu and Kaiming He. 2018. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 3–19.
- Bei Xiao, Bruce Walter, Ioannis Gkioulekas, Todd Zickler, Edward Adelson, and Kavita Bala. 2014. Looking against the light: How perception of translucency depends on lighting direction. *Journal of Vision* 14, 3 (03 2014), 17–17.
- Bei Xiao, Shuang Zhao, Ioannis Gkioulekas, Wenyan Bi, and Kavita Bala. 2020. Effect of geometric sharpness on translucent material perception. *Journal of Vision* 20, 7 (07 2020), 10–10.
- Jiyang Yu, Zexiang Xu, Matteo Mannino, Henrik Wann Jensen, and Ravi Ramamoorthi. 2016. Sparse sampling for image-based SVBRDF acquisition. In *Proceedings of the Eurographics 2016 Workshop on Material Appearance Modeling*. 19–22.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building volumetric appearance models of fabric using micro CT imaging. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–10.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-aware synthesis for predictive woven fabric appearance. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.
- Shuang Zhao, Fujun Luan, and Kavita Bala. 2016. Fitting procedural yarn models for realistic cloth rendering. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–11.
- Xilong Zhou and Nima Khademi Kalantari. 2021. Adversarial Single-Image SVBRDF Estimation with Hybrid Training. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 315–325.
- Xilong Zhou and Nima Khademi Kalantari. 2022. Look-Ahead Training with Learned Reflectance Loss for Single-Image SVBRDF Estimation. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–12.



Fig. 16. Photos and renders of several materials of our dataset. The complete set of results that includes the 36 materials with the full set of validation scenes is included in the supplementary material. We show, from left to right, photos and renders of our digitized materials on two views with directional lighting of the capture machine at the microscale, the Stretch scene, the Bias scene, and the S-validation/Specular scene. The majority of the materials seem accurately reproduced at the microscale (first two columns), as the basecolor was captured under the same illumination conditions, and directional lighting is easy to reproduce. Replicating the exact illumination conditions and camera settings in the other validation scenes is more challenging, although the results are satisfactory in most cases.