Scattering-aware Texture Reproduction for 3D Printing

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

Fig. 1. A still life photograph of our optimized printouts. The thickness of all the pictured samples is 1 cm.

Color texture reproduction in 3D printing commonly ignores volumetric light transport (cross-talk) between surface points on a 3D print. Such light diffusion leads to significant blur of details and color bleeding, and is particularly severe for highly translucent resin-based print materials. Given their widely varying scattering properties, this cross-talk between surface points strongly depends on the internal structure of the volume surrounding each surface point. Existing scattering-aware methods use simplified models for light diffusion, and often accept the visual blur as an immutable property of the print medium. In contrast, our work counteracts heterogeneous scattering to obtain the impression of a crisp albedo texture on top of the 3D print, by optimizing for a fully volumetric material distribution that preserves the target appearance. Our method employs an efficient numerical optimizer on top of a general Monte-Carlo simulation of heterogeneous scattering, supported by a practical calibration procedure to obtain scattering parameters from a given set of printer materials. Despite the inherent translucency of the medium, we reproduce detailed surface textures on 3D prints. We evaluate our system using a commercial, five-tone 3D print process and compare against the printer's native color texturing mode, demonstrating that our method preserves high-frequency features well without having to compromise on color gamut.

1 INTRODUCTION

The past decade has witnessed an unprecedented technological development of 3D printing, both in breadth and depth: consumer-grade printers are now commonly available for a range of purposes, while increasingly advanced techniques allow us to fabricate novel shapes, mechanical properties, and appearances. Notably, the printers' capabilities have improved dramatically from printing single-material objects to producing detailed structures with per-voxel material variation. Such capabilities are most prominent in printers based on the photo-polymer jetting process, which can

currently produce full-color prints with a resolution in the order of 10 μm [Sitthi-Amorn et al. 2015; Stratasys 2017].

The resins used as print materials in commercial photopolymerization printers are inherently translucent, i.e., exhibit significant sub-surface scattering. This serves effective color mixing in full-color print processes, thus commercial printer drivers offer high-quality color reproduction. At the same time, the resulting light diffusion leads to over-blurring and potential color bleeding when printing spatially-varying color textures. This translucent 'cross-talk' between surface points also strongly depends on the internal structure of the volume surrounding each surface point.

Our work aims at enabling 3D-printing of high-fidelity color textures by compensating for these artifacts through careful spatial (volumetric) assignment of color materials.

This bears resemblance to works on fabrication of custom sub-surface scattering [Dong et al. 2010; Hašan et al. 2010] and color in 3D prints [Brunton et al. 2015], which produce impressive physical renditions of naturally translucent objects, such as marble, wax and organic tissues, capitalizing on the translucency of the under-lying materials. Such works, however, use approximate models for light diffusion, and often accept the apparent blur as an immutable property of the print medium; their models work best if the target appearance lacks high spatial frequencies, conveying a sense of translucency in proportion with the actual translucency of the materials.

In contrast, our work specifically targets the challenging case where the target appearance *does contain* high-frequency details, as in the common use case of assigning a detailed color texture to a 3D object, regardless of the translucency characteristics of the printer material. Our work counteracts heterogeneous scattering to obtain the impression of a crisp diffuse albedo texture on top of the 3D print, by optimizing for a fully volumetric material distribution that preserves the target appearance.

We employ an efficient numerical optimizer on top of a general Monte-Carlo simulation of heterogeneous scattering. The simulation is supported by a practical calibration procedure to obtain scattering parameters from a given set of printer materials.

This paper presents several novel contributions:

- complete end-to-end pipeline aimed at achieving the above objective, i.e., printing a sharp surface texture approaching opaque appearance qualities, subject to physical limitations (Sec. 3);

- definition of a custom nonlinear optimization method to compute a spatial material distribution that retains high-frequency details, compensating for the inherent translucency of the print materials (Sec. 7);

- techniques of converting between different material parametrizations to facilitate the entire process up to the printing itself (Secs. 4 and 6);

- practical novel way to acquire the optical properties of the available printing materials (Sec. 5).

We demonstrate the capabilities of the resulting pipeline in a number of examples (Fig. 1, Sec. 9) and discuss limitations that translucent materials have in reproducing such opaque appearances (Sec. 10).

2 BACKGROUND AND RELATED WORK

2.1 Subsurface Scattering Effects

Established printing pipelines for color texture reproduction on 3D prints still follow strategies from 2D printing: their appearance formation models assume local mixture of primaries, without lateral cross-talk beyond *dot gain*, which also is routinely considered in 2D print production pipelines [Stollnitz et al. 1998].

Few works model the effect of lateral (subsurface) scattering in the material, and most assume it to be a spatially uniform blur that attenuates high frequencies across the surface. The resulting translucent appearance is taken as a welcomed effect [Brunton et al. 2015], kept to a minimum through specifically designed printing materials [Babaei et al. 2017], and/or is approximately compensated for by de-convolution [Babaei et al. 2017] or through an adaptation of unsharp-masking [Cignoni et al. 2008; Pintus et al. 2010].

As we will show, however, printing materials used in commercial 3D printers vary significantly in their effective scattering properties (Sec. 5). Such scattering heterogeneity implies a strong effect on the final appearance of a 3D printout. Hašan et al. [2010] and Dong et al. [2010] exploit this to control local scattering profiles by suitable volumetric multiplexing of materials. Their models, however, approximate scattering only locally, ignoring heterogeneous scattering over larger lateral distances.

Bidirectional scattering-surface reflectance distribution functions (BSSRDFs) are a common approximation for translucent appearance in computer graphics [Christensen 2015; d'Eon and Irving 2011; Donner et al. 2008; Jensen and Buhler 2002; Jensen et al. 2001; Song et al. 2009]. We however argue (Sec. 4) that, in our application, any such model that considers scattering purely locally is bound to yield visible inaccuracies, as for most commercial printing materials the effective scattering distance is substantial (e.g., on average, light travels much farther than the print's resolvable feature size before contributing to the surface's appearance). Our work hence relies on a full Monte Carlo simulation of light transport in heterogeneous media to model appearance from a given material distribution. We offset the increased computational cost by employing an inverse-rendering algorithm with a high convergence rate.

2.2 Fabrication Methodology

In computer graphics, at least two main paradigms of specifying fabricated objects are recognized: direct (forward) and constraint-driven (inverse). The OpenFab [Vidimče et al. 2013] and Spec2Fab [Chen et al. 2013] systems are good respective examples of these two types of approaches. Both OpenFab and Spec2Fab can be seen as general frameworks which can accommodate various methods of specifying and/or optimizing for desired material properties, both visual and mechanical. Examples of these include custom BRDFs and other surface properties [Lan et al. 2013; Matusik et al. 2009; Rouiller et al. 2013; Weyrich et al. 2009], voxel micro-structure [Brunton et al. 2015; Doubrovski et al. 2015; Lou and Stucki 1998], continuous mixtures of inks [Papas et al. 2013], or other, technology-specific features [Pintus et al. 2010; Reiner et al. 2014].

Within that taxonomy, our work targets the inverse problem of additive fabrication of spatially varying subsurface scattering properties from a fixed set of translucent materials. As such, it falls into

Table 1. Detailed summary of the spaces our method operates in (listed in the order of appearance in the text).

Space	Type	Definition	Description
$\mathcal{C}$	continuous	$[0, 1]^3$	Linear RGB colorspace describing the object appearance under a canonical 'white-sky' illumination. Under these conditions, we treat color as equivalent to observed reflectance, i.e., the relative portion of incident energy reflected by the object as a result of both surface and sub-surface light transport.
$\mathcal{V}$	continuous	$(\sigma_t, \alpha, g, \eta)^3$	Space of intrinsic optical parameters of the (volumetric) printing materials. For each RGB channel, it defines four parameters: $\sigma_t \in [0, \infty)$, the material's extinction coefficient (optical density); $\alpha \in [0, 1]$, the single-scattering albedo, $g \in (-1, 1)$, the scattering anisotropy; and $\eta \in (1, \infty)$, the refractive index.
$\mathcal{M}'$	discrete	{C, M, Y, K, W}	Physical space of printer-specific materials, each of which has a unique set of optical parameters $\in \mathcal{V}$ used by our prediction techniques. Each voxel of the printed object must be assigned exactly one material.
$\mathcal{M}$	continuous	$[0, 1]^5$	Virtual CMYKW tonal space of the printer. Given the properties of $\mathcal{M}'$, it is a *barycentric* space, i.e., for a given tonal mixture $w \in \mathcal{M}$ it must hold that $\sum w_i = 1$.

the same category as the works by Dong et al. [2010] and Hašan et al. [2010]. We observe, however, that once lateral subsurface scattering is taken into account, their goal of simultaneously controlling scattering *and* absorption at every point is an overdetermined problem. Instead, we forgo the control over the shape of scattering kernels and constrain diffuse albedo alone, providing the solver with the necessary degrees of freedom to match a target (albedo) texture despite lateral cross-talk through (heterogeneous) scattering. On the flip side, the lack of scattering control requires us to assume smoothly varying illumination. As we demonstrate in Sec. 9, however, under many observing conditions this does not cause any issues.

2.3 Pre-correction

Pre-correction generally refers to the modification of content to counteract subsequent degradations in the output. Most closely related to our problem is the work by Cignoni and colleagues [Cignoni et al. 2008; Pintus et al. 2010], who recover the visibility of geometric details obfuscated by light diffusion, by countering the effect through modulation of the material albedo. Their key idea is to render the target object, once with simulated subsurface scattering, and once without, to determine the shading difference that needs adding back to maintain the look of a scattering-free surface. Their method is elegant and effective, but limited to mid-tone target albedos, to allow for sufficient freedom to darken and lighten the albedo as required; furthermore, it does not consider the impact of light diffusion on the albedo texture itself, instead assuming artifact-free texturing of the 3D print.

Our method is inspired by this approach, but tries to overcome these restrictions by accurately modeling crosstalk within the heterogeneous material while controlling material deposition volumetrically, rather than applying 2D corrections only.

Other relevant works include pre-correction for light pollution (cross-talk) in display [Konrad et al. 2000] and projection systems [van Baar et al. 2011], and improving visual acuity for observers with degraded vision [Montalto et al. 2015]. These problems conceptually compare to ours, but differ in aspects which make our problem more difficult to solve. Namely, the degradation due to light pollution [van Baar et al. 2011] has a comparatively simple additive character, while blurring due to subsurface scattering in heterogeneous materials has a complex, nonlinear behavior. Also, pre-correcting for vision degradation [Montalto et al. 2015] can reasonably assume a constant point spread function (PSF), while in our case, the subsurface

scattering kernel has a high, material-dependent spatial variance. In addition, in both above cases, the degradation can be formulated as a convolution-like process with a relatively small kernel, while our problem deals with very wide-support 'kernels' with heavy tails that cannot be easily truncated.

2.4 Material Editing

Material editing often aims to solve a similar problem to ours: to change the structure of a (typically virtual) object so that it conforms to given edits (constraints). An excellent survey is provided by Schmidt et al. [2014]; here, we limit the discussion to editing of translucent objects.

Initial works use precomputed radiance transfer (PRT) to globally change bulk scattering and absorption parameters [Wang et al. 2008; Xu et al. 2007]. Extending these PRT-based approaches to heterogeneous materials, however, would lead to an intractable dimensionality increase.

Song et al. [2009; 2013] overcome this by factorizing BSSRDF profiles into individual constituent materials and (layer) geometries. Hašan and Ramamoorthi [2013] even consider generalized volumetric structures, decomposing the edited object into voxels and employing a database of albedo-space derivatives in respect to each voxel to rapidly modify the resulting appearance.

In general however, these techniques differ from ours: they aim at obtaining interactive but approximate responses to edits, be it directly in the material parameter space [Song et al. 2009; Wang et al. 2008; Xu et al. 2007] or inversely by posing constraints on the appearance [Hašan and Ramamoorthi 2013].

3 METHOD OVERVIEW

Our method uses a physically-based Monte-Carlo scattering model to predict the appearance of a heterogeneous, 3D-printed volume (Sec. 4). To obtain the optical parameters for our print materials, we employ a lightweight calibration procedure outlined in Sec. 5.

Once the optical parameters for the available set of print materials are known, the core task of our work is to map a *surface* color specification to a *volumetric* distribution of materials that matches this color specification as closely as possible.

In a first approximation (presented in Sec. 6), we achieve this through a sequence of remapping steps, using the color and materials specifications outlined in Table 1. Given a color texture specification in (linear) RGB space $\mathcal{C}$, gamut-mapped to the printer's native
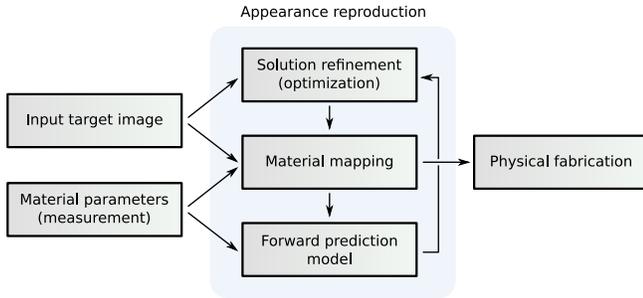
Fig. 2. High-level scheme of our pipeline, discussed in Sec. 3.

gamut constraints and extruded into the volume, we determine a tonal volume that consists of idealized mixtures of the print materials in $\mathcal{M}$ (in our case CMYKW). In this step we assume a perfect, continuous tonal mixture and no lateral cross-talk, before quantizing values in $\mathcal{M}$ to a half-toned representation $\mathcal{M}'$ of realizable material assignments, ready for print.

The resulting object, however, would still be subject to excessive blur, due to lateral scattering not having been taken into account. We hence feed the predicted residual difference into an outer optimization loop (cf. Fig. 2; detailed in Sec. 7) to obtain a refined RGB specification that, when mapped to the print materials, matches the input specification well. Key to effective convergence of that optimization is a carefully designed heuristic to drive the appearance residual into a volumetric RGB specification, prioritizing embeddings at different depths according to the *sign* of the residual texture. As such, it significantly differs from traditional error diffusion schemes (e.g., [Brunton et al. 2015]).

We evaluate our approach using a commercial poly-jetting printer (Sec. 8) and demonstrate its performance by assigning highly detailed target textures to solid slabs (Sec. 9), comparing to printouts using the printer's native color texturing mode. While our results show planar surfaces, there is no principal reason that would prevent a future extension to more general shapes. Please see section Sec. 10 for a more detailed discussion.

4 OBJECT APPEARANCE MODEL

This section describes the forward model, which is used in our method's prediction step to obtain an appearance prediction from an object's parametric description.

In spite of the popularity of analytic BSSRDFs in translucent appearance fabrication [Dong et al. 2010; Hašan et al. 2010; Papas et al. 2013] and a progress in their *spatially varying* application [Sone et al. 2017; Song et al. 2009; Song and Wang 2013], we opt *not* to use an analytic BSSRDF as a forward prediction model. The main reason is that the current methods' accuracy decreases when the transport scale exceeds the size of the target appearance features (see Fig. 3). This certainly happens in our conditions, as the target texture features are in the order of 10 μm, while our materials have mean free paths in the order of 0.1 mm (cf. Sec. 5), leading to visible scattering as far as 10 mm into the object. In addition, common spatially-varying BSSRDF models cannot fully account for volumetric transport due to lateral material heterogeneities, which is essential for our method.

Instead, we use standard path tracing (implementation provided by the Mitsuba renderer [Jakob 2010]) as the prediction model for the textured object's appearance. While computationally expensive, the accuracy of a Monte Carlo solution outweighs the cost, as this is directly correlated with the quality of the resulting color and texture reproduction. Path tracing performs well under natural ambient illumination that we target, and can be sped up by disabling next-event estimation (from within the medium) which is inefficient for non-index-matched objects we deal with here.

4.1 Forward Prediction Model

We seek to obtain the textured object's appearance, where we consider the object in a canonical position viewed from the top, illuminated by a normalized 'white sky' constant environment map. Under those conditions, the target appearance is defined as a diffuse RGB reflectance of the object in $\mathcal{C}$ (refer to Table 1 for an overview of our operating spaces). Formally we define the forward prediction model as

$$f \colon \mathcal{V} \to \mathcal{C}. \tag{1}$$

Eq. 1 maps a volume (3D matrix) of the object's optical parameters $V \in \mathcal{V}$ to a color image (2D matrix) $C \in \mathcal{C}$.

Each CMYKW material $m \in \mathcal{M}'$ has a corresponding set of physical volume parameters $v_m \in \mathcal{V}$, which we measure via the procedure detailed in Sec. 5. Since every voxel can only be assigned a single material, the object volume $V = \{v_i\}$ is eventually quantized (halftoned): $v_i \in \{v_m\}$. We detail this quantization in Sec. 6.2.

5 MATERIAL CALIBRATION

Our prediction model described in Sec. 4 relies on the knowledge of the optical data $\in \mathcal{V}$ (Table 1) of our printing materials. While state-of-the-art measurement methods can deliver excellent accuracy (see [Gkioulekas et al. 2013b] and references therein), this can come at the cost of complex design and expensive equipment.

We instead propose a novel simple and affordable measurement method that well serves our need of obtaining data that yield *visually* accurate prediction. In this section we briefly explain the concept
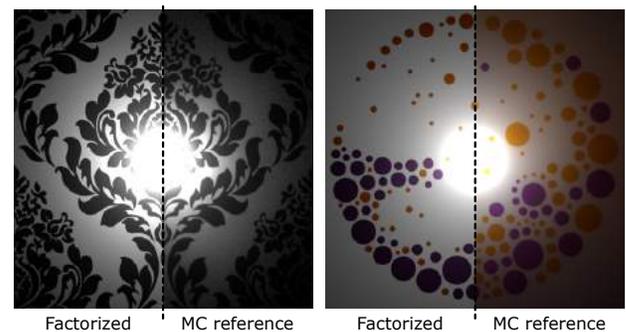
Fig. 3. Translucent transport in a material with a constant density (proportional to our printing materials) and laterally varying albedo, illuminated by a circular spotlight. We compare renderings using the factorized BSSRDF by Song et al. [2009] with a Monte Carlo path-traced reference. The BSSRDF tends to overestimate the transported energy and does not well reproduce subtle color bleeding from nearby features.

of our approach (Sec. 5.1), describe our fitting-based data extraction (Sec. 5.2) and present the resulting measurements (Sec. 5.3). A significantly expanded exposition is provided in the supplemental material.

5.1 Measurement

Since it is hard to measure the optical parameters of solid volumetric materials directly, we base our method on inverse rendering. Our approach is to acquire a complex, extrinsic interaction of light with the medium, and reproduce it in a simulation under equivalent conditions. By generating a sufficiently dense dictionary of virtual observations using different optical parametrizations $\in \mathcal{V}$, we obtain the measurements by finding a fit that matches the acquired physical observation [Gkioulekas et al. 2013b].
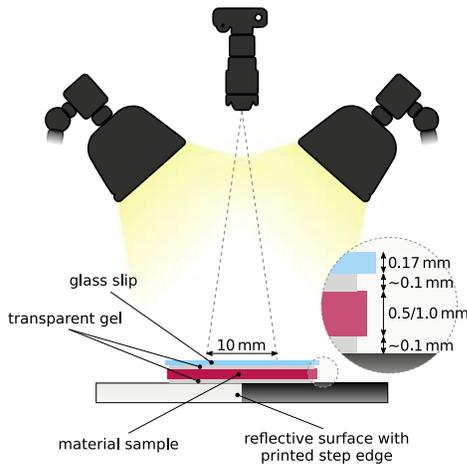
Fig. 4. Sketch of our acquisition setup (not to scale).

Our acquisition setup observes a simple printed step edge through a thin slice of the printing material, as sketched in Fig. 4. The setup consists of six main components:

- *Sample.* A thin slice with the dimensions 40×40×0.5/1 mm, 3D-printed with one of the materials from $\mathcal{M}'$. Surface finish as smooth as possible.

- *Signal modulation.* Reflective surface containing two black and white regions separated by a straight step edge, placed under the sample.

- *Camera.* Located about 1 m above the sample, vertically aligned and laterally centered with the observed step edge, imaging the sample at 28 μm resolution.

- *Lighting.* Two diffuse, neutral illuminants, symmetrically placed around the sample. Positioned roughly along the bisector between the camera's optical axis and the sample plane, to minimize the influence of single-scattering events while still providing sufficiently strong and even illumination.

- *Thin glass slip.* To provide a flat, clear, smooth air interface with roughly known BRDF on top of the sample.

- *Transparent gel.* For (approximate) index matching between the components, to suppress internal reflections and the influence of dust or residual roughness on the sample surface.
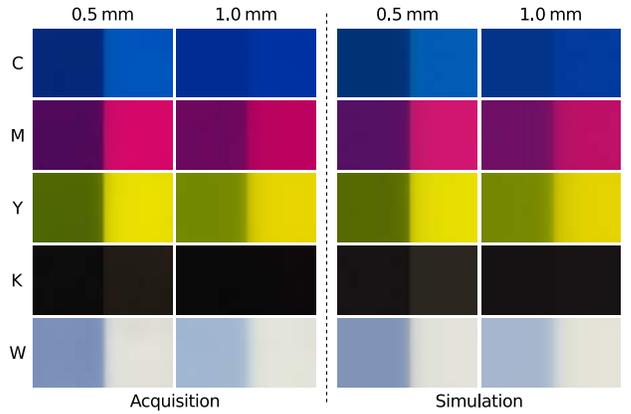
Fig. 5. Comparison of the step edges captured through the printing materials $\mathcal{M}$ = {C, M, Y, K, W} by our physical (**left**) and virtual (**right**) setups, at two different sample thicknesses. The images cover an area of 2×1 cm.

We built a physical version of this setup using a consumer DSLR camera (Canon EOS 700D body, EF-S 18–135 mm lens set to 135 mm), two 55 W, 5500 K-equivalent fluorescent lamps, thin microscope glass slides and a clear ultrasound transmission gel, chosen for its bubble-free applicability [Donner et al. 2008]. We printed the fitting edge on a Toughprint Waterproof paper using an office laser printer. The resulting edge appearances captured by our physical setup are shown in Fig. 5, left.

As for the virtual counterpart of this setup, we use the same environment as for the prediction (Sec. 4): a hand-modeled scene for the Mitsuba renderer, with path-tracing as a simulation framework. We generated the fitting dictionary for two different sample thicknesses – 0.5 mm and 1 mm – by sampling the space $\mathcal{V}$ as follows: $\sigma_t \in \{1, \ldots, 30\} \, \text{mm}^{-1}$, $\alpha \in \{0.05, \ldots, 0.99999\}$ and $g \in \{0, \ldots, 0.8\}$ (see the supplemental material for the complete enumeration). For the refractive index we used the value of $\eta = 1.5$ measured separately by optical ellipsometry. The resulting dictionary contains about 34k records and took two days on a 500-node CPU cluster to compute. We will make the dictionary publicly available, along with the detailed documentation of the data format.

A desired property of our setup is that it uses the same simulation framework (Monte Carlo path tracing) as a prediction model for the measurement as well as our optimization itself. This guarantees consistency of the obtained data and the resulting visual match.

5.2 Model Fit

Our fitting procedure operates in the space of 1D *edge profiles* $F \in C$ resulting from vertically averaging the captured edge appearances. We then find the *most similar* simulated edge profile F_S for each profile F_A acquired by our physical setup. This is done jointly over both acquired sample thicknesses $t \in T$ (in our case $T = \{0.5, 1\}$ mm).

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

$$\underset{(\sigma_t, \alpha, g) \in \mathcal{V}}{\arg \min} \sum_{t \in T} d\left(\mu \cdot F_{A|t}, \, F_{S|t}(\sigma_t, \alpha, g)\right), \qquad (2)$$

where μ is a scaling coefficient which we explain below. The function d is a distance metric in the edge profile space, and we define it as

$$d(F_A, F_S) = \left\| F_A - F_S \right\|_2 + \lambda \cdot \left\| |\nabla F_A| - |\nabla F_S| \right\|_2, \quad (3)$$

with λ being a scalar parameter. Solving Eq. 2 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 (Sec. 5.1), taking only a few minutes in Matlab on a desktop PC.

To obtain robust results, we regularize the fitting via the two scalar parameters μ and λ. While λ regularizes the fit locally by seeking a match in the profile gradients (in addition to their intensities), μ is a global regularizer that compensates for any possible global bias resulting from modeling mismatches between the physical and virtual setups. We determined $\lambda = 20$ empirically and $\mu = 0.913$ systematically – please refer to the supplemental material for details.

The simulated edge appearances of the best dictionary fits according to Eq. 2 are shown in Fig. 5, right.

5.3 Results

To our surprise, we found that the edge appearance profiles follow the predictions of the similarity theory [Wyman et al. 1989; Zhao et al. 2014] to a good accuracy, even in our thin-slice conditions. The main implication is that we can find a reliable fit of extinction coefficient σ_t and albedo α for *any* desired anisotropy g present in the dictionary. This benefits us greatly, because constraining g significantly simplifies the construction of our tonal mapping, as described later in Sec. 6.1.

We therefore performed a search across all g values in our dictionary and found that $g = 0.4$ yields the lowest overall fitting error. The resulting measurements are summarized in Table 2 and validated in Fig. 5 for the thin slices and later in Sec. 9 for complex heterogeneous material compositions.

The main limitation of our setup is the measurement of strongly absorbing materials, as the resulting edge profiles have almost no discriminability in that part of $\mathcal{V}$. This turned out to be an issue for the R channel of cyan, G channel of magenta and B channel of the yellow material (marked bold in Table 2). We have resolved this issue by printing a color chart with patches of all possible *equal* combinations of the CMYKW materials and their subsets (shown later in Fig. 8). We then ran a brute-force search in the problematic channels to identify their values, by seeking a best fit of our color

Table 2. Measured optical parameters for the Stratasys Vero Opaque family of materials, constrained for scattering anisotropy $g = 0.4$.

Material	Ext. coef. σ_t [mm^{-1}]			Scattering albedo α		
	R	G	B	R	G	B
Cyan	**9.0**	4.5	7.5	**0.05**	0.7	0.98
Magenta	2.5	**3.0**	10.0	0.98	**0.1**	0.9
Yellow	2.25	3.75	**19.0**	0.997	0.995	**0.15**
Black	5.0	5.5	6.5	0.35	0.35	0.35
White	6.0	9.0	24.0	0.9991	0.9997	0.999

predictions to the printed color chart, using our tonal mapping from Sec. 6.1 as a prediction model for efficiency sake.

6 MATERIAL MAPPING

At the center of our method is the requirement to repeatedly transform between colorspace appearance definitions and material definitions, during the overarching optimization (Sec. 7). Here, we present efficient means to accurately convert between the different working spaces listed in Table 1.

For conceptual reasons, we describe our pipeline in terms of two modes of operation, where the former is a subset of the latter as illustrated in Fig. 6:

- a direct, 'single-pass' appearance reproduction mode, agnostic of the global light transport within the printing materials, and
- an iterative, 'multi-pass' optimization-driven mode, which makes use of the direct mode to produce an improved match between material composition and the target appearance.

Qualitatively the direct mode is compatible with other state-of-the-art 3D color reproduction pipelines [Babaei et al. 2017; Brunton et al. 2015], although it needs be noted that our solution has a markedly different technical basis consistent within our scattering-aware optimization as a whole.

We now delineate the steps involved in our reproduction pipeline (also refer to Fig. 6), followed by detailed descriptions of the more complex mapping procedures in Secs. 6.1 and 6.2.

1. The target image T is vertically extruded down to a desired depth in the RGB working volume $X \in C$; the remainder of X is filled with white. The working volume X is then per-voxel transformed into the tonal space volume M using the $C \rightarrow \mathcal{M}$
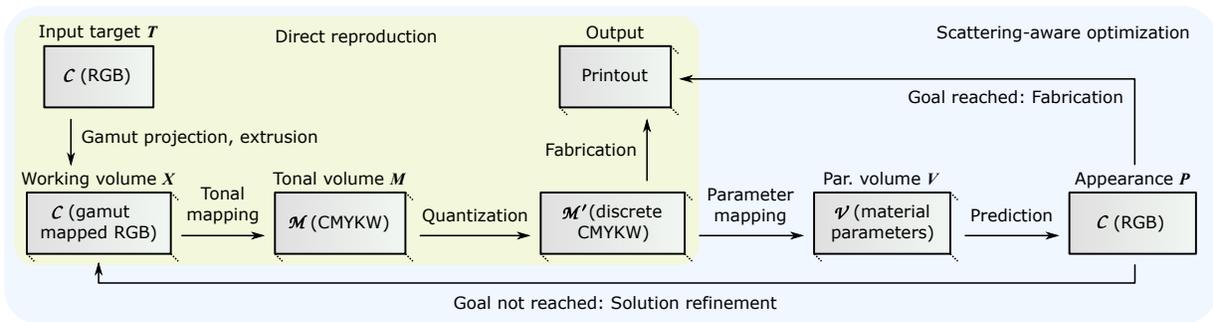
Fig. 6. Our reproduction pipeline, described in detail in Secs. 6 and 7. Note the schematic distinction between 2D and 3D working data.
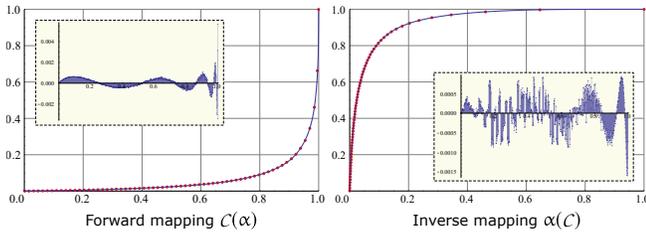
Fig. 7. Plots of the bi-directional mapping between scattering albedo α and color C (solid blue lines). The red points show the tabulated fitting data (sub-sampled by the factor of 16 for clarity). Note that the plots are normalized to [0,1], i.e., show only the volumetric portion of the reflected energy. The insets show the residual errors of the fits.

mapping detailed in Sec. 6.1. Note this also implicitly performs the gamut projection of the target image T itself, which then becomes the actual target in the algorithm.

2. The working volume $M \in \mathcal{M}$ is quantized (half-toned) so that each voxel is exclusively assigned a single material from the working set $\mathcal{M}'$, as detailed in Sec. 6.2. This already constitutes a printable result, concluding the direct reproduction mode.

3. Each voxel $m \in M$ is assigned three optical parameter vectors $v_m \in \mathcal{V}$ (one for each RGB channel) measured for its respective material $m \in \mathcal{M}'$.

4. The forward prediction model $f(V) \to C$ (described in Sec. 4) is used to compute the appearance of the working solution.

5. The predicted appearance image C is used to adjust the current solution in the optimization phase, which we describe in Sec. 7. This process is iteratively repeated until convergence, as sketched in Fig. 6.

6.1 Tonal Mapping

Tonal mapping (also known as color separation) $C \to \mathcal{M}$ is an important component in any appearance reproduction pipeline, since it amounts to a colorimetric characterization of the physical output device. Previous, dictionary-based approaches to establishing the mapping between the tonal space $\mathcal{M}$ and the input space C (e.g., [Babaei and Hersch 2016; Brunton et al. 2015]), however, require a high number of measurements, becoming impractical for higher-dimensional tonal spaces. In addition, as Brunton et al. explain, the translucent nature of our materials complicates the measurements. We therefore opt for a computational *predictive* approach instead.

We determine the mapping $C \to \mathcal{M}$ through inversion of the mapping $\mathcal{M} \to C$. Although this mapping corresponds to our forward appearance model (Sec. 4), its application would be very expensive for this purpose, for virtually the same reasons why the data-driven approaches are. We have therefore developed a tailored, analytic prediction model instead.

Albedo Mapping

We can observe that for a semi-infinite, homogeneous medium with a smooth interface, its diffuse reflectance (i.e., color C) depends on the refractive index η, scattering anisotropy g, and most importantly,

the scattering albedo α. Crucially, the diffuse reflectance is *independent* on the optical density σ_t of the medium, as this only scales the sub-surface transport uniformly in all directions.

Consequently, by fixing the values of η and g we can obtain a bijective mapping $\alpha \leftrightarrow C$. We use the measured $\eta = 1.5$, and $g = 0.4$ as justified in Sec. 5.3 (and further in the supplemental material).

Using the above canonical parametrization, we obtained the reflectance data via an analog brute-force Monte Carlo simulation. We computed 1024 data points (distributed uniformly in the α domain), and numerically fit them to an empirically chosen functional model, using simulated annealing under the L^2 norm. The resulting forward mapping $\alpha \to C$ (plotted in Fig. 7) is

$$C(\alpha) = C_s + (1 - C_s) \cdot \sum_{k=1}^{K} a_k \cdot \alpha^{b_k}, \qquad (4)$$

where C_s is the reflectance of the material's interface (i.e., the portion of energy that does not participate in the sub-surface transport) and equals $C_s = 0.04526$ under the chosen parametrization.

For $K = 5$, we obtain the average absolute error of $3.9 \cdot 10^{-4}$ and a peak absolute error of $5.7 \cdot 10^{-3}$ with the parameters

$$a_k = \{0.065773, 0.201198, 0.279264, 0.251997, 0.201767\},$$
$$b_k = \{1.569383, 6.802855, 28.61815, 142.0079, 1393.165\}.$$

For completeness we also computed the inverse mapping $C \to \alpha$ (see Apx. A.1).

Our approach is conceptually similar to [Papas et al. 2013], who however rely on tabulated data for the mapping. We found our analytic fit to be more accurate, particularly in high-α regions where the non-linearity of the space is especially high. It is worth noting that Pharr et al. [2016] also obtain an analytic expression for the
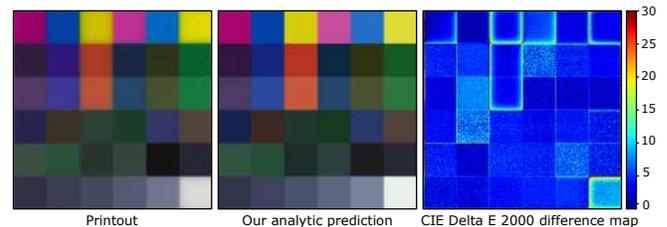
Fig. 8. Sampling of the CMYKW space with equal weights w_m for all material subsets, plus a grayscale palette. Our analytic prediction matches the printout well, especially when accounting for the different perception of the printout arising from the color bleeding between the patches.
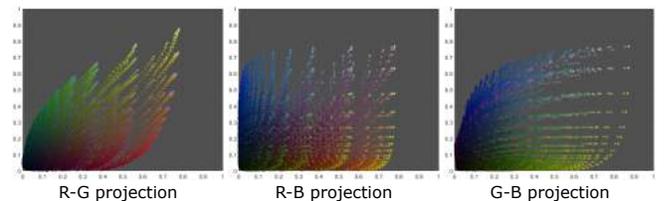
Fig. 9. Visualization of the CMYKW gamut in linear RGB (i.e., the mapping $\mathcal{M} \to C$ described in Sec. 6.1). Shown is a small set of 30k points, out of the 32M points we used to construct the mapping.

| (a) Target (full RGB) | (b) Target (gamut-mapped) | (c) Default printer output | (d) Our direct mapping | (e) Our optimized result |
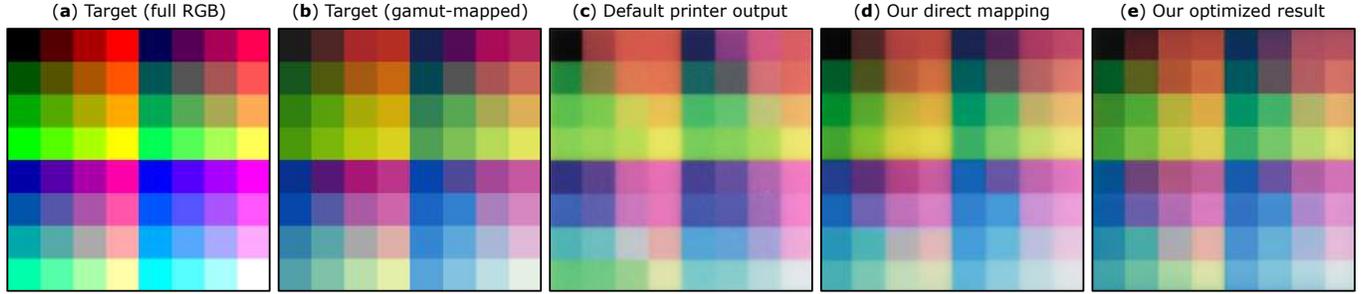
Fig. 10. Reproduction of the RGB gamut. We regularly sample the full RGB space C (a) and print it in the default printer color reproduction mode (c). We then print the gamut-mapped target (b) using both our direct reproduction method from Sec. 6 (d) and the optimized method from Sec. 7 (e). Our direct reproduction achieves a more consistent target match than the default print, further improved by the optimization in terms of both structure.
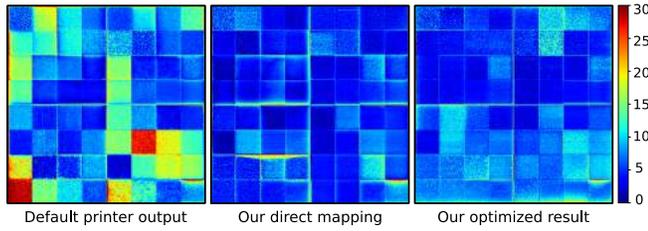
Fig. 11. Difference maps between the RGB color cube printouts (Fig. 10, c, d and e) and the gamut-mapped target (Fig. 10, b). We use the CIE Delta E 2000 color difference metric. The corresponding average color differences are 9.2, 4.1 and 5.7, respectively.

mapping, but use photon beam diffusion [Habel et al. 2013] to obtain the fitting data, which in itself is an approximate technique.

Mixture Mapping

The mapping $\alpha \to C$ (Eq. 4) enables us to define the forward mapping $\mathcal{M} \to C$ by using the measured optical parameters $\{v_m\}$ of the printing materials $m \in \mathcal{M}'$. For this we convert the optical parameters of the continuous tonal mixture $w \in \mathcal{M}$ into the *linear space of absorption and scattering coefficients* σ_a and σ_s (which fundamentally express the respective particle concentrations) as

$$\sigma_{a|w} = \sum_{m \in \mathcal{M}'} w_m \cdot (1 - \alpha_m) \cdot \sigma_{t|m} \quad (5)$$

$$\sigma_{s|w} = \sum_{m \in \mathcal{M}'} w_m \cdot \alpha_m \cdot \sigma_{t|m}, \quad (6)$$

where $\alpha_m, \sigma_{t|m} \in v_m$, i.e., the measured material albedos and extinction coefficients. The resulting mixture albedo is then

$$\alpha_w = \frac{\sigma_{s|w}}{\sigma_{a|w} + \sigma_{s|w}}. \quad (7)$$

Applying Eq. 4 yields the prediction for the mixture color, the accuracy of which is demonstrated in Fig. 8.

Tonal Mapping

The ability to efficiently predict the color for any continuous material mixture $w \in \mathcal{M}$ allows us to build the inverse mapping $C \to \mathcal{M}$. We densely sample the 5D redundant space $\mathcal{M}$ by first sampling the weights for the CMYK materials and then computing the weight

for the white material as $w_W = 1 - \sum_{m \in \text{CMYK}} w_m$, discarding all mixtures with $\sum w_m > 1$. For efficiency sake, we found it advantageous to bias the sampling of the CMYK dimensions towards lower concentrations, as the naïve uniform sampling yields a highly non-uniform distribution of the mapped points in C as a result of the strong non-linearity of the albedo space.

The mapping $C \to \mathcal{M}$ then amounts to the application of Eq. 7 and Eq. 4 to the values $\in \mathcal{M}$ obtained by the above sampling procedure, and finding the closest point in the cloud of the resulting values $\in C$. This step also implicitly performs the gamut mapping of the initial target image. The distance metric we use here is a Euclidean distance in the sRGB space (instead of the linear RGB space C), due to its better perceptual uniformity. We cache this mapping in a 3D table with 256 bins in each of the RGB dimensions, and accelerate the point cloud search by building a kD-tree over it. A sparse subsampling of the mapped point cloud is shown in Fig. 9.

The ability of our $C \to \mathcal{M}$ mapping to evenly cover and reproduce the RGB color gamut is demonstrated in Fig. 10. As we discuss further in Sec. 7 the accuracy of color versus structure reproduction is, to a certain extent, a matter of tradeoff. This is directly visible in our results in Fig. 10, but also in the corresponding numerical difference figures presented separately in Fig. 11. In any case, both our direct-mapped and optimized results achieve a better color reproduction accuracy than the default print.

6.2 Tonal Quantization

Tonal quantization essentially amounts to half-toning. We have compared the per-layer 2.5D approach of Brunton et al. [2015] with a full 3D half-toning [Lou and Stucki 1998]. We found that while the 3D half-toning approach leads to a slightly better structure preservation, the 2.5D variant leads to better colors and presents an overall better tradeoff in our specific case. We therefore use the method of Brunton et al. [2015] in all our results.

An interesting alternative would be combining our method with the recent contoning technique of Babaei et al. [2017], to suppress the quantization noise inherent to any half-toning. We see this as an interesting direction for future work.

7 STRUCTURAL OPTIMIZATION

In this section we describe the core, optimization-driven part of our pipeline (cf. Fig. 6), which builds on the direct reproduction part

detailed in Sec. 6. To recap, the purpose of this part is to iteratively refine the solution towards a better structural match with the target, by countering the distortions caused by sub-surface light transport.

7.1 Problem Definition

We seek a volumetric assignment of print materials $M \in \mathcal{M}'$ that minimizes the distance from an opaque diffuse appearance goal represented by the target image T (Sec. 6):

$$\arg\min_{M \in \mathcal{M}'} d\left[f(M \rightarrow V), T\right], \qquad (8)$$

where d is a distance metric. Because of the complex nature of the prediction operator f (Sec. 4), it is challenging to solve this problem using standard optimization techniques.

First, in spite of some similarities it is important to point out that our problem cannot be cast as a deconvolution. The principal difference is that deconvolution methods *remove* distortions that have already impacted the signal in question, while our solution has to embed the correction into the signal *a priori*. While these operations commute in the mathematical sense, they certainly do not in physical conditions where the signal is limited to a certain narrow gamut. A good demonstration of this problem has been given by [Montalto et al. 2015] in the context of visual acuity enhancement. In comparison to them, we face two additional significant issues: our 'blurring kernel' is typically very asymmetric, and has a strong material-dependent spatial variation.

The complex non-linear nature of our appearance formation process, combined with the sheer size of our search space, also makes the application of generic optimization solvers very difficult. For instance, even a small 5×5 cm printout will have the lateral resolution of 600^2, with up to 50 layers modified by the optimization, corresponding to 18M voxels. This led us to develop a more specialized approach built around the properties of light transport in our conditions. Specifically, we seek physically justified heuristics that will significantly prune our search space yet leave enough freedom to find meaningful solutions.

7.2 Solution Concept

Our solution's central idea comes from the observation that color fidelity and structural fidelity are actually competing requirements. This is demonstrated in Fig. 12, showing a colored binary pattern vertically extruded into different depths (which we will refer to as

0.5mm   1mm   2mm
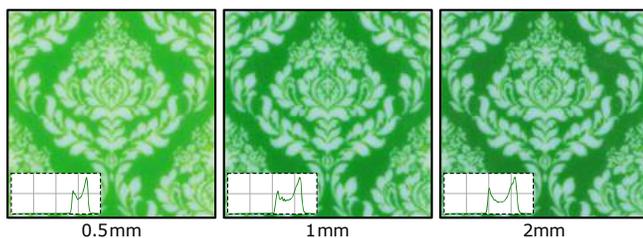
Fig. 12. The effect of varying the thickness of the colored layer (underlied by pure white material). Notice how the material is overall darkening and the undesirable color bleeding into white is increasing when more absorbing material is present. This manifests in shifting the green channel histogram towards lower values and spreading its width as shown in the insets.

characteristic depth). We can clearly see that increasing the characteristic depth trades structural acuity (which would peak in the extreme case of a single colored layer, as this would lead to the smallest possible color bleeding) for color saturation (which is directly proportional to the amount of applied colored material). The latter observation agrees with Brunton et al. [2015] who report an expansion of the color gamut with the characteristic depth increase.

The challenge lies in generalizing the above qualitative view to arbitrary target structures. It is clear that we need to consider a unique vertical material variation ('stacking') for every surface point. However, the appearance at every surface point is influenced by a large neighborhood, due to the relatively large diffusion distance of our materials. Consequently, we cannot optimize for each stacking in isolation (as was possible in previous approaches [Dong et al. 2010; Hašan et al. 2010] due to band-limited input specifications), but instead optimize *jointly* for the full 3D structure (Eq. 8) and iteratively refine it based on the global appearance prediction given by the operator f (Eq. 1).

7.3 Optimization

Based on the discussion in Sec. 7.2 the following guidelines apply while optimizing the 3D voxel structure of a printed object. First, we minimize the amount of absorbing material in each channel to reduce the undesired color bleeding and excessive material darkening (Fig. 12). We refer to absorbing materials as 'dark', since they reduce the amount of light either selectively per channel (CMY materials) or overall (K material). Second, we concentrate the dark materials as close as possible to the object surface to minimize their impact on neighboring regions' colors. This effectively means that darker materials are always placed atop of lighter ones. Third, at edges between dark and light materials, we *erode* the dark material as a function of depth to further limit the lateral energy absorption (cf. Fig. 13). As a result a tapered volume of the dark material is obtained, which helps retaining the contrast and symmetricity of the edge at the intended position.

Outline. Accounting for these guidelines, our solution starts from the top material layer and proceeds to subsequent layers as follows:

1. Assign a default material (typically plain white).
2. For a given iteration, express the current solution error as a signed energy difference between the target appearance and the prediction from the current iteration. If the error did not decrease since the last iteration, output the best solution so far.
3. Propagate the error vertically into the object, starting at the top layer. To account for the energy interchange between the neighboring stackings, diffuse the undeposited positive error by a small erosion kernel that increases with depth to achieve the tapered volume effect (Fig. 13).
4. If a negative error cannot be deposited into the current layer ('too dark' case), and vice-versa for a positive error ('too light' case), proceed to the next lower layer.
5. If all error energy has been deposited, initiate a new iteration and continue to Step 2.

We now detail the key aspects of the procedure.

ALGORITHM 1: Solution refinement (optimization core). All the operations are applied on entire matrices (images), i.e., separately for each channel and each vertical stack of voxels corresponding to a given surface pixel.

Input:

X : current working (proxy) volume $\in \mathcal{C}$

T : gamut-projected target image $\in \mathcal{C}$

P : current appearance prediction image $\in \mathcal{C}$

Result:

X : updated working volume $\in \mathcal{C}$

begin

    // Calculate positive and negative differences from target.

    $D = T - P$

    $D^+ = \max(D, 0)$

    $D^+_{\mathrm{res}} = 0$

    $D^-_{\mathrm{res}} = \min(D, 0)$

    // Loop over all proxy volume layers, starting from the topmost one, down to the last layer Z.

    for $z = 0..Z$, $L_z \in X$ **do**

        // Conservatively darken the layer proportionally to c_{p}.

        $L_z \mathrel{+}= c_{\mathrm{p}} \cdot D^-_{\mathrm{res}}$

        // Lighten the layer in the full extent.

        $L_z \mathrel{+}= D^+$

        // Add any remaining residual energy from the above layers, diffusing it laterally with the 2D kernel k_e.

        $L_z \mathrel{+}= k_e(D^+_{\mathrm{res}}, z)$

        // Calculate residuals for the next layer. D^-_{res} is the residual darkening energy from the current layer compensated for the absorbance of the above layers with c_{a}. The lower bound for D^-_{res} is chosen to enable enough energy propagation while avoiding divergence.

        $D^+_{\mathrm{res}} = \max(L_z, 1) - 1$

        $D^-_{\mathrm{res}} = \mathrm{clamp}(c_{\mathrm{a}} \cdot L_z, -2, 0)$

        // Ensure the layer stays in the gamut.

        $L_z = \mathrm{clamp}(L_z, 0, 1)$

    end

end

Fig. 13. Erosion during traversal. The benefit of progressively eroding the darker materials (**right**) compared to simple extrusion (**left**) is a better balance of the appearance 'profiles' around edges: the result overall matches the light side of the edge much better, at the cost of just a slightly worse match in the dark side.

Representation. We represent the working solution by an RGB proxy volume $X \in \mathcal{C}$ (also see Fig. 6). The advantage over representing it in the tonal space $\mathcal{M}$ (or even its discrete version $\mathcal{M}'$) is its linearity, and the fact that being expressed in the same space as both the target and the appearance prediction allows us to project the prediction gradient directly back into the solution. We store the proxy and hence compute the entire optimization in the native printer resolution.

Solution Refinement. The prediction gradient (i.e., the residual difference between the prediction and the target) is used to adjust the solution. Since we represent the solution in the linear space $\mathcal{C}$, the adjustment is directly proportional to the gradient. The main challenge is the dimensionality mismatch: the gradient is a 2D image, while the solution a 3D volume.

In the refinement procedure (Algorithm 1), we divide the residual energy into its positive (the solution needs to be lightened) and negative (the solution needs to be darkened) parts. We then traverse the solution vertically downwards, and add the positive residual aggressively to each layer, while the negative residual is applied conservatively and only up to the minimal necessary depth to minimize color bleeding (in line with the heuristics from the beginning of this section). We also slightly diffuse the positive residual during the traversal, which preserves the dark composition in homogeneous areas, but erodes it in the vicinity of structural gradients where an excess of dark material would cause the majority of color bleeding.

The described behavior is controlled by three empirical parameters (cf. Algorithm 1).

- $c_{\mathrm{p}} \in (0, 1)$: proportionality coefficient. Conservatively tempers the darkening of the layering. We have experimentally determined $c_{\mathrm{p}} = 0.5$ to yield good and stable results.

- $c_{\mathrm{a}} \in [1, \inf)$: absorbance coefficient. Approximates the lower visibility of a given layer through the stack above and compensates for it. Since by design it cannot happen that a dark voxel is placed underneath a stack of lighter voxels, we apply this compensation to the negative energy only. We obtain the value of c_{a} by averaging the absorption coefficients of all significantly absorbing print materials $\in \mathcal{M}'$ and by computing the resulting transmittance of a single layer of such virtual material.

- k_e: 2D erosion kernel. Diffuses the positive residual energy to progressively suppress dark material assignment near structural gradients. For simplicity we use a small Gaussian kernel with the standard deviation equal to half the respective layer depth z.

We use the same set of parameters to generate all our results.

Stopping Criterion. We stop the outer loop of the optimization if the value of our distance metric d (see Eq. 8) does not change significantly between two successive iterations. We use the structural similarity index (SSIM, [Wang et al. 2004]) as a metric, as its performance is usually not influenced by the Monte Carlo noise inherently present in the appearance prediction P.

7.4 Discussion

Thanks to the simplicity of our core refinement procedure, it is very efficient. Since the radius of the erosion kernel k_e is bounded by the characteristic depth, the complexity of the refinement is in fact linear w.r.t. number of voxels in the working volume X. The overall optimization process typically converges in about 20−25 iterations; we show an example of the convergence process for a complex input in Fig. 14.

While we require the specification of some empirical parameters, they themselves have an intuitive meaning and, once set up, worked well for all our results without any modifications. We hypothesize
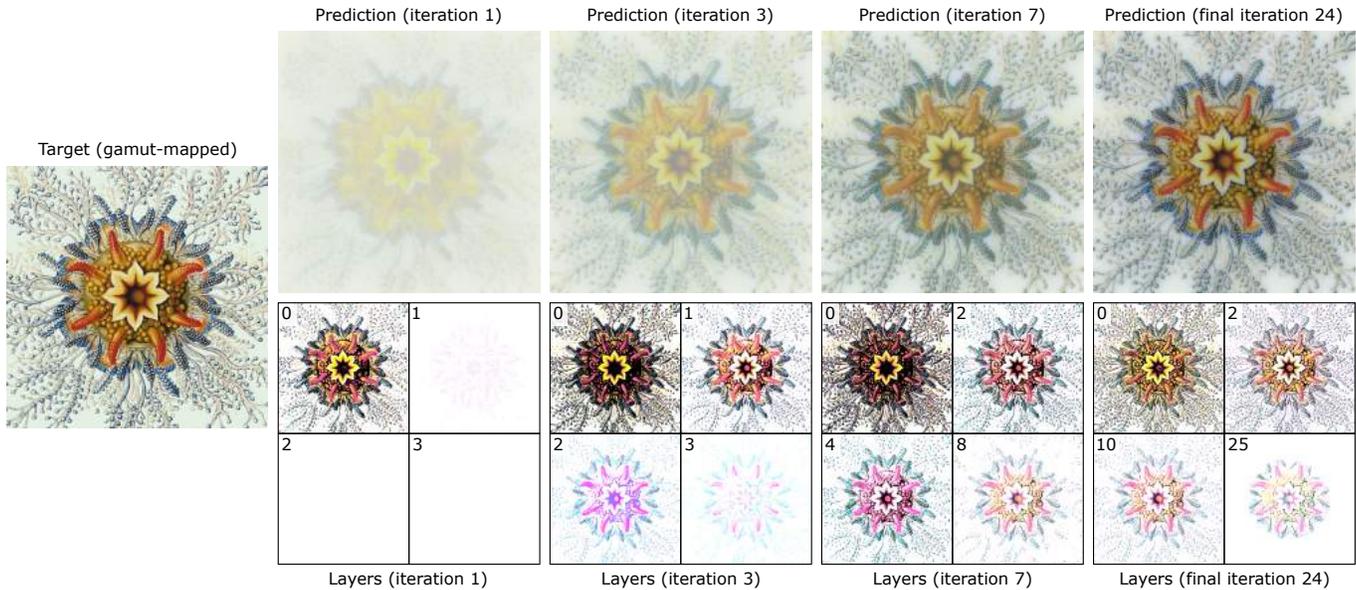
Fig. 14. Convergence of the optimization (Sec. 7) for a complex target. For a few selected iterations we visualize the distributions of the CMYKW materials in several selected layers under the surface (the z-index of each layer is shown in the upper-left corners). In spite of some residual color bleeding that could not have been compensated for, the optimization finds an overall good match to the target.

their values could be automatically derived from the optical parameters of the printing materials (as we already do for c_a), but leave this to future investigation.

8 FABRICATION

We printed all our results on the Stratasys J750 poly-jetting printer [Stratasys 2017] using a manual voxel printing mode provided by the manufacturer that allows us to address native output voxels directly. The tonal space $\mathcal{M}'$ corresponds to the CMYKW printing materials from the Vero Rigid Opaque family (specifically BlackPlus and PureWhite for KW). For the comparisons in Sec. 9, we additionally use the standard texture print mode of the machine, which can emulate up to 360k unique tonal mixtures (dubbed 'digital multimaterials' by the manufacturer). The lateral printing resolution was 300 DPI (i.e., 84.7 μm per voxel), and the vertical resolution 27 μm per layer. The total thickness of all our samples was 10 mm, which is the peak penetration distance we determined for our prints (by shining a white LED light through them, in a dark room).

While the manual voxel print gives us full control over the material composition, it produces raw printouts without a controlled surface finish. The surface is also polluted by the support material the printer automatically applies, which significantly distorts the appearance. To counteract this, we print a few extra transparent layers (using the Vero Clear material) on the sample surface, and scrub the residual support material with a grade-1000 sandpaper. We then spray each printout with a transparent lacquer to achieve a smooth and reproducible finish.

9 RESULTS

We first list several parameters and statistics of our method.

- Our entire pipeline operates on the native printer resolution, which is 300 DPI laterally and 900 DPI vertically. The resulting voxel size is therefore $85 \times 85 \times 27$ μm.
- Our Monte Carlo prediction model takes about three minutes to compute for a 500×500 image with 500 samples per pixel, on a cluster with 100 quad-core Intel Xeon E5620 CPUs. The remaining steps of our pipeline take less than a minute to compute in a hybrid Python/C++ implementation. For a volume with $500 \times 500 \times 370$ voxels ($42 \times 42 \times 10$ mm) the optimization takes an hour to compute, on average, corresponding to approximately 20 iterations.
- Photographs of the printouts were taken using the same setup (camera, lens, lights) as for the calibration (Sec. 5). The white balance was set up using X-Rite ColorChecker Passport. We ensured consistency among all the printouts' photos, but the capturing system as a whole was not fully color calibrated, that is why the actual color perception in the images might differ slightly from the targets.

Prediction Accuracy. As a first step we need to validate the accuracy of our prediction (Sec. 4) using the material data measured in Sec. 5. In Fig. 15 we show several examples that we can indeed predict the appearance of our prints. A similar match is achieved for the remaining results in the paper.

Optimization Consistency. Given that our target inputs cannot often be reproduced perfectly (i.e., there will always be some amount of color diffusion in translucent materials), we want to verify that our optimization can reach a goal that is guaranteed to be achievable. For this purpose we generate a very specific volumetric composition: randomized, colored characters with different thicknesses assigned

Prediction   Printout
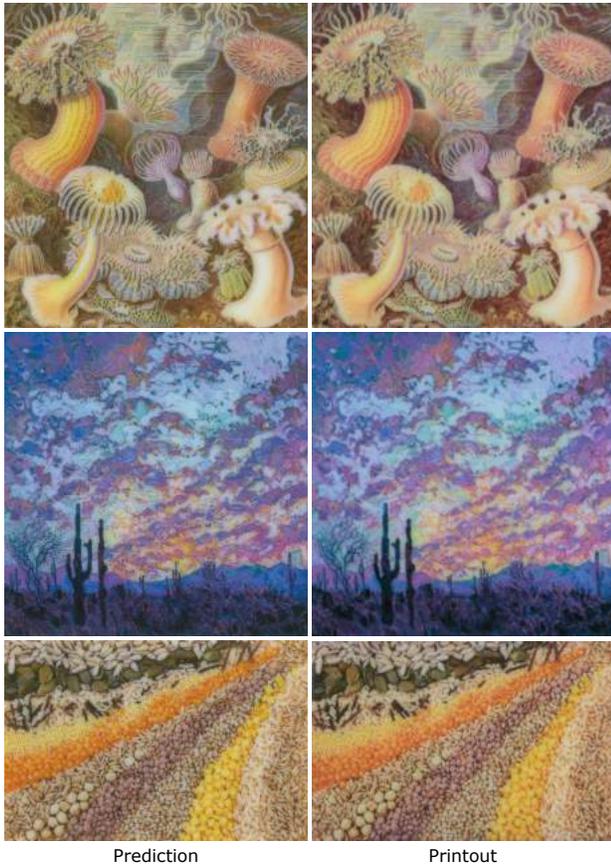
Fig. 15. Predictions of our printouts' appearance using the measured data from Table 2. The slight blurriness of the printouts' images is due to the band-limitedness of our optical capture system, plus any residual surface roughness which our surface treatment (Sec. 8) cannot eliminate entirely.

Target   Optimization match   Target   Optimization match

Fig. 16. Validation of our optimization. Our method can reproduce target appearances arising from complex, randomly assigned material distributions, even if they by themselves are not explored by the optimization.

up to 60 layers under the surface. We then render our prediction of this volume and use *this* image as a target for our optimization (which gives us the guarantee such appearance is in fact reproducible). Fig. 16 shows that we can in fact achieve such appearances, even though they arise from a complex volumetric structure that by itself is not explored by our optimization.

Reproduction Evaluation. The main results of our method are presented in Fig. 17. We aimed to reproduce various different motifs. In the order of listing: two detailed drawings (*Flower* and *Marine*), two impressionist paintings with structured strokes (*Dusk* and *Tree*), and two highly heterogeneous natural materials (*Leaves* and *Cork*). The reference method we compare against is the default texture printing mode of our machine, which functionally corresponds to the direct reproduction part of our pipeline (Sec. 6), i.e., focuses only on color reproduction. To provide a comparison for our optimized results, we use unsharp masking to enhance the targets and print them in the default texture print mode. This is compatible with the treatment of Babaei et al. [2017]; we also use an equivalent method to estimate the sharpening kernel.

We draw several conclusions here. First, our direct reproduction can achieve deeper colors compared to the default texture print. Part of the reason is that the default mode applies a relatively thin layer of the colored materials (usually below 1 mm, while we use 2.5 mm). In any case, our proposed material mapping generally reproduces the target colors at least as well as the default texture print.

Second, our optimization always leads to a better match of both structure and color than any of the compared methods. In some cases, the unsharp-masking solution yields almost comparable results (especially in lower-contrast prints where the color bleeding tends to be less noticeable). It is important to note however, that the basic property of unsharp masking (and other such enhancement methods) is to increase contrast by adding dark material around edges. This in turn causes additional undesired color bleeding, which our method avoids by design (see, e.g., the first row in Fig. 17).

Non-uniform illumination. Although our method assumes a uniform 'white-sky' illumination during the prediction, we show that the quality of the obtained results holds well outside those conditions. First, in Fig. 18 we show several samples illuminated by direct sunlight, which is strongly directionally non-uniform. Second, to visualize their local diffusion characteristics, we illuminated a subset of our samples with a laser line in Fig. 19. In both scenarios, our results optimized for uniform lighting (Fig. 17, column **f**) are superior to the unsharp-masked default prints (Fig. 17, column **e**) in terms of structure preservation, agreeing with the main results.

Monochromatic Printing. For grayscale targets, our method can sometimes produce chromatic quantization noise, especially for light tones. This is mainly because our tonal mapping (Sec. 6.1) can sometimes map grayscale tones to mixtures of CMY. We can get around this by explicitly constructing the tonal mapping for just the black and white materials, as demonstrated in Fig. 20. A more general future solution would be based on a controlled black separation, akin to methods used in 2D printing.

10 DISCUSSION

Convergence. The relative simplicity of the proposed optimization procedure makes it particularly easy to implement. Thanks to the efficient design of our heuristics and the high accuracy of the Monte Carlo forward prediction (Sec. 4), the optimization also converges in a modest number of iterations (cf. Sec. 7.4).

The prediction itself is currently the biggest bottleneck of the pipeline (cf. timing figures in Sec. 9), but this is mostly due to the general CPU implementation of path tracing we use. Given the
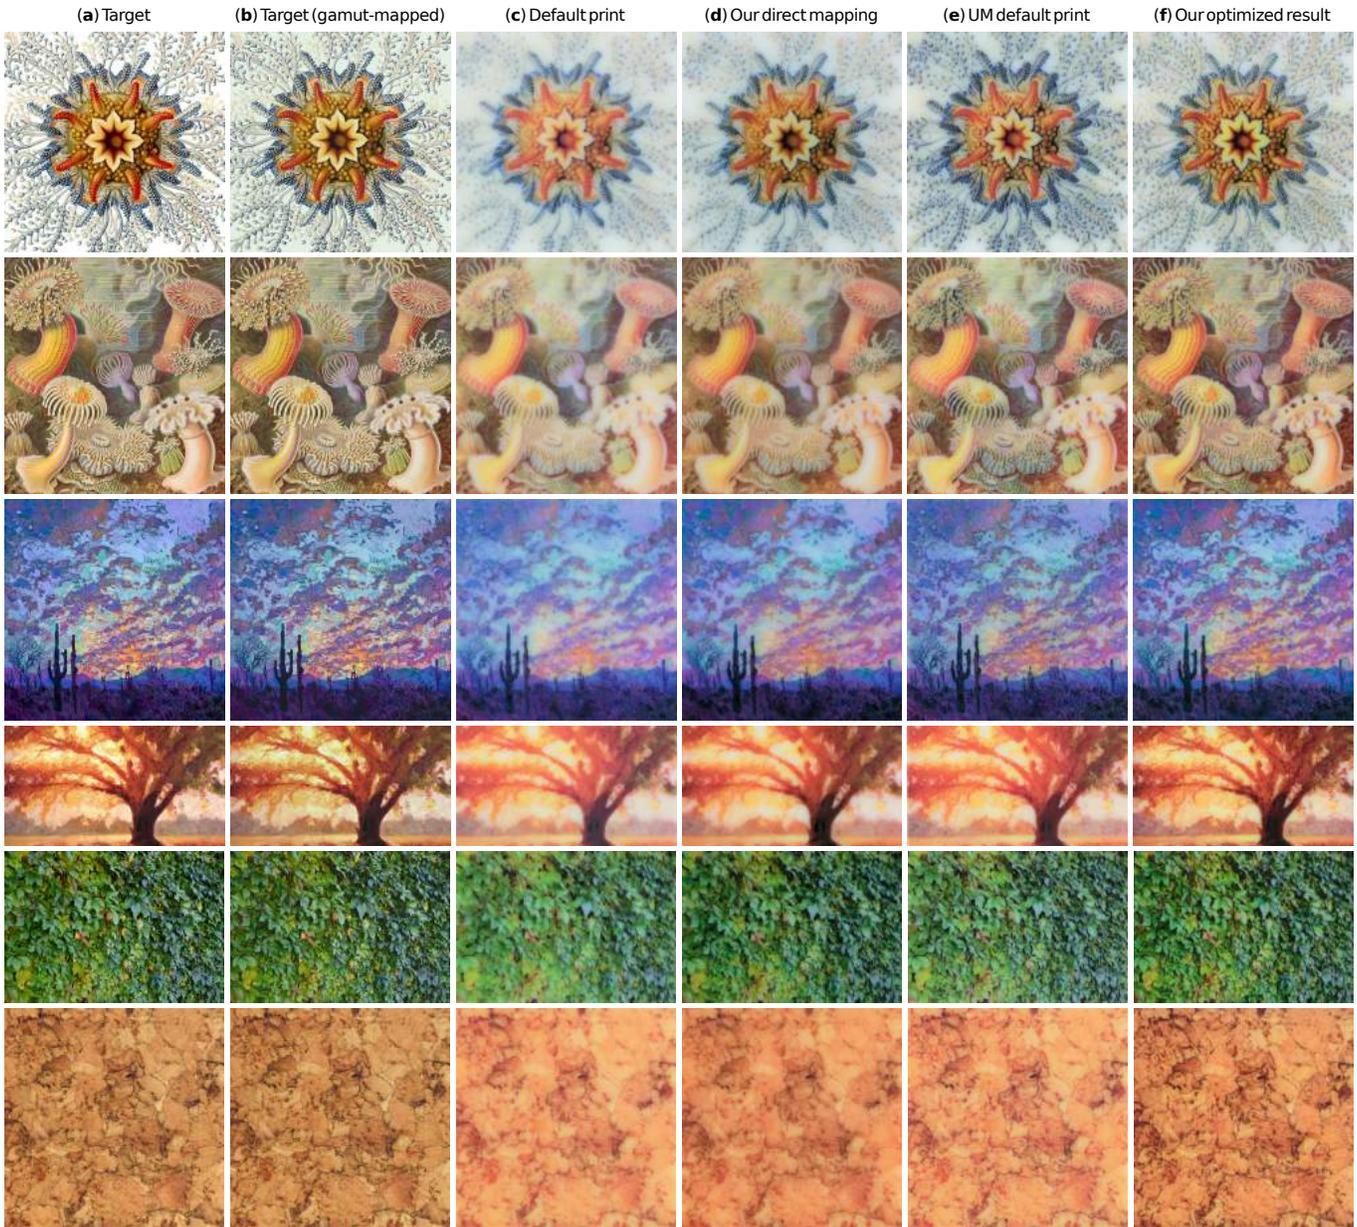
Fig. 17. Evaluation of the main results of our method. We compare the default printer output (**c**) and its sharpened version (**e**) which uses unsharp masking (UM), to the results of our method when using the direct reproduction mode (**d**) and the scattering-aware optimization (**f**). Results (**c**) and (**e**) were produced using the original target (**a**) as an input, while for the results (**d**) and (**f**) we used our gamut-mapped target (**b**).

good parallelizability of path tracing, an optimized multi-GPU implementation could bring at least an order-of-magnitude speedup (or achieve comparable timings on a single high-end GPU). Orthogonally to that, a multi-resolution approach could converge in a comparable number of iterations, but could dramatically decrease the computation time for each of them in an amortized sense. Finally, given the high optical density of the print materials (Sec. 5), the simulated path lengths can easily reach thousands of scattering

interactions. Since many of these occur in the underlying homogeneous white material, we hypothesize that applying the similarity theory [Wyman et al. 1989; Zhao et al. 2014] could also significantly decrease the simulation costs without compromising its accuracy.

General Geometry. Because our solution primarily focuses on the input textures' complexity, we operate in the canonical setting of planar slabs. On the other hand, we believe a generalization of our results to full 3D geometry is perfectly plausible, and would involve

Fig. 18. A subset of our samples illuminated by direct sunlight on a clear day, comparing the unsharp-masked default prints (**top row**) with our optimized results (**bottom row**). The light is arriving from the top-right, with the zenith angle of approximately 50°. Although our results are not optimized for such directionally non-uniform light, their structure is preserved comparatively better (cf. respective targets in Fig. 17).
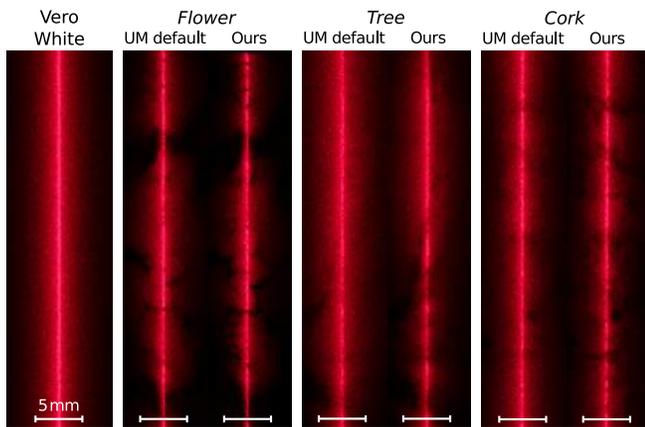
Fig. 19. A subset of our samples illuminated by a 650 nm laser line at matching locations, comparing the unsharp-masked default prints with our optimized results. For reference we also include a clear white slab in the comparison. Although our results are not optimized for this type of lighting, the lateral scattering is still inhibited to a greater degree. Please note that the graininess in the images arises from the camera's sensor noise.
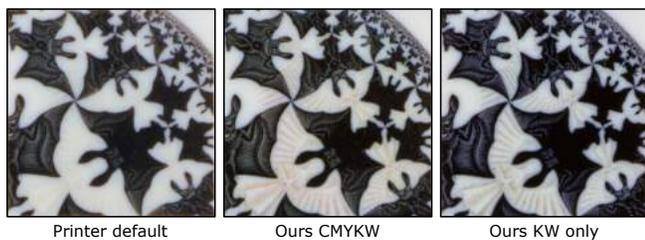
Fig. 20. Comparison of our full CMYKW result and the KW (black and white) only, to the default printer output.

two key steps: first, aligning the solution domain with the texture space of the input 3D object (including a correct handling of boundary conditions), and second, propagating the prediction gradient in an arbitrary direction in the object's voxel grid (similarly to the

approach Brunton et al. [2015] chose for the 3D generalization of layered half-toning). That however still carries the assumption that the object's geometric features are at least as large as the materials' diffusion distance; if that were not the case, an additional redesign of our optimization heuristics would be necessary.

Volumetric Blockers. The lateral transport could alternatively be prevented by placing opaque barriers near salient edges. We however do not have such an optically dense material at our disposal. More importantly, color bleeding influences all features, from edges to smooth gradients, and such barriers cannot be placed everywhere as they would negatively influence the overall light transport and cause unwanted darkening.

Perceptual Considerations. We considered implementing our optimization in a perceptually uniform colorspace (such as La*b*), but applying some of the optimization steps in such a space is not trivial. Orthogonally, a perceptually motivated global [Montalto et al. 2015] or, preferably, local contrast modification implemented via an adaptive tradeoff between structure and color (as hinted on in Sec. 7) could subjectively improve our results even further. In general, the perception of heterogeneous translucency is largely an open topic [Fleming and Bülthoff 2005; Fleming et al. 2004; Gkioulekas et al. 2013a] and deserves a dedicated treatment.

11 CONCLUSION

This work presents a self-contained, end-to-end system to calibrate, model, and suitably compensate for lateral scattering in heterogeneous translucent 3D-printed volumes.

Our method involves a practical print material calibration, a hybrid material mapping scheme based on physical Monte-Carlo simulation and analytical prediction, and a carefully designed optimization loop that reliably leads to high-fidelity texture appearance.

We evaluate our system using a commercial, five-color 3D printing process where we compare favorably against the printer's native color texturing mode. Our method preserves high-frequency features well, without having to compromise on color gamut. While all of the results are demonstrated on planar surfaces, we are confident that our method can translate to more general shapes.

As a result, our system enables, for the first time, faithful reproduction of high-frequency color textures in 3D prints, despite the inherent blur of the translucent print materials.

REFERENCES

Vahid Babaei and Roger D. Hersch. 2016. N-ink printer characterization with barycentric subdivision. *IEEE Transactions in Image Processing* 25, 7 (2016).

Vahid Babaei, Kiril Vidimče, Michael Foshey, Alexandre Kaspar, Piotr Didyk, and Wojciech Matusik. 2017. Color contoning for 3D printing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 36 (2017). Issue to appear.

Alan Brunton, Can Ates Arikan, and Philipp Urban. 2015. Pushing the limits of 3D color printing: error diffusion with translucent materials. *ACM Transactions on Graphics* 35, 1 (December 2015), 4:1–4:13.

Desai Chen, David I. W. Levin, Piotr Didyk, Pitchaya Sitthi-Amorn, and Wojciech Matusik. 2013. Spec2Fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32, 4 (2013), 135:1–135:10.

Per H. Christensen. 2015. An approximate reflectance profile for efficient subsurface scattering. In *ACM SIGGRAPH Talks*. ACM.

Paolo Cignoni, Enrico Gobbetti, Ruggero Pintus, and Roberto Scopigno. 2008. Color enhancement for rapid prototyping. In *Proc. of International Symposium on Virtual Reality, Archaeology and Cultural Heritage*. Eurographics, 9–16.

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

Craig Donner, Tim Weyrich, Eugene d'Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz. 2008. A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 27 (2008), 140:1–140:12. Issue 5. https://doi.org/10.1145/1409060.1409093

E. L. Doubrovski, Elizabeth Yinling Tsai, Daniel Dikovsky, Jo M. P. Geraedts, Hugh Herr, and Neri Oxman. 2015. Voxel-based fabrication through material property mapping: A design method for bitmap printing. *Computer-Aided Design* 60 (2015).

Roland W. Fleming and Heinrich H. Bülthoff. 2005. Low-level image cues in the perception of translucent materials. *ACM Trans. Appl. Percept.* 2, 3 (2005).

Roland W. Fleming, Henrik Wann Jensen, and Heinrich H Bülthoff. 2004. Perceiving translucent materials. In *Proc. of ACM Symposium on Applied Perception in Graphics and Visualization*.

Ioannis Gkioulekas, Bei Xiao, Shuang Zhao, Edward Adelson, Todd Zickler, and Kavita Bala. 2013a. Understanding the role of phase function in translucent appearance. *ACM Transactions on Graphics* 32, 5 (2013), 147:1–147:19.

Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013b. Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics* 32, 6 (2013), 162:1–162:13. https://doi.org/10.1145/2508363.2508377

Ralf Habel, Per H. Christensen, and Wojciech Jarosz. 2013. Photon beam diffusion: A hybrid Monte Carlo method for subsurface scattering. In *Proc. of EGSR*. 27–37.

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.

Miloš Hašan and Ravi Ramamoorthi. 2013. Interactive albedo editing in path-traced volumetric materials. *ACM Transactions on Graphics* 32, 2 (2013), 11:1–11:11.

Wenzel Jakob. 2010. Mitsuba renderer. (2010). http://www.mitsuba-renderer.org.

Henrik Wann Jensen and Juan Buhler. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 21, 3 (2002).

Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. In *Proc. SIGGRAPH*.

J. Konrad, B. Lacotte, and E. Dubois. 2000. Cancellation of image crosstalk in time-sequential displays of stereoscopic video. *IEEE Trans. Image Processing* 9, 5 (2000).

Yanxiang Lan, Yue Dong, Fabio Pellacini, and Xin Tong. 2013. Bi-scale appearance fabrication. *ACM Transactions on Graphics* 32, 4 (2013).

Qun Lou and Peter Stucki. 1998. Fundamentals of 3D halftoning. In *Electronic Publishing, Artistic Imaging, and Digital Typography*, RogerD. Hersch, Jacques André, and Heather Brown (Eds.). Vol. 1375. Springer Berlin Heidelberg.

Wojciech Matusik, Boris Ajdin, Jinwei Gu, Jason Lawrence, Hendrik P.A. Lensch, Fabio Pellacini, and Szymon Rusinkiewicz. 2009. Printing spatially-varying reflectance. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 28, 5 (2009), 40:1–40:7.

Carlos Montalto, Ignacio Garcia-Dorado, Daniel Aliaga, Manuel M. Oliveira, and Feng Meng. 2015. A total variation approach for customizing imagery to improve visual acuity. *ACM Transactions on Graphics* 34, 3, Article 28 (May 2015).

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

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically based rendering: From theory to implementation* (3rd ed.). Morgan Kaufmann.

R. Pintus, E. Gobbetti, P Cignoni, and R. Scopigno. 2010. Shape enhancement for rapid prototyping. *The Visual Computer* 26, 6–8 (2010).

Tim Reiner, Nathan Carr, Radomír Měch, Ondřej Šťáva, Carsten Dachsbacher, and Gavin Miller. 2014. Dual-color mixing for fused deposition modeling printers. *Computer Graphics Forum (Proc. of Eurographics)* 33, 2 (2014).

O. Rouiller, B. Bickel, J. Kautz, W. Matusik, and M. Alexa. 2013. 3D-printing spatially varying BRDFs. *IEEE Computer Graphics and Applications* 33, 6 (2013).

Thorsten-Walther Schmidt, Fabio Pellacini, Derek Nowrouzezahrai, Wojciech Jarosz, and Carsten Dachsbacher. 2014. State of the art in artistic editing of appearance, lighting, and material. In *Eurographics 2014 - State of the Art Reports*.

Pitchaya Sitthi-Amorn, Javier E. Ramos, Yuwang Wangy, Joyce Kwan, Justin Lan, Wenshou Wang, and Wojciech Matusik. 2015. MultiFab: A Machine Vision Assisted Platform for Multi-material 3D Printing. *ACM Transactions on Graphics* 34, 4 (July 2015), 129:1–129:11. https://doi.org/10.1145/2766962

Hiroki Sone, Toshiya Hachisuka, and Takafumi Koike. 2017. Parameter estimation of BSSRDF for heterogeneous materials. In *Eurographics Short Papers*.

Ying Song, Xin Tong, Fabio Pellacini, and Pieter Peers. 2009. SubEdit: A Representation for Editing Measured Heterogeneous Subsurface Scattering. *ACM Transactions on Graphics* 28, 3 (July 2009), 31:1–31:10.

Y. Song and W. Wang. 2013. A data-driven model for anisotropic heterogeneous subsurface scattering. In *Proc. of Signal and Information Processing Association Annual Summit and Conference*.

Eric J. Stollnitz, Victor Ostromoukhov, and David H. Salesin. 1998. Reproducing Color Images Using Custom Inks. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. ACM, New York, NY, USA.

Stratasys. 2017. J750 Printer. (2017). http://www.stratasys.com/3d-printers/production-series/stratasys-j750.

Jeroen van Baar, Steven Poulakos, Wojciech Jarosz, Derek Nowrouzezahrai, Rasmus Tamstorf, and Markus Gross. 2011. Perceptually-based compensation of light pollution in display systems. In *Proc. of ACM Symposium on Applied Perception in Graphics and Visualization*. https://doi.org/10.1145/2077451.2077460

Kiril Vidimče, Szu-Po Wang, Jonathan Ragan-Kelley, and Wojciech Matusik. 2013. OpenFab: a programmable pipeline for multi-material fabrication. *ACM Transactions on Graphics* 32, 4 (July 2013). https://doi.org/10.1145/2461912.2461993

Rui Wang, Ewen Cheslack-Postava, Rui Wang, David Luebke, Qianyong Chen, Wei Hua, Qunsheng Peng, and Hujun Bao. 2008. Real-time editing and relighting of homogeneous translucent materials. *The Visual Computer* 24 (2008).

Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–12.

Tim Weyrich, Pieter Peers, Wojciech Matusik, and Szymon Rusinkiewicz. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28 (2009), 32:1–32:6. https://doi.org/10.1145/1531326.1531338

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).

Kun Xu, Yue Gao, Yong Li, Tao Ju, and Shi-Min Hu. 2007. Real-time homogenous translucent material editing. *Computer Graphics Forum* 26, 3 (2007).

Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order Similarity Relations in Radiative Transfer. *ACM Transactions on Graphics* 33, 4 (July 2014).

A APPENDIX

A.1 Inverse Albedo Mapping

In Sec. 6.1 we present an analytical mapping from the space of scattering albedos to colors, i.e., $\alpha \rightarrow C$. For completeness we also computed the inverse mapping $C \rightarrow \alpha$:

$$\alpha(C) = \sum_{k=1}^{K} c_k \left(1 - \left(1 - \frac{C - C_s}{1 - C_s}\right)^{d_k}\right). \quad (9)$$

For $K = 5$, we obtain the average absolute error of $3.5 \cdot 10^{-4}$ and the peak absolute error of $1.6 \cdot 10^{-3}$ with the parameters

$$c_k = \{0.163581, 0.391943, 0.029277, 0.316847, 0.098352\},$$
$$d_k = \{3.969542, 15.94272, 46.26871, 59.95706, 206.5716\}.$$

Note that Eq. 9 is not a simple inversion of Eq. 4; rather, a separate fit is necessary to obtain it. This inverse mapping is plotted alongside the forward one in Fig. 7.