

Scattering-aware Texture Reproduction for 3D Printing: Supplemental Document – Measurement

OSKAR ELEK*, Charles University, Prague, Czech Republic
DENIS SUMIN*, Max Planck Institute for Informatics, Saarbrücken, Germany
RAN ZHANG, Institute for Science and Technology Austria
TIM WEYRICH, University College London, United Kingdom
KAROL MYSZKOWSKI, Max Planck Institute for Informatics, Saarbrücken, Germany
BERND BICKEL, Institute for Science and Technology Austria
ALEXANDER WILKIE, Charles University, Prague, Czech Republic
JAROSLAV KŘIVÁNEK, Charles University, Prague, Czech Republic

ACM Reference format:

Oskar Elek*, Denis Sumin*, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Krivánek. 2017. Scattering-aware Texture Reproduction for 3D Printing: Supplemental Document – Measurement. *ACM Trans. Graph.* 36, 6, Article 241 (November 2017), 11 pages. DOI: 10.1145/3130800.3130890

1 INTRODUCTION

This document is intended to provide a full, detailed exposition of the material measurement part of our pipeline. Please refer to the main article for the description of the remaining pipeline components and how the measured data are utilized in them.

The document is structured as follows. After stating the measurement problem (Sec. 2) we explain our design and position it within the most related works in the area (Sec. 3). Our setup and measurement procedure are then detailed in Secs. 4 and 5. The resulting measured data are listed and validated in Sec. 6. We then wrap up with a short discussion (Sec. 7).

2 PROBLEM STATEMENT

The primary objective of our method is the reproduction of a target appearance given by a user-specified color texture. We reproduce this appearance using the photo-polymer-based working materials of the Stratasys J750 printer. Without loss of generality we define the discrete space of these materials as $\mathcal{M}' \equiv \{C, M, Y, K, W\}$, standing for cyan, magenta, yellow, black, and white respectively. In order to predict the appearance of an object fabricated with these materials we ultimately need to obtain their intrinsic volumetric optical parameters defined in the continuous space $\mathcal{V} \equiv (\sigma_t, \alpha, g, \eta)$, where

- $\sigma_t \in [0, \infty)$ is the extinction coefficient,
- $\alpha \in [0, 1]$ is the single-scattering albedo,

*Oskar Elek and Denis Sumin share the first authorship of this work. This is a supplemental material.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 0730-0301/2017/11-ART241 \$15.00
DOI: 10.1145/3130800.3130890

- $g \in (-1, 1)$ is the scattering anisotropy, and
- $\eta \in (1, \infty)$ is the refractive index (IOR) of the material.

For dielectric materials, the refractive index η can be determined directly via optical ellipsometry. Using this method, we have measured the black material's IOR to be $\eta_K = 1.48$ and the IORs of the remaining four materials $\eta_{\{C, M, Y, W\}} = 1.53$.

For the remaining three parameters— σ_t , α and g —a direct measurement is unfortunately complicated by the fact that these parameters relate to a light interacting with a single particle of the material. While dilution can enable approximate isolation of single light-particle interactions for some fluid and soluble solid media [Narasimhan et al. 2006], for most solids including our printing materials this is currently infeasible. We therefore have to rely on *indirect* measurements (also known as inverse rendering) of these parameters, i.e., measurements based on observations of many cumulative light-particle interactions.

Indirect measurements can be performed in isolation or jointly. To mention a few instances of the first approach, the extinction coefficient σ_t can be inferred via Beer-Lambert-Bouguer law, and the scattering albedo α can be directly related to the bulk reflectance (i.e., color) of a sufficiently large piece of material under the assumption of diffusive conditions. However, the scattering anisotropy g is more difficult to obtain in this fashion. Moreover, an isolated inference of the parameters can lead to biased measurements, since the errors of each partial measurement are unlikely to cancel each other out. As described in Sec. 3, we therefore opt for the latter, joint approach of obtaining all three parameters by a single measurement procedure.

3 DESIGN PRINCIPLES

Our primary design objective has been to provide a pragmatic way to measure the photo-polymers or similar materials used by the current 3D fabrication technologies. To this end we followed these general design principles:

- *Simplicity.* Avoid the reliance on overly advanced or specific theory to enable easy diagnostics. Avoid complicated design to decrease the probability of introducing errors and hence improve measurement repeatability.
- *Accessibility.* Use the smallest number of restrictive or hard-to-achieve design assumptions. Ensure robustness in case of

weak violation of the assumptions. Provide a thorough process documentation to enable the reproducibility of the setup.

- *Affordability*: Reduce the number of specialized or expensive hardware components. Use open-source or free software whenever possible. Minimize operator time in exchange for machine time.

First we briefly review the landscape of the existing volume acquisition methods and then provide rationale for developing a new approach in the perspective of the above principles.

3.1 State of the Art in Volumetric Media Acquisition

Gkioulekas et al. [2013] provide an excellent overview of the currently available approaches. We therefore review only the works most relevant for our context, i.e., the acquisition of multiple-scattering solid materials. All the mentioned works rely on indirect inference of the material parameters.

Both Hašan et al. [2010] and Dong et al. [2010] focus on reproducing spatially varying BSSRDFs of translucent materials, and as such need to measure their properties in some way. Hašan et al. acquire the diffuse albedo and scattering profile for both reflected and transmitted light through the measured material slab. Since their reproduction is data-driven, they continue to operate in the space of these profiles, and therefore do not convert to the primary material parameters which we require.

The reproduction approach of Dong et al. is related in its aim to the one of Hašan et al. – however, they do perform a forward simulation step in their pipeline. Nevertheless, both their dipole-based acquisition as well as the simulation itself rely on the assumption of diffusive transport, operating with only the *reduced* extinction coefficient σ_t' and scattering albedo α' . One of the known limitations of the diffusion dipole is its decreased accuracy for media with significant absorption [d'Eon and Irving 2011], which—contrary to Dong et al. who only operate with high-albedo materials—is a significant limitation for us. In addition, the assumption of diffusive transport that allows the use of reduced material parameters is not feasible for us, since sharp material transitions common for our target appearances violate this assumption.

The work of Papas et al. [2013] also describes an optimization-driven appearance reproduction framework, however with the aim to reproduce an arbitrary homogeneous material by *continuously* mixing silicone pigments. Similar to Hašan et al. [2010] they acquire the diffuse albedos and scattering profiles of the pigments and the target material, but just as Dong et al. [2010] they rely on the *reduced* material parameters obtained from the acquired data by fitting. Their method also becomes quite involved as significant refinement is required for adjusting the measured data to reproduce the absorptive nature of their materials with sufficient fidelity.

In contrast to the above works, Gkioulekas et al. [2013] focus purely on material acquisition. It is currently the most advanced method in the field with the ability to obtain the full set of material parameters, including an approximate shape of the material's phase function (not only its anisotropy g that we operate with). Their approach uses collimated laser light to capture reflective and transmissive profiles of the material, and a dictionary fit to optimize for

the parameters themselves. The approach is thoroughly theoretically justified, which on the other hand causes it to be very complex, plus it relies on accurately calibrated specialized equipment. It also requires solid samples to be molded into a glass containment cell, which is not very practical for hard polymer materials.

3.2 Discussion

As a general principle, in order to perform an indirect joint measurement of the volumetric optical parameters, a method needs to rely on some kind of signal being present in the environment that the sample optically interacts with. We will refer to it as the *modulation signal*. The parameters are then inferred from the *observed distortion* of this signal by following the first principles of volumetric light transport in a Monte Carlo simulation, and *fitting* the observation to the results of the simulation which samples material parameter space.

All the aforementioned approaches use the (spatial or directional) structure of the illuminant itself as a source of the modulation signal: Papas et al. [2013] use LEDs as approximate point sources, while the remaining works rely on collimated light from lasers (Dong et al. [2010] and Gkioulekas et al. [2013]) or a projector (Hašan et al. [2010]). Note that, with the exception of lasers, it is difficult to characterize or control light sources to fulfill the assumptions of the measurement, which then potentially limits the discriminability of the setup.

However, laser light is strictly monochromatic, while our method operates in the RGB colorspace (for perceptual reasons, but also for simplicity and efficiency sake). Identifying the number of wavelengths required for colorimetrically capturing the printing materials' unknown spectra with sufficient fidelity is unfortunately far from trivial. In addition, increasing the number of acquired wavelengths also increases the build costs, as well as the requirements on alignment accuracy (for illustration, both Dong et al. [2010] and Gkioulekas et al. [2013] only operate with three wavelengths). On the other hand, a broad-band RGB acquisition setup can always be upgraded to a spectral one through band-pass filters, if needed.

Conclusion. Based on the above discussion, we propose to use a source of the modulation signal different from the illuminant. This can either be a feature *embedded* in the material, or an *external* one.

The simplest feature embedded in the measured object itself would be a boundary (an edge) between two materials, for instance black (K) and white (W). Such an approach is appealing, as transitions between materials are a key feature in our target appearance space. However, this would mean that the fitting space would have six dimensions, instead of three when performed for a single isolated material. Given the cost of Monte Carlo simulation, a sufficiently dense sampling of this space would be hardly achievable. And since (as you can see later in Fig. 5) there is no clear ordering in the space of observations that would reflect the structure of the intrinsic parameter space, any adaptive sampling (e.g., based on binary search) is problematic.

We therefore opt for an approach based on an external source of the modulation signal. As we detail in the following section, we chose a simple 1D step edge printed on a reflective material placed

underneath a thin material sample. To see the results produced by our design, please refer to Sec. 6.

4 SETUP DESCRIPTION

We first explain the measurement setup on a conceptual level (Sec. 4.1), followed by the descriptions of the real and virtual setups we constructed to realize it (Secs. 4.2 and 4.3) Then, based on these setups, we provide a detailed exposition of the measurement procedure itself in Sec. 5.




Fig. 1. Sketch of our acquisition setup (not up scale). Please refer to Sec. 4.1 for details.

4.1 Concept

Based on the reasoning in Sec. 3 and extensive experimentation, we have arrived to a measurement setup design that is sketched in Fig. 1. Its six key components are as follows:

- *Sample*. A thin slab with the dimensions $40 \times 40 \times 0.5/1$ mm, 3D-printed with one of the materials from \mathcal{M} . Surface finish as smooth as possible.
- *Modulation signal source*. Reflective surface containing two regions (achromatic peak dark and light colors) separated by a straight vertical step edge. Placed under the sample.
- *Camera*. Placed about 1 m above the sample. Vertically aligned and laterally centered with the observed step edge (in camera coordinates).
- *Light*. Symmetric around the sample (rotationally or axially), diffuse, neutral illuminant. Positioned roughly along the bisector between the camera optical axis and the sample plane, to minimize direct reflection off the sample into the camera, but still provide sufficiently strong illumination.
- *Thin glass slip*. Intended to provide a flat, clear, specular interface on top of the sample to minimize bias in the acquisition.

- *Transparent gel*. Intended for suppressing the reflections from any potential residual roughness on the sample surface. Applied thinly and evenly between the sample/signal source and the sample/glass slip surfaces to minimize the gradients between their refractive indices and the atmospheric one.

For each material in \mathcal{M} the output of this setup is a *normalized, linear* RGB image, i.e., a 2D matrix of pixels with values $\in C \equiv [0, 1]^3$. From this image we crop an area of 1×1 cm with the modulation step edge running vertically through the center (see Fig. 2). Under the above assumption of spatially uniform illumination, this cropped sub-image can be considered vertically uniform as well, therefore we vertically average it for denoising purposes (but also to minimize the impact of any residual surface artifacts on the measurement). We denote the resulting 1D vector as the *edge profile* F , on which our measurement back-end further operates (more in Sec. 5.5).

Before moving onto describing our physical and virtual realizations of this setup, it first needs to be demonstrated that the proposed design offers adequate discriminability, i.e., that the resulting profiles are clearly distinguishable for different points in the volume parameters space \mathcal{V} to which we want to map the printing materials. In Fig. 3 we demonstrate that this is indeed so, for two different sample points from \mathcal{V} and both in terms of the profile shape and its gradient. This experiment also explains the motivation to use (at least) two sample thicknesses in the measurement, since the discriminability of the setup in different regions of \mathcal{V} varies with the thickness.

The choice of the sample thicknesses should in general follow a simple guideline: the samples should be thick enough so that the transport has a significant multiple-scattering component, but thin enough that the signal can still be distinguished in at least on of the RGB channels. That is, roughly in the order of $O(1/\sigma_t)$, or multiple *mean free paths* in other words.




Fig. 2. Step edges captured through the printing materials $\mathcal{M} = \{C, M, Y, K, W\}$ by our physical (left) and virtual (right) setups, at two different sample thicknesses. The images cover the area of 2×1 cm (which is twice as wide as the cropped version our fitting operates with, cf. Sec. 4.1). Please note that these images (as well as all other images of this kind in the document) use standard gamma compression for proper viewing.

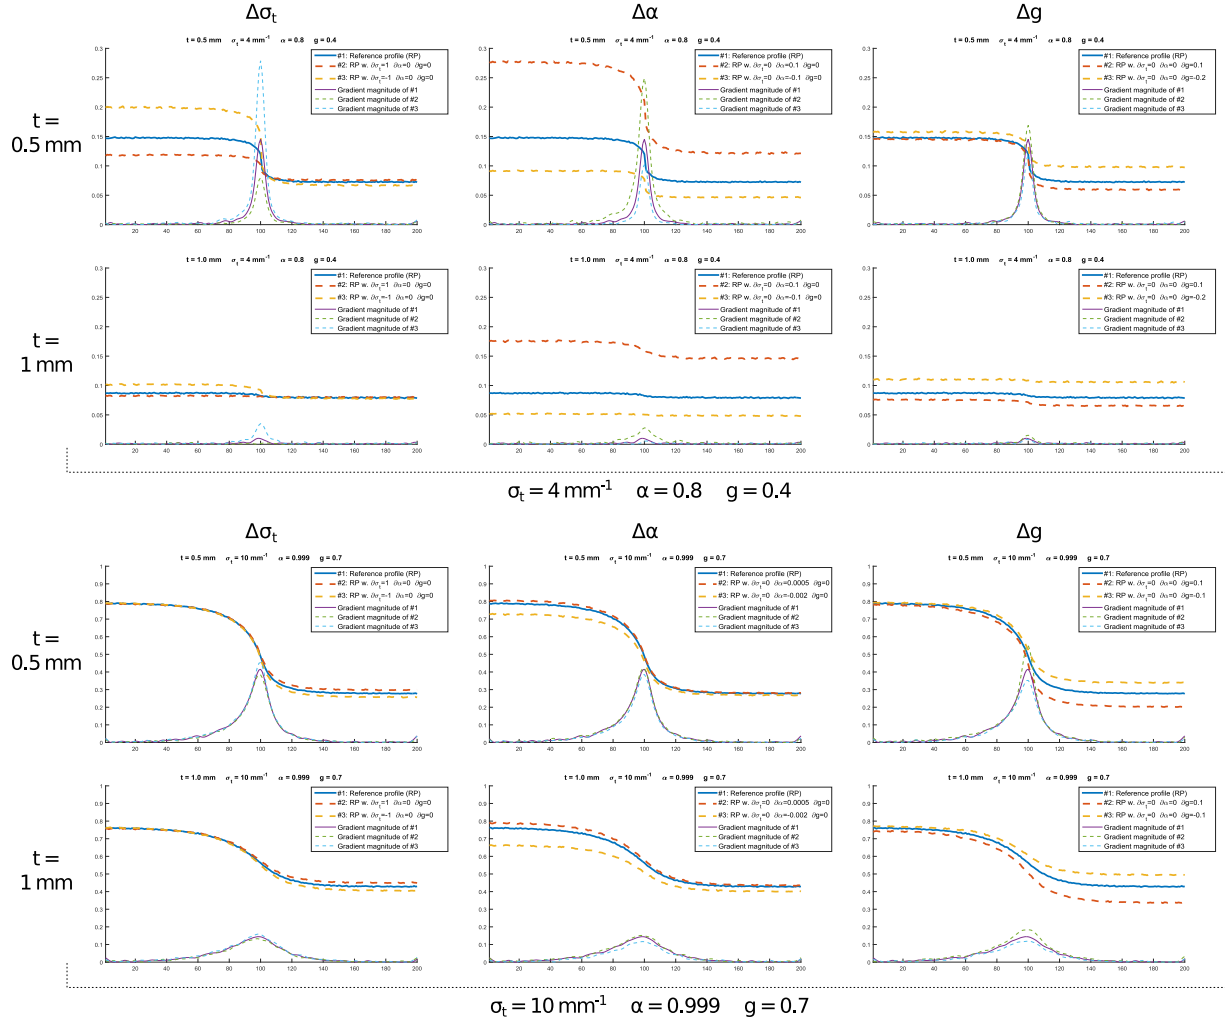


Fig. 3. Edge profiles simulated with our virtual setup (Sec. 4.3) for two different parameter sets $\in \mathcal{V}$, each for two sample thicknesses $t \in \{0.5, 1\}$ mm. Starting from a reference profile (solid blue), each plot shows two profiles resulting from a small change along one of the three dimensions of \mathcal{V} (dashed red and yellow). We also compare the respective profile gradient magnitudes, since these are taken into account by our fitting (Sec. 5.5). Details are listed in the corresponding plot legends. Please also note the different ranges in the two plot sets.

4.2 Physical Setup

In agreement with our design principles described in Sec. 3, we have made an effort to build the physical setup from commonly available, inexpensive parts:

- *Sample.* Ten samples (five materials in \mathcal{M} , each at thicknesses $t = \{0.5, 1\}$ mm) 3D-printed on the Stratasys J750 machine, using the standard Vero Opaque materials family. We added a small delta to the reference thicknesses to compensate for the material lost by polishing, see Sec. 5.3.
- *Modulation signal source.* Black and white checkerboard with 20 mm square size, printed with a laser printer on a Toughprint waterproof paper. We have opted for a pattern instead of a single step edge between two homogeneous regions to facilitate the

setup calibration (especially to calculate the pixel-to-mm ratio in the acquired images).

- *Camera.* Canon EOS 700D body with EFS 18–135 mm lens fixed at maximum zoom, mounted on a heavy tripod. Manually focused on the signal source (not the top of the sample), optical stabilization disabled. Exposed in manual mode at 1/20 s and f-11, triggered remotely via standard Canon utility software.
- *Light.* Two simple studio lamps with 55 W, 5500 K-equivalent fluorescent bulbs, covered with stock transmissive diffusers. The lights were positioned symmetrically above and below the sample (from the camera’s viewpoint), at a height to ensure an approximately $0^\circ/45^\circ$ acquisition geometry.
- *Thin glass slip.* Microscope slides $50 \times 24 \times 0.17$ mm made of BK7 borosilicate glass, $\eta = 1.52$.

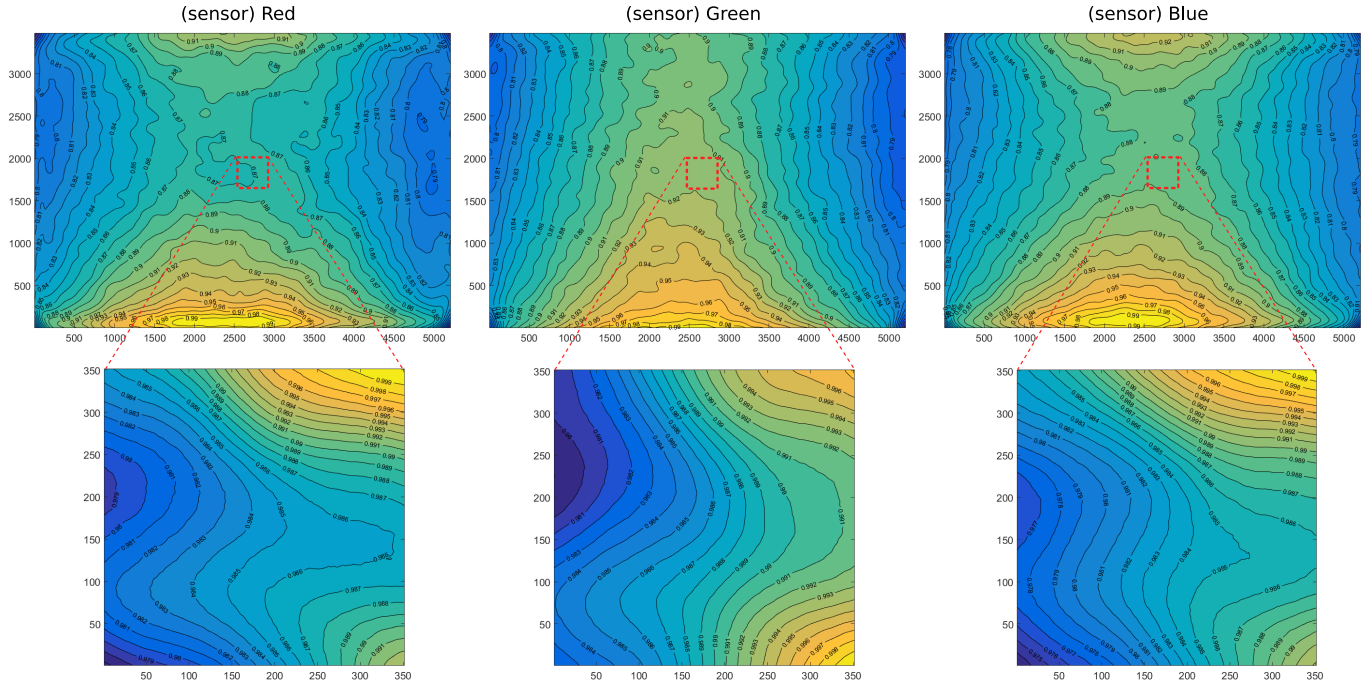


Fig. 4. False color map of the per-channel *relative* pixel intensity captured by our physical setup. The maps were convolved with an appropriate Gaussian filter to suppress noise. The iso-line spacing in the full images (**top**) is 1%, and 0.1% in the crops of the measurement region (**bottom**).

- *Transparent gel.* Clear Anagel ultrasound transmission gel with $\eta = 1.33$. The advantage of this type of gel is its low air bubbles content, which can pollute the measurement and are hard to notice by naked eye. Due to the gel’s low refractive index, a better alternative would be using an optical glue index-matched to the samples. This would however permanently bind the entire measurement ‘stack’ including the signal source, making the measurement alignment less reliable.

We have built the setup in a neutrally colored room (dark room is not necessary given the assumption of diffuse illumination) and eliminated all vibration and other light sources.

4.3 Virtual Setup

After building and testing the physical setup (Sec. 4.2) we have proceeded to modeling its virtual counterpart. Same as for the prediction step in the core algorithm (see the main text) we have used the open-source rendering package Mitsuba by Wenzel Jakob [Jakob 2010]. We modeled the scene manually (instead of using a 3D modeling software) from simple planar elements to guarantee proper dimensions, orientation and alignment of the optical interfaces in the measurement ‘stack’.

For the signal source, the RGB reflectances of the black and white squares are obtained during the physical setup calibration (see Sec. 5.1). We estimated the roughness β of both the polished samples ($\beta \approx 0.05$) and the signal source ($\beta \approx 0.15$) by visually comparing to rendered templates with different values of β (using Mitsuba’s roughdielectric BRDF with the GGX micro-facet distribution).

For the illuminant, we used a toroidal diffuse emitter. This configuration (from the perspective of the measurement region) has the same symmetry as the illumination in our physical setup, but leads to a more efficient simulation as more paths successfully find the source.

To ensure robust and unbiased results, we use (volumetric) path tracing as the integrator. We empirically determined that 250k samples per pixel yields an amount of noise comparable to our physical acquisition. We setup the virtual sensor to only render a single linear profile F matching the ones extracted from the physical setup (as explained in Sec. 4.1).

5 MEASUREMENT PROCEDURE

Having described our measurement setups, we now proceed to detail the measurement procedure itself. It consists of the following steps:

- setup calibration (Sec. 5.1),
- generation of material dictionary (Sec. 5.2),
- physical acquisition (Sec. 5.3),
- data processing and extraction (Sec. 5.4), and finally
- fitting-based measurement (Sec. 5.5).

The output of the procedure is an *optical characterization* of the material space \mathcal{M} , i.e., a set of parameters (σ_t, α, g) for every material $m \in \mathcal{M}$ and RGB channel.

5.1 Setup Calibration

We have extracted and developed all acquired raw images using the ImageJ [et al. 2016; Schindelin et al. 2015] and DCRaw [Coffin




Fig. 5. Visualization of 288 (about 0.8%) of the records in our material profile dictionary (see Sec. 5.2), for $t \in \{0.5, 1\}$ mm. Optical thickness σ_t varies horizontally, scattering albedo α vertically and the anisotropy g within each sub-plot. The vertical blue line in each plot signals the step edge position.

2016] open source tools. As a first step, we ensured a linear, neutral camera response. To verify linearity we used the Xrite Passport

color checker, specifically its six neutral patches with provided reference reflectance values. We have then calculated the scene's white

point using the brightest neutral patch (mean reference reflectance 0.956 with a high degree of spectral uniformity) and normalized all captured images by its value.

Second, to verify our assumption of spatial illumination uniformity over the measurement region of interest, we captured an image of a white office paper placed in the imaging plane. As Fig. 4 demonstrates, while the captured intensity varies significantly across the image plane (most prominently due to the lighting geometry and camera lens’ vignetting), the central measurement region is within a 1% average deviation. We have therefore made sure that all calibration readings and measurements themselves are taken in this region.

The above steps then allowed us to measure the signal source’s RGB reflectances to be (0.985, 0.972, 0.971) for the white squares and (0.024, 0.023, 0.026) for the black squares (averaged over an area of roughly 1×1 cm). These values are then used in the virtual setup as a prerequisite to the dictionary generation.

5.2 Material Dictionary Generation

With our virtual setup as a back-end (Sec. 4.3), we scripted the dictionary generation using Mitsuba’s Python interface. After several refinements we settled on using the following sampling of the parameter space \mathcal{V} :

$$\begin{aligned} \sigma_t &\in \{1.0, 1.25, 1.5, 1.75, 2.0, 2.25, 2.5, 2.75, 3.0, \\ &3.25, 3.5, 3.75, 4.0, 4.5, 5.0, 5.5, 6.0, 6.5, 7.0, \\ &7.5, 8.0, 9.0, 10.0, 11.0, 12.0, 13.0, 14.0, 15.0, \\ &16.0, 17.0, 18.0, 19.0, 20.0, 22.0, 24.0, 26.0, \\ &28.0, 30.00\} \text{ mm}^{-1} \\ \alpha &\in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, \\ &0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.875, 0.9, 0.915, 0.92, 0.93, \\ &0.94, 0.95, 0.96, 0.97, 0.98, 0.99, 0.993, 0.995, 0.996, 0.997, \\ &0.998, 0.999, 0.9991, 0.9993, 0.9995, 0.9997, 0.9999, \\ &0.99993, 0.99999\} \\ g &\in \{0.0, 0.1, 0.2, 0.3, 0.4, 0.45, 0.5, 0.55, 0.6, 0.7, 0.8\} \\ t &\in \{0.5, 1.0\} \text{ mm} \end{aligned}$$

Generating the resulting 34.2k-record dictionary took approximately two days on a 500-core CPU cluster. You can find a visualization of its small subset in Fig. 5.

5.3 Acquisition

Samples preparation. After the set of 10 samples had been 3D-printed and cleaned up, we followed by polishing them. This is necessary since the printer creates small but visible structures on a printout’s surface. We therefore manually polished them in a two-step procedure: with a grade-1000 sandpaper until reaching an even matte finish, then with an electric drill using a soft bit coated with plastic polishing paste, until achieving a glossy finish.

Since the heat buildup from the polishing can cause small bending of the samples, we flatten them by briefly placing each in boiling water, then for about 30 s between two flat heavy objects. At this temperature the materials become pliable but do not lose structural integrity or change optical properties.

Setup initialization. We turned on both the lights and camera at least 30 minutes prior to the acquisition, to provide a warm-in period for the components. We then verified the camera is aligned with the imaging plane in all three dimensions, and that the signal source is in focus. This was followed by taking several mockup images and then an image of the white patch of the Xrite color checker, to obtain the scene white point for the current run (since, as we found out, the white point can slightly change between consecutive runs).

Capturing procedure. For capturing each of the ten samples we have followed the same procedure:

- Wipe the signal source clean. Then apply a small amount of the clear gel onto its surface.
- Press the sample firmly until only a minimal amount of gel is left underneath.
- Apply a small amount of the gel on top the sample.
- Press a blank glass slide flatly into the gel, again pushing as much as possible out at the sides. Take care not to contaminate the top surface by the gel.
- Visually inspect the sample for any air bubbles remaining in the gel. If found, repeat the entire procedure for that sample.
- Capture two images of the sample, slightly shifting the entire stack (but not the signal source) in between to decrease the likelihood of the data being contaminated by any kind of remaining surface artifacts.

5.4 Data Extraction

After developing the acquired raw images with ImageJ and DCRaw (using the white point obtained during the calibration, Sec. 5.1), *all further processing* is done in Matlab. We mark the edge location manually in each image, and if necessary, we rectify them automatically using Radon transform [Radon 1986] so that the signal edge appears vertical. After that the measurement region is extracted and averaged as described in Sec. 4.1 to yield a single edge profile F per material sample.

5.5 Fitting

As the last step of our measurement pipeline, the fitting procedure finds the *most similar* simulated edge profile F_S for each profile F_A acquired by our physical setup. This is done jointly over all acquired sample thicknesses $t \in T$ (in our case $T = \{0.5, 1\}$ mm).

Definition. We formally define the fitting as a minimization across the volume parameter space \mathcal{V} :

$$\arg \min_{v \in \mathcal{V}} \sum_{t \in T} d(\mu \cdot F_{A|t}, F_{S|t}(v)), \quad (1)$$

where $v \equiv (\sigma_t, \alpha, g)$ and μ is a scaling coefficient which we elaborate below. The function d is a distance metric in the edge profile space, and we define it as

$$d(F_A, F_S) = \|F_A - F_S\|_2 + \lambda \cdot \left\| |\nabla F_A| - |\nabla F_S| \right\|_2, \quad (2)$$

with λ being a scalar parameter. Solving Eq. 1 for every material—and for each RGB channel—leads to an optical characterization of the whole material space \mathcal{M} . The solution itself is efficiently implemented as an exhaustive search in the material dictionary

Table 1. Measured optical parameters for the Stratasys Vero Opaque family of materials, constrained for scattering anisotropy $g = 0.4$. See Sec. 6 for details.

Material	Extinction coef. σ_t [mm^{-1}]			Scattering albedo α			Fit error		
	R	G	B	R	G	B	R	G	B
Cyan (C)	9.0	4.5	7.5	0.05	0.7	0.98	N/A	0.009	0.011
Magenta (M)	2.5	3.0	10.0	0.98	0.1	0.9	0.039	N/A	0.008
Yellow (Y)	2.25	3.75	19.0	0.997	0.995	0.15	0.041	0.022	N/A
Black (K)	5.0	5.5	6.5	0.35	0.35	0.35	0.003	0.002	0.002
White (W)	6.0	9.0	24.0	0.9991	0.9997	0.999	0.024	0.017	0.041




Fig. 6. The aggregate fitting error (y-axis) in dependence on the scaling factor $\mu \in [0.8, 1]$ (x-axis). As you can see, the error drops quickly from $\mu = 1$ to the global minimum at $\mu = 0.913$, and then grows steadily again. In our opinion, this supports the existence of a bias in the data and the effectiveness of our approach to counteract it.

(Sec. 5.2), taking only a few minutes on a desktop PC. You can find the results in Sec. 6.

However, because of the relative simplicity of our measurement setup, inherent acquisition inaccuracies as well as any modeling inconsistencies between the setups would lead to an imprecise measurement. To minimize that, we sought a fitting procedure with a high degree of robustness against the above mentioned issues. We have therefore developed three provisions in this direction: joint fitting over multiple sample thicknesses, regularization of the fitting, and regularization of the acquired data.

Joint fitting. The advantage of our setup is that it directly allows us to build a customizable degree of meaningful redundancy into the acquired data, by varying the sample thickness for each given material. While this does not safeguard against ‘global’ biases present in either dataset (acquired or simulated), it increases robustness against ‘local’ errors in each sample (e.g., surface scratches). We therefore seek a joint fit across all acquired thicknesses $t \in T$ (Eq. 1). For sake of efficiency we choose to work with two different thicknesses, although using more would likely increase the fitting precision even further, if desired.

Fitting regularization. The design of our profile distance metric d was motivated by the fact that we want to find a fit that is most similar in terms of:

- overall intensity of the profile, especially its ‘tails’ (predominantly determined by absorption, i.e., correlated with α);
- distortion induced to the step edge (mostly linked to scattering, i.e., correlated with σ_t and g).

The two terms of Eq. 2 directly correspond to these two criteria. Since both F_A and F_S are inherently noisy, we robustly estimate their gradient magnitudes $|\nabla F_A|$ and $|\nabla F_S|$ using the method of Luo et al. [2006] based on discrete wavelet decomposition. The regularization parameter λ controls the relative importance of the above criteria; we have used $\lambda = 20$ in all our measurements.

Acquisition regularization. While the previous two measures treat possible inaccuracies in the data or acquisition, the last one addresses any remaining inconsistencies in the setup geometry and lighting. This complex issue is simplified by the fact that the measurement takes place in a very small region, hence we assume a constant bias in the data and model it by a single multiplicative factor μ , see Eq. 1.

To find the value of μ , we seek to obtain the lowest possible *aggregate error* when characterizing the entire material space \mathcal{M} . We therefore perform a meta-optimization on top of Eq. 1, across all materials and RGB channels:

$$\arg \min_{\mu \in \mathbb{R}^+} \frac{1}{\mu} \left[\sum_{m \in \mathcal{M}} \sum_{c \in \text{RGB}} \min_{v \in \mathcal{V}} \sum_{t \in T} d(\mu \cdot F_A|_{m,c,t}, F_S|_{c,t}(v)) \right]. \quad (3)$$

The division by μ is necessary to normalize the resulting fitting error, which would otherwise tend to zero for $\mu \rightarrow 0$ since the dictionary in fact does contain profiles with virtually zero intensity (refer to Fig. 5).

Given that we expect $\mu \approx 1$, we quickly determined that $\mu \in [0.8, 1]$. A linear search in this interval with the step of 10^{-3} yielded $\mu = 0.913$ (see Fig. 6), and took several hours on a desktop PC.

6 RESULTS

In this section, we present the measured optical parameters for the Stratasys Vero Opaque family, comprising Cyan (C), Magenta (M), Yellow (Y), Black (K) and White (W) materials. The resulting parameters are listed in Table 1.

For a detailed demonstration of the fitting, we document the results for the white material in Fig. 7. Here we show, for every RGB channel, the resulting joint fit (obtained with Eq. 1) of the Vero White material across the two sample thicknesses we operate with. Each plot also shows the profile gradient magnitude curves, and for illustration also the neighboring dictionary records adjacent to the best fitting one. As you can see, our procedure obtains a good match

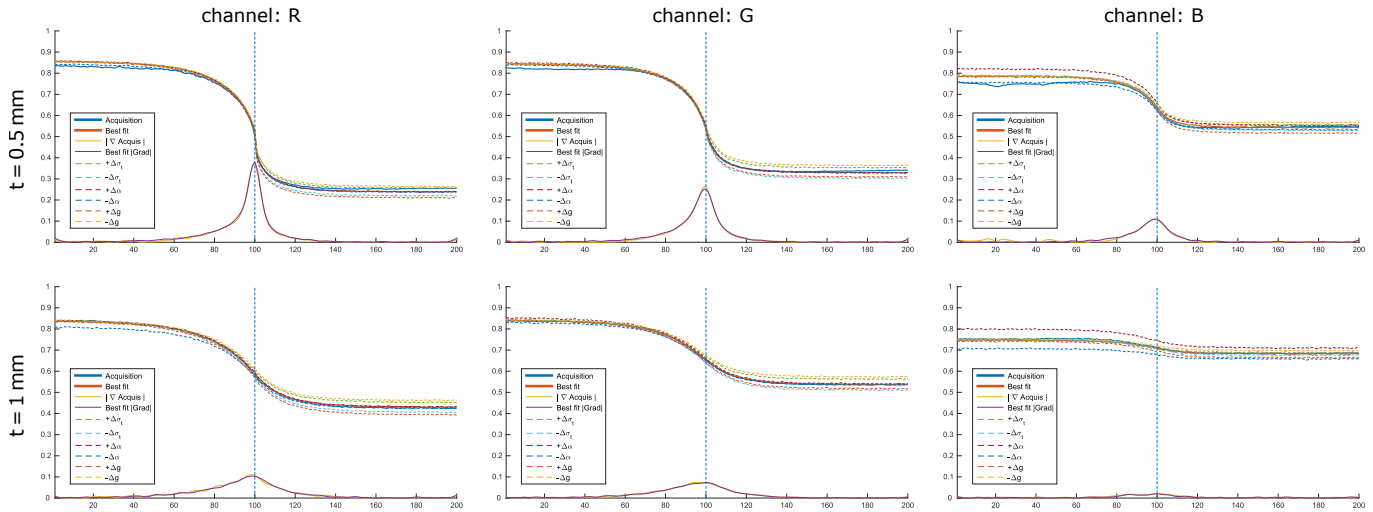


Fig. 7. Fitting results for the Vero White material. The plots show the best-fit profiles, as well as the corresponding gradient magnitude curves and adjacent dictionary records (see the corresponding legends).

for the profile itself, but also its gradient. We obtain similarly good fits for the remaining materials (see the fitting errors in Table 1), with the following exception.

The main limitation of our setup turned out to be the measurement of strongly absorbing materials, as the resulting edge profiles have very low discriminability in that part of \mathcal{V} . This has been an issue R channel of cyan, G channel of magenta and B channel of the yellow material (marked bold in Table 1). We have resolved this issue by printing a color chart with patches of all possible *equal* combinations of the CMYKW materials. We then ran a brute-force search in the problematic channels to identify their values, by seeking a best fit of our color predictions to the printed color chart, using our analytical tonal mapping (described in the paper) as a prediction model.

6.1 Validation

We present three different ways to validate our data: thin-material setup, bulk-material setup, and cross-validation against the similarity theory [Wyman et al. 1989; Zhao et al. 2014].

Thin-material validation. As the first, most direct step, we validated the measured data in the same configuration in which the fitting operates. This is the setting presented already in Fig. 2. As you can see, the simulated results render a high-fidelity impression of their acquired counterparts. Notice this is not a trivial consequence of the fitting, since this operates on averaged 1D profiles and has to jointly fulfill the constraints of the available sample thicknesses.

Structured validation. We use the measured data to compute predictions of the optimized printouts produced by our method. In Fig. 8 we show several examples, demonstrating the accuracy of predicting the appearance of the resulting prints.

Theoretical validation. Similarity theory (originating in general particle transport, brought to medical imaging by Wyman et al. [1989] and recently gaining attention in computer graphics [Zhao

et al. 2014]) states that under certain conditions the volume parameters space will contain *equivalence classes* which can lead to equivalent (or “similar”) transport solutions. While these conditions are not guaranteed to apply in our case, we can still attempt to evaluate our measurements within the scope of this framework.

According to the simplest (first-order) similarity relation, the scattering coefficient σ_s and absorption coefficient σ_a (where $\sigma_t = \sigma_s + \sigma_a$ and $\alpha = \sigma_s/\sigma_t$) will lead to a similar transport solution as for some other medium with σ'_s and σ'_a , if

$$\sigma_a = \sigma'_a \quad (4)$$

$$\sigma_s(1-g) = \sigma'_s(1-g'). \quad (5)$$

Expressing the latter relationship for σ'_s as

$$\sigma'_s = \frac{\sigma_s(1-g)}{1-g'} \quad (6)$$

therefore gives us a way to derive the parameters of a different, *similar* material, given a target g' .

This is exactly demonstrated in Fig. 9 for the Vero White material. For every fit we marked its location in the volume parameters space by a solid red crosshair (overlaid in the distance metric maps). Then, according to Eqs. 4 and 6, we computed the parameters of the similar media for all other anisotropy values (i.e., representing g' in Eq. 6). As you can see, the resulting ‘extrapolated’ parametrizations (marked with dashed crosshairs) tend to be well aligned with the local minima of our distance metric, across the entire anisotropy sub-space.

This property is actually beneficial, as it enables us to constrain the fitting process to a specific anisotropy g . We found that the best such constrained fit is obtained for $g = 0.4$, which is reflected in the measured data in Table 1. We use this property in our material mapping method, as described in the main paper.



Fig. 8. Predictions of the appearance of our printouts using the measured data from Table 1.

7 DISCUSSION

In this document, we have provided an in-depth description of our approach to measure intrinsic optical parameters of solid translucent materials used, e.g., in poly-jetting 3D printing processes. We believe our approach provides a viable alternative to current measurement methods, as it achieves visually plausible results, while keeping the overall design accessible without compromising its reliability. This is mainly achieved by a robust joint fitting procedure, which ensures that a reasonable set of parameters is found even for imperfect data.

The main limitation have turned out to be very absorbing materials, which we resolved by an independent fit as explained in Sec. 6. A more reliable measurement of these materials with our setup would require much thinner samples than we used, they are however difficult to polish properly. We have considered embedding a single absorbing layer into a block of clear material, but a further investigation is needed here.

Future improvements. In order to increase the overall precision of the acquisition, several refinements could be made, including:

- more consistent polishing of the samples,

- modulation signal source made from a more rigid, stable material,
- higher illumination quality and uniformity,
- better calibration target (Spectralon instead of a standard color checker),
- better index-matched optical binding medium, and ultimately,
- dense spectral measurement to avoid potential issues with color metamerism.

Aside from these, other, qualitative improvements are left to be examined. For instance, in spite of the density of our dictionary, we could still benefit from a finer sampling. As this is costly, an alternative could be to ‘interpolate’ several closest fits instead of selecting only the single closest one. This would require re-designing the profile distance metric (Eq. 2) to yield a *signed* distance, in order to provide an implicit ordering needed for such interpolation.

REFERENCES

- Dave Coffin. 1997–2016. DCRaw. (1997–2016). <https://www.cybercom.net/~dcoffin/dcrav/>.
- Eugene d’Eon and Geoffrey Irving. 2011. A Quantized-diffusion Model for Rendering Translucent Materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30, 4 (July 2011), 56:1–56:14. <https://doi.org/10.1145/2010324.1964951>
- Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. 2010. Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 4 (2010), 62:1–62:10. <https://doi.org/10.1145/1778765.1778799>
- Wayne Rasband et al. 1997–2016. ImageJ. (1997–2016). <https://imagej.net>.
- Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics* 32, 6 (2013), 162:1–162:13. <https://doi.org/10.1145/2508363.2508377>
- Miloš Hašan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 3 (2010), 61:1–61:10.
- Wenzel Jakob. 2010. Mitsuba renderer. (2010). <http://www.mitsuba-renderer.org>.
- J. W. Luo, J. Bai, and J. H. Shao. 2006. Application of the wavelet transforms on axial strain calculation in ultrasound elastography. *Progress in Natural Science* 16, 9 (2006), 942–47.
- Srinivasa G. Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, and Henrik Wann Jensen. 2006. Acquiring Scattering Properties of Participating Media by Dilution. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (July 2006), 1003–12. <https://doi.org/10.1145/1141911.1141986>
- Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. 2013. Fabricating translucent materials using continuous pigment mixtures. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32, 4 (July 2013), 146:1–146:12. <https://doi.org/10.1145/2461912.2461974>
- J. Radon. 1986. On the determination of functions from their integral values along certain manifolds. *IEEE Transactions on Medical Imaging* 5, 4 (1986), 170–76.
- Johannes Schindelin, Curtis T. Rueden, Mark C. Hiner, and Kevin W. Eliceiri. 2015. The ImageJ ecosystem: An open platform for biomedical image analysis. *Molecular Reproduction and Development* 82, 7-8 (2015), 518–29.
- Douglas R. Wyman, Michael S. Patterson, and Brian C. Wilson. 1989. Similarity relations for the interaction parameters in radiation transport. *Applied Optics* 28, 24 (1989).
- Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order Similarity Relations in Radiative Transfer. *ACM Transactions on Graphics* 33, 4 (July 2014).




Fig. 9. Maps of the distance metric (Eq. 2) for the entire optical parameter space \mathcal{V} for the Vero White material: each sub-plot corresponds to a single value of g , with the co-albedo $\tilde{\alpha} = 1 - \alpha$ varying along the x-axis and σ_t along the y-axis. The red crosshairs then show the positions of the fits for different anisotropy g values: the solid one pinpoints the best fit, and the dashed ones the remaining ones for the respective g . The red dashed lines show the predictions according to the *similarity theory* – please refer to Sec. 6.1 for details on this.